The breakdown point of signal subspace estimation

Raj Rao Nadakuditi
Dept. of Electrical Engineering & Computer Science
University of Michigan, Ann Arbor, Michigan 48104
Email: rajnrao@umich.edu

Florent Benaych-Georges
Laboratoire de Probabilites et Modeles Aleatoires
Universite Paris 6, place Jussieu, 75252 Paris Cedex 05
Email: florent.benaych@upmc.fr

***Abstract*—The breakdown point of signal subspace methods, which is the SNR below which the algorithm's performance deteriorates dramatically, is intimately related to the breakdown point of PCA based signal subspace estimation. We shed new light on this breakdown point for a broad class of signal-plus-noise models, provide a transparent derivation that highlights its origin and verify the accuracy of the high-dimensional predictions with numerical simulations for a moderately sized system.**

I. Introduction

The breakdown in performance of signal subspace methods due to low SNR has been well-documented in the signal processing literature (see e.g. [1], [2], [3]). Advances in random matrix theory [4] have deepened the understanding of the connection between the deterioration in performance of subspace based algorithms and the breakdown in the principal component signal subspace estimates at low SNR in high dimensional settings. This paper extends this line of inquiry by providing 1) a succinct, closed-form expression via the T-transform of the noise-eigenspectrum for the breakdown point for a broader class of models that goes well beyond the Wishart models considered in [5], [6], [7] (see Section III) and 2) presents the simplest derivation yet that brings into sharp focus the mechanism that induces this phenomenon (see Section IV). Applications and extensions are discussed in Section V.

II. A general sample covariance matrix model

Consider the standard model for the snapshot vectors:

$$y_i = As_i + z_i \quad \text{for } i = 1, \ldots, m. \tag{1}$$

In the setting where the $n \times r$ matrix A has full column rank and we the model the $r \ll n$ dimensional signal vector s_i and the n dimensional noise-only vector z_i as mutually independent and i.i.d. Gaussian with mean zero and covariance, Ψ and I_n , respectively, the snapshot vector $y_i \sim \mathcal{N}(0, P_n + I_n)$ where $P_n = A\Psi A$ is the rank-r signal covariance matrix.

The additivity of the Gaussian signal and noise implies that the signal-plus-noise sample covariance matrix $S_n = 1/m \sum_i y_i y_i'$ can be decomposed as $S_n = (P_n + I_n)^{1/2} X_n (P_n + I_n)^{1/2}$ where $X_n = 1/m \sum_{i=1}^m x_i x_i'$ with i.i.d. $x_i \sim \mathcal{N}(0, I_n)$ can be thought of as the noise-only random sample covariance matrix. Since x_i, thus defined, is a spherically symmetric Gaussian random vector, it is invariant under distribution to unitary (or orthogonal) transformations and so will[1] the matrix X_n.

[1]In other words, the joint probability distribution $f(X_n) = f(Q' X_n Q)$ for arbitrary orthogonal or unitary matrix Q.

A. The model & assumptions

The preceding argument shows that a Wishart distributed sample covariance matrix (SCM) $S_n = 1/m \sum_i y_i y_i'$ formed assuming the generative model for the snapshots y_i as in (1) can be represented as

$$S_n = (P_n + I_n)^{1/2} X_n (P_n + I_n)^{1/2} \tag{2}$$

where the signal covariance matrix P_n has low rank r.

This motivates our general model for the signal-plus-noise SCM in (2) with the assumptions:

- X_n is some unitarily invariant noise-only random matrix
- $P_n = \sum_{i=1}^r \sigma_i u_i u_i'$ is the "signal" covariance matrix with "signal" eigenvalues $\{\sigma_i\}$ and eigenvectors $\{u_i\}$

The unitarily invariant noise assumption is common in the literature (see [8]) and constitutes a natural class of models wherever PCA based signal subspace estimation is justified. This setup has the advantage that it allows us to treat the SCM in (2) as a noisy, matrix-valued signal in and of itself and examine the statistical properties of its eigen-structure independent of any explicit or implicit generative model such as (1) for the snapshots used to form the SCM.

Additionally we make two assumptions on the statistics of the eigen-spectrum of X_n: 1) that the empirical distribution μ_{X_n} of the eigenvalues $\{\lambda_j(X_n)\}$ of X_n, defined as the probability measure:

$$\mu_{X_n} = \frac{1}{n} \sum_{j=1}^n \delta_{\lambda_j(X_n)},$$

converges almost surely weakly, as $n \to \infty$ to a non-random, compactly supported probability measure μ_X and 2) that the largest (resp. smallest) eigenvalue of X_n converges almost surely to b (resp. a). These are also the supremum (resp. infimum) of the support of μ_X.

Recall that the signal subspace estimate formed from the SCM S_n is the span of the r eigenvectors $\{\widetilde{u}_i\}$ corresponding to the r (assuming r is known) largest eigenvalues $\{\lambda_i(S_n)\}$ of S_n.

We now provide a precise and succinct characterization, summarized in Figure 1, of the biases and breakdown point of subspace estimation - this is the point at which the sample "signal" eigenvalues and eigenvectors are "noise-like". In stating the results, it will initially prove to be more convenient to analyze the matrix $\widetilde{S}_n = X_n(I_n + P_n)$ instead which is closely related to S_n.

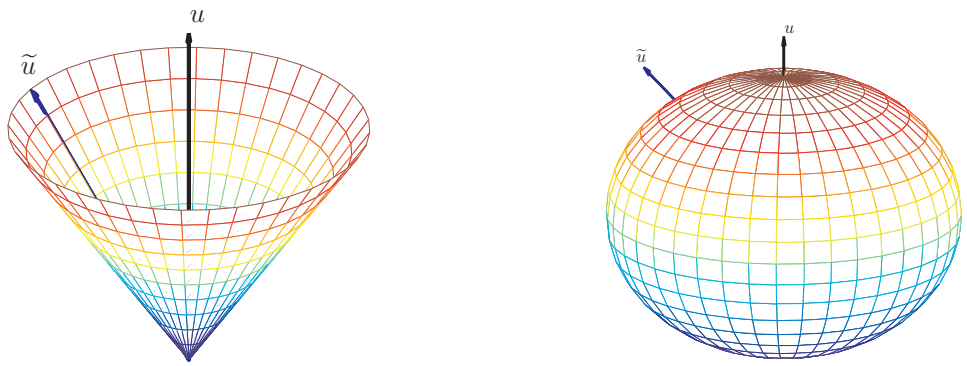

(a) When $\sigma > \sigma_{\text{BREAK.}} := 1/T_{\mu_X}(b^+)$ (b) When $\sigma \leq \sigma_{\text{BREAK.}} := 1/T_{\mu_X}(b^+)$

Fig. 1. (left panel) The signal subspace/eigenvector estimate $\widetilde{u}$ is biased relative to the true eigenvector u with the magnitude of the bias given by Theorem 3.2 that is entirely characterized by the T-transform of noise eigen-spectrum μ_X and its largest eigenvalue b . (right panel) When the signal eigenvalue is below the breakdown point $\sigma_{\text{BREAK.}}$, the signal eigenvectors are "noise-like".

III. Characterizing the breakdown point

Theorem 3.1 (Eigenvalue bias and breakdown): The "signal" eigenvalues of $\widetilde{S}_n$ exhibit the following behavior as $n \longrightarrow \infty$. We have that for each $i = 1, \ldots, r$

$$\lambda_i(\widetilde{S}_n) \xrightarrow{\text{a.s.}} \begin{cases} T^{-1}_{\mu_X}(1/\sigma_i) & \text{if } \sigma_i > 1/T_{\mu_X}(b^+) = \sigma_{\text{BREAK.}}, \\ b & \text{otherwise,} \end{cases}$$

where the T-transform of μ_X is given by

$$T_{\mu_X}(z) = \int \frac{t}{z-t} \mathrm{d}\mu_X(t) \qquad \text{for } z \notin \operatorname{supp} \mu_X,$$

and $\xrightarrow{\text{a.s.}}$ denotes almost sure convergence.

Theorem 3.2 (Bias in sample eigenvectors): Consider indices $i_0 \in \{1, \ldots, r\}$ such that $1/\sigma_{i_0} \in (T_{\mu_X}(a^-), T_{\mu_X}(b^+))$. For each n, let $\widetilde{v}_i$ be a unit norm eigenvector of $\widetilde{S}_n$ associated with the eigenvalue $\lambda_i(\widetilde{S}_n)$. Then we have
a)

$$|\langle \widetilde{v}_i, \ker(\sigma_{i_0} I_n - P_n)\rangle|^2 \xrightarrow{\text{a.s.}} \frac{-1}{\sigma^2_{i_0} \rho T'_{\mu_X}(\rho) + \sigma_{i_0}},$$

where $\rho = T^{-1}_{\mu_X}(1/\sigma_{i_0})$ is the limit of $\lambda_i(\widetilde{S}_n)$;
b)

$$|\langle \widetilde{v}_i, \oplus_{j \neq i_0} \ker(\sigma_j I_n - P_n)\rangle| \xrightarrow{\text{a.s.}} 0,$$

as $n \longrightarrow \infty$, where ker is the kernel or nullspace.

Theorem 3.3 (Breakdown of sample eigenvectors): When $r = 1$, let the sole non-zero eigenvalue of P_n be denoted by σ. Suppose that

$$\frac{1}{\sigma} \notin (T_{\mu_X}(a^-), T_{\mu_X}(b^+)),$$

and

$$T'_{\mu_X}(b^+) = -\infty$$

Then, we have

$$|\langle \widetilde{v}_i, \operatorname{range}(P_n)\rangle| \xrightarrow{\text{a.s.}} 0$$

as $n \longrightarrow \infty$.

Under additional natural conditions (omitted here for brevity - see [9]), Theorem 3.3 holds for general $r \geq 1$.

A. Extension to Sample Covariance Matrix Form

Recall that we were interested in the sample covariance matrix $S_n = (I_n + P_n)^{1/2} X_n (I_n + P_n)^{1/2}$ which is related to $\widetilde{S}_n = X_n(I_n + P_n)$ by a similarity transformation so that they share the same eigenvalues and consequently the same breakdown point as in Theorem 3.1.

Additionally, if $\widetilde{u}_i$ is an eigenvector of S_n associated with $\lambda_i(S_n)$ and when $1/\sigma_{i_0} < T_{\mu_X}(b^+)$, we have:

$$|\langle \widetilde{u}_i, \ker(\sigma_{i_0} I_n - P_n)\rangle|^2 \xrightarrow{\text{a.s.}} -\frac{\sigma_{i_0} + 1}{\sigma_{i_0} T'_{\mu_X}(\rho)}$$

and

$$|\langle \widetilde{u}_i, \oplus_{j \neq i_0} \ker(\sigma_j I_n - P_n)\rangle| \xrightarrow{\text{a.s.}} 0.$$

The results hold for general $r \geq 1$ under additional natural conditions that have been omitted here for brevity [9].

B. Connection with Free Probability Theory

The T-transform of the noise eigen-spectrum plays in important role in Theorems 3.1-3.3 in determining the breakdown point and the biases of signal subspace estimation. The S-transform defined as

$$S_\mu(z) := \frac{1+z}{z} \frac{1}{T^{-1}_\mu(z)}, \tag{3}$$

is closely related and is the analogue of the Mellin transform[2] in free probability theory [10] in the sense described next.

The operation free multiplicative convolution $\boxtimes$, is defined as the binary operation on the set of probability measures on $[0, +\infty)$. Given two independent, unitarily invariant random matrices A_n, B_n, with empirical eigenvalue distributions $\mu_{A_n} \xrightarrow{\text{a.s.}} \mu_A$ and $\mu_{B_n} \xrightarrow{\text{a.s.}} \mu_B$, the empirical eigenvalue distribution of their product $\mu_{A_n B_n} \xrightarrow{\text{a.s.}} \mu_A \boxtimes \mu_B$ will be characterized by the fact that $S_{\mu_A \boxtimes \mu_B} = S_{\mu_A} S_{\mu_B}$.

This connection between free multiplicative convolution and T^{-1}_μ, via the S-transform, is what makes its manifestation in Theorem 3.1 natural since we were investigating a signal-plus-noise random matrix of the form $\widetilde{S}_n = X_n(I_n + P_n)$.

[2]Recall that the Mellin transform of the product of two independent, scalar valued random variables is the product of the respective Mellin transforms.

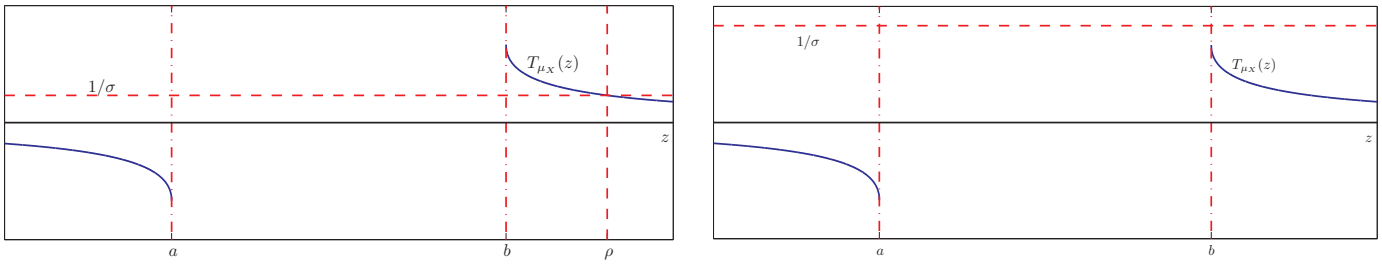

(a) When $\sigma > \sigma_{\text{BREAK.}} := 1/T_{\mu_X}(b^+)$.

(b) When $\sigma \leq \sigma_{\text{BREAK.}} := 1/T_{\mu_X}(b^+)$.

Fig. 2. Relating the biases and breakdown in signal subspace estimates to the signal eigenvalue σ and the T-transform of the noise-eigenspectrum μ_X.

IV. A Simple Derivation of the Breakdown Point

Consider the setting where $r = 1$, so that $P = \sigma\, uu^H$, with u being a fixed, unit norm column vector. Since X is assumed to be invariant, in law, under unitary conjugation, the eigendecomposition of $X = Q\,\mathrm{diag}(\lambda_1, \ldots, \lambda_n)Q^H = Q\Lambda Q^H$ will produce isotropically random eigenvectors Q.

A. Largest Eigenvalue Phase Transition

For any $z \in \mathbb{C}$ such that z is not an eigenvalue of X, we have

$$z - X(I+P) = (z - X) \times (I - (z-X)^{-1}XP),$$

so that z is an eigenvalue of $\widetilde{S} = X(I+P)$ if and only if 1 is an eigenvalue of $(z-X)^{-1}XP$. But $(z-X)^{-1}XP = Q(z-\Lambda)^{-1}\Lambda Q^H \sigma\, uu^H$ has rank one, so its only non-zero eigenvalue will equal its trace, which in turn is equal to $\sigma T_{\mu_n}(z)$, where μ_n is a "weighted" spectral measure of X, defined by

$$\mu_n = \sum_{k=1}^{n} |w_k|^2 \delta_{\lambda_k}, \tag{4}$$

where the w_k's are the coordinates of $w = Q^H u$. Thus any z out of the spectrum of X is an eigenvalue of $\widetilde{S} = X(I+P)$ if and only if

$$\sum_{k=1}^{n} \frac{|w_k|^2}{z - \lambda_k} =: T_{\mu_n}(z) = \frac{1}{\sigma}. \tag{5}$$

Equation (5) describes the relationship between the eigenvalues of $X(I+P)$ and the eigenvalues of X and the dependence on the coordinates of the vector $w = Q^H u$ (via the measure μ_n).

This is where randomization simplifies analysis. Since Q is isotropically random, the vector $w = Q^H u$ is a random vector with uniform distribution on the unit n-sphere. Hence, we have that for large n, $|w_k|^2 \approx \frac{1}{n}$ with high probability. Consequently, we have $\mu_n \approx \mu_X$ so that $T_{\mu_n}(z) \approx T_{\mu_X}(z)$. Inverting equation (5) after substituting these approximations yields the location of the largest "signal" eigenvalue to be $T^{-1}_{\mu_X}(1/\sigma)$ as in Theorem 3.1 and Figure 2-(a).

The breakdown point exists because under our assumption that the limiting probability measure μ_X is compactly supported on $[a, b]$, the T-transform T_{μ_X} is defined *outside* $[a, b]$ and unlike what happens for T_{μ_n}, we do not always have $T_{\mu_X}(b^+) = +\infty$. Consequently, when $1/\sigma < T_{\mu_X}(b^+)$, we have that $\lambda_1(\widetilde{S}) \approx T^{-1}_{\mu_X}(1/\sigma)$ as before. However, when $1/\sigma \geq T_{\mu_X}(b^+)$ then the phase transition manifests and $\lambda_1(\widetilde{S}) \approx \lambda_1(X) = b$, as in Figure 2-(b).

B. Eigenvectors Phase Transition

Let $\widetilde{v}$ be a unit eigenvector of $\widetilde{S} = X(I+P)$ associated with the eigenvalue z that satisfies (5). From the relationship $(X + XP)\widetilde{v} = zv$, we deduce that, for $P = \sigma\, uu^H$,

$$(zI - X)\widetilde{v} = XP\widetilde{v} = \sigma X uu^H v = (\sigma u^H \widetilde{v})Xu,$$

where the last equality follows because $u^H x$ is a scalar thereby implying that $\widetilde{v}$ is proportional to $(zI-X)^{-1}Xu$. Since $\widetilde{v}$ has norm one, we have

$$\widetilde{v} = \frac{(zI - X)^{-1}Xu}{\sqrt{u^H X(zI-X)^{-2}Xu}} \tag{6}$$

and

$$|\langle \widetilde{v}, u\rangle|^2 = \frac{(u^H(zI-X)^{-1}Xu)^2}{u^H(zI-X)^{-2}u} \tag{7}$$

$$= \frac{w^H(zI-\Lambda)^{-1}\Lambda w}{w^H\Lambda(zI-\Lambda)^{-2}\Lambda w} \tag{8}$$

Equation (8) describes the relationship between the eigenvectors of $\widetilde{S} = X(I+P)$ and the eigenvalues of X and the dependence on the coordinates of the vector $w = Q^H u$

Here too, randomization simplifies analysis since for large n, we have $\mu_n \approx \mu_X$ and so that when $1/\sigma < T_{\mu_X}(b^+)$, which implies that $\rho = T^{-1}_{\mu_X}(1/\sigma) > b$, we have

$$|\langle v, u\rangle|^2 \xrightarrow{\text{a.s.}} \frac{-1}{\sigma^2 \rho T'_{\mu_X}(\rho) + \sigma},$$

whereas when $1/\sigma \geq T_{\mu_X}(b^+)$ and T_{μ_X} has infinite derivative at b, we have

$$\langle v, u\rangle^2 \xrightarrow{\text{a.s.}} 0.$$

An extension of these arguments for $r > 1$ yields the general result. We rely on concentration inequalities to make the arguments rigorous [9].

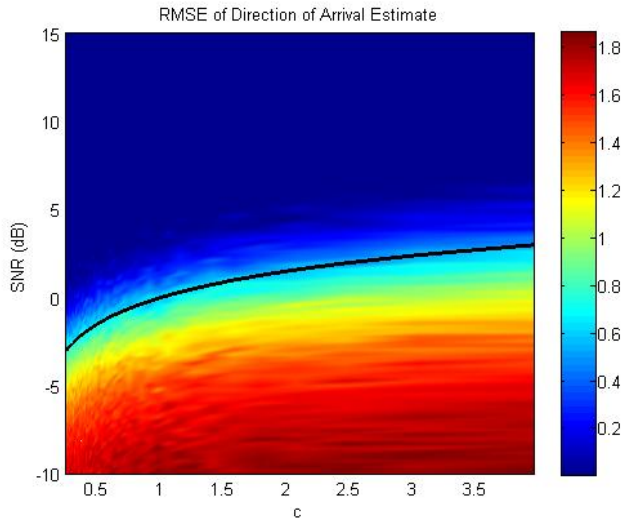

Fig. 3. A heat map of the empirical root mean squared error (RMSE) of the signal subspace MUSIC direction-of-arrival (DOA) algorithm [14] in degrees versus c = number of sensors to number of signal-plus-noise snapshots phase space illustrating the performance breakdown and the accuracy of the theoretical predicted solid black line. In this example, we chose a uniform line array with half-wavelength spacing, number of sensors $n = 250$, with $N = n/c$ i.i.d. normally distributed snapshots and signal-plus-noise covariance $I + P = \text{diag}(\sigma_1 = 1 + \text{SNR}, 1, \ldots, 1)$. We placed one source at broadside and searched over a ± 3.5 degree window. The RMSE was averaged over 1000 Monte-Carlo trials and a grid of 100 equally spaced points in the -10 dB to 15 dB (eigen) SNR range and 100 equally spaced points in the c space. The values of the colormap at each of the 100×100 faces were interpolated across each line segment and face to obtain the above plot.

V. Applications and Extensions

In Section III, we provided a succinct, closed-form expression for the breakdown point in high-dimensions for a general class of signal-plus-noise models. The signal-plus-noise models were characterized by the fact that the noise-only component had isotropically random eigenvectors. This is a very natural assumption in settings where the use of signal subspace methods (or PCA) is appropriate.

This assumption allows us to treat the sample covariance matrix as a noisy, matrix-valued signal in and of itself and examine the statistical properties of its eigen-structure independent of any explicit or implicit generative model for the snapshots used to form the sample covariance matrix. Consequently, our results are very general and go well beyond the Wishart matrix-based generative models previously considered in the literature [5], [6], [7], [11], [12], [13].

Moreover, since the breakdown point (see Theorem 3.1) is entirely characterized by the T-transform of the noise eigen-distribution, kernel based noise eigen-distribution estimation methods can be used to predict the subspace breakdown points from data. In [9] we extend this to other signal-plus-noise models. Figure 3 illustrates the agreement between the theoretical predictions in high-dimensions (expression omitted for brevity) and numerical simulations for moderate dimensions for the MUSIC DOA algorithm in the Wishart setting that has also been considered in [3].

VI. Conclusion

In this paper, we used random matrix theory [9] to predict the breakdown point for PCA-based signal subspace estimation methods for a broad, natural class of signal-plus-noise models. The utility of the asymptotic predictions for moderate system performance breakdown analysis was demonstrated in Figure 3 for the MUSIC DOA algorithm. The generality of the results presented provides a principled common basis for evaluating the performance of various signal subspace-based algorithms and facilitates the comparison of subspace-based versus alternative algorithms in the weak signal regime.

In that regard, this work represents the signal subspace breakdown counterpart to the sparse approximation breakdown work of Donoho and Tanner [15]. In contrast to [15], an important feature of the signal subspace breakdown point is its non-universality, *i.e.*, the explicit dependence on the noise-only eigen-spectrum via its T-transform (see Theorem 3.1). The emergence of the T-transform as a predictive statistic emerges from the intimate connection with (non-commutative) free probability theory as discussed in Section III-B.

References

[1] J. Thomas, L. Scharf, and D. Tufts, "The probability of a subspace swap in the SVD," *IEEE Trans. on Signal Proc.*, vol. 43, no. 3, pp. 730–736, 1995.

[2] M. Hawkes, A. Nehorai, and P. Stoica, "Performance breakdown of subspace-based methods: prediction and cure," in *Proc. of ICASSP*, 2001, pp. 4005–4008.

[3] B. Johnson, Y. Abramovich, and X. Mestre, "MUSIC, G-MUSIC, and maximum-likelihood performance breakdown," *IEEE Trans. on Signal Proc.*, vol. 56, no. 8 Part 2, pp. 3944–3958, 2008.

[4] I. Johnstone, "High dimensional statistical inference and random matrices," in *Intl. Cong. of Math., Vol. I.*, 2007, pp. 307–333.

[5] D. C. Hoyle and M. Rattray, "Statistical mechanics of learning multiple orthogonal signals: asymptotic theory and fluctuation effects," *Phys. Rev. E (3)*, vol. 75, no. 1, pp. 016 101, 13, 2007.

[6] D. Paul, "Asymptotics of sample eigenstructure for a large dimensional spiked covariance model," *Statist. Sinica*, vol. 17, no. 4, pp. 1617–1642, 2007.

[7] B. Nadler, "Finite sample approximation results for principal component analysis: a matrix perturbation approach," *Ann. Statist.*, vol. 36, no. 6, pp. 2791–2817, 2008.

[8] Y. Chen, A. Wiesel, Y. Eldar, and A. Hero III, "Shrinkage Algorithms for MMSE Covariance Estimation," 2009, available at arxiv.org/pdf/0907.4698.

[9] F. Benaych-Georges and R. R. Nadakuditi, "Eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices," 2009.

[10] D. V. Voiculescu, K. J. Dykema, and A. Nica, *Free random variables*, ser. CRM Monograph Series. Providence, RI: Amer. Math. Soc., 1992, vol. 1.

[11] J. Baik, G. Ben Arous, and S. Péché, "Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices," *Ann. Probab.*, vol. 33, no. 5, pp. 1643–1697, 2005.

[12] J. Baik and J. W. Silverstein, "Eigenvalues of large sample covariance matrices of spiked population models," *J. Multivariate Anal.*, vol. 97, no. 6, pp. 1382–1408, 2006.

[13] N. El Karoui, "Tracy-Widom limit for the largest eigenvalue of a large class of complex sample covariance matrices," *Ann. Probab.*, vol. 35, no. 2, pp. 663–714, 2007.

[14] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. on Ant. and Prop.*, vol. 34, no. 3, pp. 276–280, 1986.

[15] D. Donoho and J. Tanner, "Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing," *Phil. Trans. of the Roy. Soc. A: Math., Phys. and Engg. Sciences*, vol. 367, no. 1906, pp. 4273–4293, 2009.