

Tightness of the maximum likelihood semidefinite relaxation for angular synchronization

Afonso S. Bandeira¹ · Nicolas Boumal² ·
Amit Singer²

Received: 5 December 2014 / Accepted: 26 July 2016 / Published online: 8 August 2016
© Springer-Verlag Berlin Heidelberg and Mathematical Optimization Society 2016

Abstract Maximum likelihood estimation problems are, in general, intractable optimization problems. As a result, it is common to approximate the maximum likelihood estimator (MLE) using convex relaxations. In some cases, the relaxation is tight: it recovers the true MLE. Most tightness proofs only apply to situations where the MLE exactly recovers a planted solution (known to the analyst). It is then sufficient to establish that the optimality conditions hold at the planted signal. In this paper, we study an estimation problem (angular synchronization) for which the MLE is not a simple function of the planted solution, yet for which the convex relaxation is tight. To establish tightness in this context, the proof is less direct because the point at which to verify optimality conditions is not known explicitly. Angular synchronization consists in estimating a collection of n phases, given noisy measurements of the pairwise relative phases. The MLE for angular synchronization is the solution of a (hard) non-bipartite Grothendieck problem over the complex numbers. We consider a stochastic model for the data: a planted signal (that is, a ground truth set of phases) is corrupted with non-adversarial random noise. Even though the MLE does not coincide with the

✉ Nicolas Boumal
nboumal@princeton.edu
<http://www.math.princeton.edu/~nboumal/>

Afonso S. Bandeira
bandeira@cims.nyu.edu
<http://cims.nyu.edu/~bandeira/>

Amit Singer
amits@math.princeton.edu
<http://www.math.princeton.edu/~amits/>

¹ New York University, New York City, NY, USA

² Princeton University, Princeton, NJ, USA

planted signal, we show that the classical semidefinite relaxation for it is tight, with high probability. This holds even for high levels of noise.

Keywords Angular synchronization · Semidefinite programming · Tightness of convex relaxation · Maximum likelihood estimation

Mathematics Subject Classification 90C22 (Semidefinite programming) · 90C26 (Nonconvex programming, global optimization) · 62F10 (Point estimation)

1 Introduction

Recovery problems in statistics and many other fields are commonly solved under the paradigm of maximum likelihood estimation, partly due to the rich theory it enjoys. Unfortunately, in many important applications, the parameter space is exponentially large and non-convex, often rendering the computation of the maximum likelihood estimator (MLE) intractable. It is then common to settle for heuristics, such as expectation-maximization algorithms to name but one example. However, it is also common for such iterative heuristics to get trapped in local optima. Furthermore, even when these methods do attain a global optimum, there is in general no way to verify this.

A now classic alternative to these heuristics is the use of convex relaxations. The idea is to maximize the likelihood in a larger, convex set that contains the parameter space of interest, as (well-behaved) convex optimization problems are generally well understood and can be solved in polynomial time. The downside is that the solution obtained might not be in the original feasible (acceptable) set. One is then forced to take an extra, potentially suboptimal, rounding step.

This line of thought is the basis for a wealth of modern *approximation algorithms* [58]. One preeminent example is Goemans and Williamson’s treatment of Max-Cut [33], an NP-hard combinatorial problem that involves segmenting a graph in two clusters to maximize the number of edges connecting the two clusters. They first show that Max-Cut can be formulated as a semidefinite program (SDP)—a convex optimization problem where the variable is a positive semidefinite matrix—with the additional, non-convex constraint that the sought matrix be of rank one. Then, they propose to solve this SDP while relaxing (removing) the rank constraint. They show that the obtained solution, despite typically being of rank strictly larger than one, can be rounded to a (suboptimal) rank-one solution, and that it provides a guaranteed approximation of the optimal value of the hard problem. Results of the same nature abound in the recent theoretical computer science literature [10, 11, 19, 29, 34, 41, 52, 59, 59].

In essence, approximation algorithms insist on solving *all* instances of a given NP-hard problem in polynomial time, which, unless $P = NP$, must come at the price of accepting some degree of sub-optimality. This worst-case approach hinges on the fact that a problem is NP-hard as soon as every efficient algorithm for it can be hindered by at least one pathological instance.

Alternatively, in a non-adversarial setting where “the data is not an enemy,” one may find that such pathological cases are not prominent. As a result, the applied

mathematics community has been more interested in identifying regimes for which the convex relaxations are tight, that is, admit a solution that is also admissible for the hard problem. When this is the case, no rounding is necessary and a truly optimal solution is found in reasonable time, together with a certificate of optimality. This is sometimes achieved by positing a probability distribution on the instances of the problem and asserting tightness *with high probability*, as for example in compressed sensing [21, 31, 54], matrix completion [20], clustering [1] and inverse problems [7, 25]. In this approach, one surrenders the hope to solve all instances of the hard problem, in exchange for true optimality with high probability.

The results presented in this work are of that nature. As a concrete object of study, we consider the *angular synchronization problem* [11, 49], which consists in estimating a collection of n phases $e^{i\theta_1}, \dots, e^{i\theta_n}$, given noisy measurements of pairwise relative phases $e^{i(\theta_k - \theta_\ell)}$ —see Sect. 2 for a formal description. This problem notably comes up in time-synchronization of distributed networks [32], signal reconstruction from phaseless measurements [5, 9], ranking [27], digital communications [51], and surface reconstruction problems in computer vision [4] and optics [45]. Angular synchronization serves as a model for the more general problem of synchronization of rotations in any dimension, which comes up in structure from motion [36, 42], surface reconstruction from 3D scans [57] and cryo-electron microscopy [50], to name a few.

The main contribution of the present paper is a proof that, even though the angular synchronization problem is NP-hard [59], its MLE in the face of Gaussian noise can often be computed (and certified) in polynomial time. This remains true even for entry-wise noise levels growing to infinity as the size of the problem (the number of phases) grows to infinity. The MLE is obtained as the solution of a semidefinite relaxation described in Sect. 2. This arguably striking phenomenon has been observed empirically before [14] (see also Fig. 2), but not explained.

Computing the MLE for angular synchronization is equivalent to solving a non-bipartite Grothendieck problem. Semidefinite relaxations for Grothendieck problems have been thoroughly studied in theoretical computer science from the point of view of approximation ratios. The name is inspired by its close relation to an inequality of Grothendieck [35]. We direct readers to the survey by Pisier [44] for a discussion.

The proposed result is qualitatively different from most tightness results available in the literature. Typical results establish either exact recovery of a planted signal [1, 8, 37] (mostly in discrete settings), or exact recovery in the absence of noise, joint with stable (but not necessarily optimal) recovery when noise is present [22, 23, 28, 29, 57, 60]. In contrast, this paper shows optimal recovery even though exact recovery is not possible. In particular, Demanet and Jugnon showed stable recovery for angular synchronization via semidefinite programming, under adversarial noise [29]. We complement this by showing tightness in a non-adversarial setting, meaning the actual MLE is computed.

A similar semidefinite relaxation was studied in a digital communications context, where the parameters to estimate are m th roots of unity [51]. There, non-asymptotic results show the relaxation approximates the MLE within some factor, with high probability. The present paper is related to the limit $m \rightarrow \infty$, although the noise model considered is different, and we focus on exact MLE computation.

Our proof relies on verifying that a certain candidate dual certificate is valid with high probability. The main difficulty comes from the fact that the dual certificate

depends on the MLE, which does not coincide with the planted signal, and is a nontrivial function of the noise. We use necessary optimality conditions of the hard problem to both obtain an explicit expression for the candidate dual certificate, and to partly characterize the point whose optimality we aim to establish. This seems to be required since the MLE is not known in closed form.

In the context of sparse recovery, a result with similar flavor is support recovery guarantee [54], where the support of the estimated signal is shown to be contained in the support of the original signal. Due to the noise, exact recovery is also impossible in this setting. Another example is a recovery guarantee in the context of latent variable selection in graphical models [24].

Besides the relevance of angular synchronization in and of its own, we are confident this new insight will help uncover similar results in other applications where it has been observed that semidefinite relaxations can be tight even when the ground truth cannot be recovered. Notably, this appears to be the case for the Procrustes and multi-reference alignment problems [13, 14]).

The crux of our argument concerns the rank of the solutions of an SDP. We mention in passing that there are many other deterministic results in the literature pertaining to the rank of solutions of SDP's. For example, it has been shown that, in general, an SDP with only equality constraints admits a solution of rank at most (on the order of) the square root of the number of constraints, see [16, 43, 48]. Furthermore, Sagnol [47] shows that under some conditions (that are not fulfilled in our case), certain SDP's related to packing problems always admit a rank-one solution. Sojoudi and Lavaei [53] study a class of SDP's on graphs which is related to ours and for which, under certain strong conditions on the topology of the graphs, the SDP's admit rank-one solutions—see also applications to power flow optimization [40].

1.1 Contribution

We frame angular synchronization as the problem of estimating $z \in \mathbb{C}^n$ with $|z_i| = 1 \forall i$, given measurements $C = zz^* + \sigma W$, where $W = W^*$ has i.i.d. complex standard Gaussian entries above its diagonal. The MLE x maximizes x^*Cx over the parameter space (QP). This is relaxed to $\max \text{Tr}(CX)$ s.t. $\text{diag}(X) = \mathbb{1}$, $X \succeq 0$ (SDP). After proving exact recovery for *real* x and X (Fig. 1), we focus on the *complex* case (Fig. 2). Our main result is that, if $\sigma \leq \frac{1}{18}n^{1/4}$, then, with high probability, (SDP) admits a unique solution of rank 1, revealing the global optimum of (QP) (Thm. 2.1). In order to do so, we characterize global optimizers of (QP) as points x which satisfy first- and second-order necessary optimality conditions and perform at least as well as z in terms of the cost function of (QP) (under conditions on the perturbation σW).

1.2 Notation

For $a \in \mathbb{C}$, \bar{a} denotes its complex conjugate and $|a| = \sqrt{a\bar{a}}$ its modulus. For $v \in \mathbb{C}^n$, v^* denotes its conjugate transpose, $\|v\|^2 = \|v\|_2^2 = v^*v$, $\|v\|_\infty = \max_i |v_i|$ and $\|v\|_1 = \sum_i |v_i|$. $\mathbb{1}$ is a vector of all-ones. $\text{diag}(v)$ is a diagonal matrix with $\text{diag}(v)_{ii} = v_i$. For a matrix M , $\|M\|_{\text{op}}$ is the maximal singular value and $\Re(M)$, $\Im(M)$ extract the real and

Proportion of exact recovery (real case)

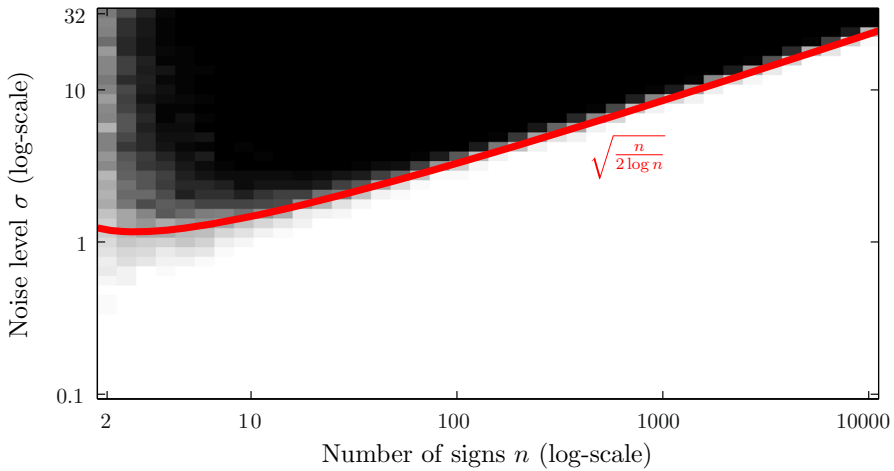


Fig. 1 In the real case (QPR), one can hope to recover the signal $z \in \{\pm 1\}^n$ exactly from the pairwise sign comparisons $zz^T + \sigma W$. This figure shows how frequently the semidefinite relaxation (SDPR) returns the correct signal z (or $-z$). For each pair (n, σ) , 100 realizations of the noise W are generated independently and the dual certificate for the true signal (2.1) is verified (it is declared numerically positive semidefinite if its smallest eigenvalue exceeds $-10^{-14}n$). The frequency of success is coded by intensity (*bright* for 100 % success, *dark* for 0 % success). The results are in excellent agreement with the theoretical predictions (2.2) (*red curve*)

imaginary parts; $\text{diag}(M)$ extracts the diagonal of M into a vector and $\text{ddiag}(M)$ sets all off-diagonal entries of M to zero. We use the real inner product $\langle u, v \rangle = \Re\{u^*v\}$ over \mathbb{C}^n . \mathbb{E} denotes mathematical expectation.

2 The angular synchronization problem

We focus on the problem of *angular synchronization* [11,49], in which one wishes to estimate a collection of n phases ($n \geq 2$) based on measurements of pairwise phase differences. We restrict our analysis to the case where a measurement is available for every pair of phases. More precisely, we let $z \in \mathbb{C}^n$ be an unknown, complex vector with unit modulus entries, $|z_1| = \dots = |z_n| = 1$, and we consider measurements of the form $C_{ij} = z_i \bar{z}_j + \varepsilon_{ij}$, where $\varepsilon_{ij} \in \mathbb{C}$ is noise affecting the measurement. By symmetry, we define $C_{ji} = \overline{C_{ij}}$ and $C_{ii} = 1$, so that the matrix $C \in \mathbb{C}^{n \times n}$ whose entries are given by the C_{ij} 's is Hermitian.

Further letting the noise ε_{ij} be i.i.d. (complex) Gaussian variables for $i < j$, it follows that an MLE for z is any vector of phases x minimizing $\sum_{i,j} |C_{ij}x_j - x_i|^2$. Equivalently, an MLE is a solution of the following quadratically constrained quadratic program (sometimes called the complex constant-modulus QP [41, Table 2] in the optimization literature, and non-bipartite Grothendieck problem in theoretical computer science):

$$\max_{x \in \mathbb{C}^n} x^* C x, \quad \text{subject to } |x_1| = \dots = |x_n| = 1, \tag{QP}$$




Fig. 2 In the complex case (QP), exact recovery of the phases z from the pairwise relative phase measurements $zz^* + \sigma W$ is hopeless as soon as $\sigma > 0$. Nevertheless, below the *middle line*, the perturbation is smaller than the signal (in operator norm, assuming z -discordance), and one expects to be able to recover z reasonably well. And indeed, the maximum likelihood estimator (MLE) for z is close to z , as per Lemmas 4.1 and 4.2. Computing the MLE is hard in general, but solving the semidefinite relaxation (SDP) is tractable. When (SDP) has a rank-one solution, that solution coincides with the MLE. This figure shows, empirically, how frequently the (SDP) admits a unique rank-one solution (same color code as Fig. 1). For each pair (n, σ) , 100 realizations of the noise W are generated independently and (SDP) is solved using a complex version of the low-rank algorithm in [17]. The (SDP) appears to be tight for remarkably large levels of noise. Theorem 2.1 (our main contribution) partly explains this phenomenon, by showing that the SDP is tight with high probability below the *bottom line*. We further note that, above the *top line*, no unbiased estimator for z performs better than a random guess [18]. Note that, for $n \leq 3$, a complex version of [43, Thm. 2.1] guarantees deterministic tightness of the relaxation, in accordance with observation here

where x^* denotes the conjugate transpose of x . This problem can only be solved up to a global phase, since only relative information is available. Indeed, given any solution x , all vectors of the form $xe^{i\theta}$ are equivalent solutions, for arbitrary phase θ . Solving (QP) is, in general, an NP-hard problem [59, Prop. 3.5]. It is thus unlikely that there exists an algorithm capable of solving (QP) in polynomial time for an arbitrary cost matrix C . In response, building upon now classical techniques, various authors [1, 2, 10, 49, 52, 59] study the following convex relaxation of (QP). For any admissible x , the Hermitian matrix $X = xx^* \in \mathbb{C}^{n \times n}$ is Hermitian positive semidefinite, has unit diagonal entries and is of rank one. Conversely, any such X may be written in the form $X = xx^*$ such that x is admissible for (QP). In this case, the cost function can be rewritten in linear form: $x^*Cx = \text{Tr}(x^*Cx) = \text{Tr}(CX)$. Dropping the rank constraint then yields the relaxation we set out to study:

$$\max_{X \in \mathbb{C}^{n \times n}} \text{Tr}(CX), \quad \text{subject to } \text{diag}(X) = \mathbb{1} \text{ and } X \succeq 0. \quad (\text{SDP})$$

Such relaxations *lift* the problem to higher dimensional spaces. Indeed, the search space of (QP) has dimension n (or $n - 1$, discounting the global phase) whereas

the search space of (SDP) has real dimension $n(n - 1)$. In general, increasing the dimension of an optimization problem may not be advisable. But in this case, the relaxed problem is a semidefinite program. Such optimization problems can be solved to global optimality up to arbitrary precision in polynomial time [55].

It is known that the solution of (SDP) can be rounded to an approximate solution of (QP), with a guaranteed approximation ratio [52, § 4]. But even better, when (SDP) admits an optimal solution X of rank one, then no rounding is necessary: the leading eigenvector x of $X = xx^*$ is a global optimum of (QP), meaning we have solved the original problem exactly. Elucidating when the semidefinite program admits a solution of rank one, i.e., when the relaxation is *tight*, is the focus of the present paper.

2.1 A detour through synchronization over $\mathbb{Z}_2 = \{\pm 1\}$

Problem (QP) is posed over the complex numbers. As a result, the individual variables x_i in (QP) live on a continuous search space (the unit circle). One effect of this is that even small noise on the data precludes exact recovery of the signal z (in general). This is the root of most of the complications that will arise in the developments hereafter. In order to first illustrate some of the ideas of the proof in a simpler context, this section proposes to take a detour through the real case. Besides this expository rationale, the real case is interesting in and of itself. It notably relates to correlation clustering [1] and the stochastic block model [2]. The analysis proposed here also appears in [12].

Let $z \in \{\pm 1\}^n$ be the signal to estimate and let $C = zz^T + \sigma W$ contain the measurements, with $W = W^T$ a Wigner matrix: its above-diagonal entries are i.i.d. (real) standard normal random variables, and its diagonal entries are zero. Each entry C_{ij} is a noisy measurement of the relative sign $z_i z_j$. For example, z_i could represent political preference of an agent (left or right wing) and C_{ij} could be a measurement of agreement between two agents' views [2]. Consider this MLE problem:

$$\max_{x \in \mathbb{R}^n} x^T C x, \quad \text{subject to } |x_1| = \cdots = |x_n| = 1. \quad (\text{QPR})$$

Thus, $x_i \in \{\pm 1\}$ for each i . The corresponding relaxation reads:

$$\max_{X \in \mathbb{R}^{n \times n}} \text{Tr}(CX), \quad \text{subject to } \text{diag}(X) = \mathbb{1} \text{ and } X \succeq 0. \quad (\text{SDPR})$$

Certainly, if (SDPR) admits a rank-one solution, then that solution reveals a global optimizer of (QPR). Beyond rank recovery, because the variables x_i are discrete, it is expected that, if noise is sufficiently small, then the global optimizer of (QPR) will be the true signal z . Figure 1 confirms this even for large noise, leveraging the tightness of the relaxation (SDPR). Note that exact recovery is a strictly stronger requirement than rank recovery; for example, for $n = 2$, the relaxation is always tight [43], but it is not always exact. We now investigate this exact recovery phenomenon for (SDPR).

The aim is to show that $X = zz^T$ is a solution of (SDPR). Semidefinite programs admit a dualproblem [55], which for (SDPR) reads:

$$\min_{S \in \mathbb{R}^{n \times n}} \text{Tr}(S + C), \text{ s.t. } S + C \text{ is diagonal and } S \succeq 0. \tag{DSDP\mathbb{R}}$$

Strong duality holds, which implies that a given feasible X is optimal if and only if there exists a dual feasible matrix S such that $\text{Tr}(CX) = \text{Tr}(S + C)$, or, in other words, such that $\text{Tr}(SX) = 0$.¹ Since both S and X are positive semidefinite, this is equivalent to requiring $SX = 0$ (a condition known as *complementary slackness*), and hence requiring $Sz = 0$.

For ease of exposition, we now assume (without loss of generality) that $z = \mathbb{1}$. Tentatively, let $S = nI + \sigma \text{diag}(W\mathbb{1}) - C$ —for the complex case, we will see how to obtain this candidate without guessing. Then, by construction, $S + C$ is diagonal and $S\mathbb{1} = 0$.² It remains to determine under what conditions S is positive semidefinite.

Define the diagonal matrix $D_W = \text{diag}(W\mathbb{1})$ and the (Laplacian-like) matrix $L_W = D_W - W$. The candidate dual certificate is a sum of two Laplacian-like matrices:

$$S = nI - \mathbb{1}\mathbb{1}^\top + \sigma L_W. \tag{2.1}$$

The first part is the Laplacian of the complete graph. It has eigenvalues $0, n, \dots, n$, with the zero eigenvalue corresponding to the all-ones vector $\mathbb{1}$. Since $L_W\mathbb{1} = 0$ by construction, S is positive semidefinite as long as σL_W does not “destroy” any of the large eigenvalues n . This is guaranteed in particular if $\|\sigma L_W\|_{\text{op}} < n$. It is well-known from concentration results about Wigner matrices that, for all $\varepsilon > 0$, with high probability for large n , $\|L_W\|_{\text{op}} \leq \sqrt{(2 + \varepsilon)n \log n}$ (see [30, Thm. 1] and [12]). Thus, it is expected that (SDP\mathbb{R}) will yield exact recovery of the signal z (with high probability) as long as

$$\sigma < \sqrt{\frac{n}{(2 + \varepsilon) \log n}}. \tag{2.2}$$

This is indeed compatible with the empirical observation of Fig. 1.

2.2 Back to synchronization over SO(2)

We now return to the complex case, which is the focus of this paper. As was mentioned earlier, in the presence of even the slightest noise, one can no longer reasonably expect the true signal z to be an optimal solution of (QP) (this can be further quantified using Cramér-Rao bounds [18]). Nevertheless, we set out to show that (under some assumptions on the noise) solutions of (QP) are close to z and they can be computed via (SDP).

The proof follows that of the real case in spirit, but requires more sophisticated arguments because the solution is no longer known explicitly. This is important because the candidate dual certificate S itself depends on that solution. With this in mind, the proof of the upcoming main lemma (Lemma 3.2) follows this reasoning:

¹ Indeed, $S + C$ is diagonal and $X_{ij} = 1$, hence $\text{Tr}(S + C) = \text{Tr}((S + C)X) = \text{Tr}(SX) + \text{Tr}(CX)$.

² Using $C = \mathbb{1}\mathbb{1}^\top + \sigma W$, we get $S\mathbb{1} = n\mathbb{1} + \sigma W\mathbb{1} - n\mathbb{1} - \sigma W\mathbb{1} = 0$.

1. For small enough noise levels σ , any optimal solution x of (QP) is close to the sought signal z (Lemmas 4.1 and 4.2).
2. Solutions x , a fortiori, satisfy necessary optimality conditions for (QP). First-order conditions take up the form $Sx = 0$, where $S = \Re\{\text{ddiag}(Cxx^*)\} - C$ depends smoothly on x (see (4.9)). Second-order conditions will also be used.
3. Remarkably, this S can be used as a dual certificate for solutions of (SDP). Indeed, $X = xx^*$ is optimal if and only if S is positive semidefinite (Lemma 4.4). The solution is unique if $\text{rank}(S) = n - 1$ (Lemma 4.3). Thus, it only remains to study the eigenvalues of S .
4. In the absence of noise, S is a Laplacian for a complete graph with unit weights (up to a unitary transformation), so that its eigenvalues are 0 with multiplicity 1, and n with multiplicity $n - 1$. Then, $X = zz^*$ is always the unique solution.
5. Adding small noise, because of the first point, the solution x will move only by a small amount, and hence so will S . Thus, the large eigenvalues should be controllable into remaining positive (Sect. 4.4).
6. The crucial fact follows: because of the way S is constructed (using first-order optimality conditions), the zero eigenvalue is “pinned down” (as long as x is a local optimum of (QP)). Indeed, both x and S change as a result of adding noise, but the property $Sx = 0$ remains valid. Thus, there is no risk that the zero eigenvalue from the noiseless scenario would become negative when noise is added.

Following this road map, most of the work in the proof below consists in bounding how far away x can be from z , and in using that to control the large eigenvalues of S . This constructive way of identifying the dual certificate (third point in the roadmap) already appears explicitly in work by Journée et al. [38], who considered a different family of real, semidefinite programs which also admit a smooth geometry when the rank is constrained. This points to smoothness of (QP) (and non-degeneracy of (SDP) [6]) as a principal ingredient in our analysis: the KKT conditions of (QP) are a subset of the KKT conditions of its relaxation. It is because the former are explicit (rather than existential as in Lemma 4.3) that they help in identifying S .

Our main theorem follows. In a nutshell, it guarantees that: under (complex) Wigner noise W , with high probability, solutions of (QP) are close to z , and, assuming the noise level σ is smaller than (on the order of) $n^{1/4}$, (SDP) admits a unique solution, it is of rank one and identifies the solution of (QP) (unique, up to a global phase shift).

Theorem 2.1 *Let $z \in \mathbb{C}^n$ be a vector with unit modulus entries, $W \in \mathbb{C}^{n \times n}$ a Hermitian Gaussian Wigner matrix and let $C = zz^* + \sigma W$. Let $x \in \mathbb{C}^n$ be a global optimizer of (QP) satisfying $z^*x = |z^*x|$ (fixing the global phase). With probability at least $1 - \mathcal{O}(n^{-5/4})$, x is close to z in the following two senses:*

$$\|x - z\|_\infty \leq 6 \left(\sqrt{\log n} + 29\sigma \right) \sigma n^{-1/2}, \text{ and}$$

$$\|x - z\|_2 \leq 12\sigma.$$

Furthermore, if

$$\sigma \leq \frac{1}{18} n^{1/4}, \tag{2.3}$$

then the semidefinite program (SDP), given by

$$\max_{X \in \mathbb{C}^{n \times n}} \text{Tr}(CX), \text{ subject to } \text{diag}(X) = \mathbb{1} \text{ and } X \geq 0,$$

has, as its unique solution, the rank-one matrix $X = xx^*$.

The numerical experiments (Fig. 2) suggest it should be possible to allow σ to grow at a rate of $\frac{n^{1/2}}{\text{polylog}(n)}$ (as in the real case), but we were not able to establish that (see Remark 4.6). Nevertheless, we do show that σ can grow unbounded with n . To the best of our knowledge, this is the first result of this kind. We hope it might inspire similar results in other problems where the same phenomenon has been observed [14].

3 Main result

In this section we present our main technical result and show how it can be used to prove Theorem 2.1. We begin with a central definition in this paper. Intuitively, this definition characterizes non-adversarial noise matrices W .³

Definition 3.1 (*z-discordant matrix*) Let $z \in \mathbb{C}^n$ be a vector with unit modulus entries. A matrix $W \in \mathbb{C}^{n \times n}$ is called *z-discordant* if it is Hermitian and satisfies both

$$\|W\|_{\text{op}} \leq 3\sqrt{n}, \quad \text{and} \quad \|Wz\|_{\infty} \leq 3\sqrt{n \log n}.$$

The next lemma is the main technical contribution of this paper. It is a deterministic, non-asymptotic statement.

Lemma 3.2 Let $z \in \mathbb{C}^n$ be a vector with unit modulus entries, let $W \in \mathbb{C}^{n \times n}$ be a Hermitian, *z-discordant* matrix (Definition 3.1), and let $C = zz^* + \sigma W$. Let $x \in \mathbb{C}^n$ be a global optimizer of (QP) satisfying $z^*x = |z^*x|$. x is close to z in the following two senses:

$$\|x - z\|_{\infty} \leq 6 \left(\sqrt{\log n} + 29\sigma \right) \sigma n^{-1/2}, \text{ and}$$

$$\|x - z\|_2 \leq 12\sigma.$$

Furthermore, if $\sigma \leq \frac{1}{18}n^{1/4}$, then the semidefinite program (SDP) has, as its unique solution, the rank-one matrix $X = xx^*$.

We defer the proof of Lemma 3.2 to Sect. 4. The following proposition, whose proof we defer to ‘‘Appendix’’, shows how this lemma can be used to prove Theorem 2.1.

Proposition 3.3 Let $z \in \mathbb{C}^n$ be a (deterministic) vector with unit modulus entries. Let $W \in \mathbb{C}^{n \times n}$ be a Hermitian, random matrix with i.i.d. off-diagonal entries following a standard complex normal distribution and zeros on the diagonal. Thus, $W_{ii} = 0$, $W_{ij} = \overline{W_{ji}}$, $\mathbb{E}W_{ij} = 0$ and $\mathbb{E}|W_{ij}|^2 = 1$ (for $i \neq j$). Then, W is *z-discordant* with probability at least $1 - 2n^{-5/4} - e^{-n/2}$.

³ A similar but different definition appeared in a previous version of this paper.

The latter result is not surprising. Indeed, the definition of z -discordance requires two elements. Namely, (i) that W be not too large as an operator, and (ii) that no row of W be too aligned with z . For W a Wigner matrix independent of z , those are indeed expected to hold. In fact, for Gaussian noise, the constants 3 can be replaced by $2 + \varepsilon$ for any $\varepsilon > 0$, provided n is large enough.

The definition of z -discordance is not tightly adjusted to Wigner noise. As a result, it is expected that Lemma 3.2 will be applicable to show tightness of semidefinite relaxations for a larger span of noise models.

4 The proof

In this section, we prove Lemma 3.2. See Sect. 2.2 for an outline of the proof.

4.1 Global optimizers of (QP) are close to z

This subsection focuses on bounding the maximum likelihood estimation error, in ℓ_2 and ℓ_∞ norms. Later on, only the ℓ_2 bound is used to establish tightness of the SDP.

Lemma 4.1 *If W is z -discordant and $x \in \mathbb{C}^n$ verifies $\|x\|_2^2 = n$ and $z^*Cz \leq x^*Cx$ (in particular, if x is a global optimizer of (QP)), then*

$$\min_{\theta \in \mathbb{R}} \|xe^{i\theta} - z\|_2^2 = 2(n - |z^*x|) \leq 144\sigma^2.$$

Proof The ℓ_2 error and the correlation between z and x are related as follows:

$$\min_{\theta \in \mathbb{R}} \|xe^{i\theta} - z\|_2^2 = 2 \left(n - \max_{\theta \in \mathbb{R}} \Re\{e^{i\theta} z^*x\} \right) = 2(n - |z^*x|). \quad (4.1)$$

Without loss of generality, assume the global phase of x is such that $z^*x = |z^*x|$, i.e., x and z are optimally aligned. Expand the inequality $z^*Cz \leq x^*Cx$ using the data model $C = zz^* + \sigma W$ to obtain

$$n^2 + \sigma z^*Wz \leq |z^*x|^2 + \sigma x^*Wx.$$

Group noise terms as follows:

$$n^2 - |z^*x|^2 \leq \sigma (x^*Wx - z^*Wz).$$

Since $n^2 - |z^*x|^2 = (n - |z^*x|)(n + |z^*x|)$, divide both sides by $n + |z^*x| \geq n$ to obtain

$$n - |z^*x| \leq \sigma (x^*Wx - z^*Wz) n^{-1}. \quad (4.2)$$

The difference of quadratic terms should be small if x and z are close. Indeed, using $\|W\|_{\text{op}} \leq 3n^{1/2}$ (by z -discordance) and $\|x\|_2 = \|z\|_2 = n^{1/2}$:

$$\begin{aligned} x^*Wx - z^*Wz &= \Re\{(x - z)^*W(x + z)\} \\ &\leq \|x - z\|_2 \cdot \|W\|_{\text{op}} \cdot \|x + z\|_2 \\ &\leq 6n\|x - z\|_2. \end{aligned}$$

Combine this and (4.1) with (4.2) to obtain $\|x - z\|_2 \leq 12\sigma$. □

Note that the constant 12 is pessimistic, because we used $n + |z^*x| \geq n$ even though it is closer to $2n$. Assuming $\|x - z\|_2 \leq \alpha\sigma$ and $\sigma \leq t\sqrt{n}$, the argument above can be bootstrapped to reduce the constant. One obtains $\|x - z\|_2 \leq \alpha_k\sigma$ for all k , with $\alpha_0 = 12$ and $\alpha_{k+1} = 6 + (t/2)^2\alpha_k^3$. For $t < 1/6\sqrt{2}$ small enough, α_k converges arbitrarily close to 6. For example, if $\sigma \leq \sqrt{n}/12$, then $\|x - z\|_2 \leq 6.5\sigma$.

The next lemma establishes a bound on the largest individual error, $\|x - z\|_\infty$, after proper global phase alignment of x and z . Interestingly, for $\sigma \ll n^{1/4}$, the bound shows that individual errors decay, uniformly, as n increases.

Lemma 4.2 *If W is z -discordant and x is a global optimizer of (QP) with global phase such that $z^*x = |z^*x|$ (that is, the phase that minimizes the ℓ_2 distance), then*

$$\|x - z\|_\infty \leq 6 \left(\sqrt{\log n} + 29\sigma \right) \sigma n^{-1/2}.$$

Proof We wish to upper bound, for all $i \in \{1, 2, \dots, n\}$, the value of $|x_i - z_i|$. Let $e_i \in \mathbb{R}^n$ denote the i th vector of the canonical basis (its i th entry is 1 whereas all other entries are zero). Consider $\hat{x} = x + (z_i - x_i)e_i$, a feasible point of (QP) obtained from the optimal x by changing its i th entry to the corresponding entry of z . Since x is optimal, it performs at least as well as \hat{x} according to the cost function of (QP):

$$x^*Cx \geq \hat{x}^*C\hat{x} = x^*Cx + |z_i - x_i|^2C_{ii} + 2\langle z_i - x_i, (Cx)_i \rangle, \tag{4.3}$$

with the (real) inner product (4.6) over \mathbb{C} . Since $C_{ii} = 1 \geq 0$,⁴ it follows that the right-most term is nonpositive. Using $C = zz^* + \sigma W$, we have

$$\langle z_i - x_i, (z^*x)z_i + \sigma(Wx)_i \rangle \leq 0.$$

Observe that

$$|x_i - z_i|^2 = 2(1 - \langle x_i, z_i \rangle) = 2\langle z_i - x_i, z_i \rangle, \tag{4.4}$$

and combine with the assumption $z^*x = |z^*x|$ to obtain

$$|z^*x||x_i - z_i|^2 \leq 2\sigma \langle x_i - z_i, (Wx)_i \rangle \leq 2\sigma|x_i - z_i|(Wx)_i. \tag{4.5}$$

⁴ The inequality is independent of C_{ii} . Assuming nonnegativity merely eases the exposition.

This holds for all i , hence

$$|z^*x| \|x - z\|_\infty \leq 2\sigma \|Wx\|_\infty.$$

Invoke Lemma 4.1, namely, $|z^*x| \geq n - 72\sigma^2$, to get

$$n \|x - z\|_\infty \leq 2\sigma \|Wx\|_\infty + 72\sigma^2 \|x - z\|_\infty.$$

Since $\|x - z\|_\infty \leq 2$, it finally comes that

$$\|x - z\|_\infty \leq 2 (\|Wx\|_\infty + 72\sigma) \sigma n^{-1}.$$

We discuss the problem of bounding $\|Wx\|_\infty$ in more details later on. For now, we simply use the suboptimal bound (4.11) (obtained independently of the present lemma), i.e., $\|Wx\|_\infty \leq 3\sqrt{n \log n} + 36\sigma\sqrt{n}$. Then, for all $n \geq 2$, using $24/\sqrt{2} \leq 17$,

$$\|x - z\|_\infty \leq 6 \left(\sqrt{\log n} + 12\sigma + 24 \frac{\sigma}{\sqrt{n}} \right) \sigma n^{-1/2} \leq 6 \left(\sqrt{\log n} + 29\sigma \right) \sigma n^{-1/2}.$$

(For $\sigma \ll n^{1/2}$, the constant 29 could be replaced by one arbitrarily close to 12.) \square

4.2 Optimality conditions for (SDP)

The global optimizers of the semidefinite program (SDP) can be characterized completely via the Karush-Kuhn-Tucker (KKT) conditions.

Lemma 4.3 *A Hermitian matrix $X \in \mathbb{C}^{n \times n}$ is a global optimizer of (SDP) if and only if there exists a Hermitian matrix $\hat{S} \in \mathbb{C}^{n \times n}$ such that all of the following hold:*

1. $\text{diag}(X) = \mathbb{1}$;
2. $X \succeq 0$;
3. $\hat{S}X = 0$;
4. $\hat{S} + C$ is (real) diagonal; and
5. $\hat{S} \succeq 0$.

If, furthermore, $\text{rank}(\hat{S}) = n - 1$, then X has rank one and is the unique global optimizer of (SDP).

Proof These are the KKT conditions of (SDP) [46, Example 3.36]. Conditions 1 and 2 are primal feasibility, condition 3 is complementary slackness and conditions 4 and 5 encode dual feasibility. Since the identity matrix I is strictly feasible, the Slater condition is fulfilled. This ensures the KKT conditions are necessary and sufficient for global optimality [46, Theorem 3.34]. Slater's condition also holds for the dual. Indeed, let $\tilde{S} = \alpha I - C$, where $\alpha \in \mathbb{R}$ is such that $\tilde{S} \succ 0$ (such an α always exists); then $\tilde{S} + C$ is indeed diagonal and \tilde{S} is strictly admissible for the dual. This allows to use results from [6]. Specifically, assuming $\text{rank}(\hat{S}) = n - 1$, Theorem 9 in [6] implies that \hat{S} is dual nondegenerate. Then, since \hat{S} is also optimal for the dual (by

complementary slackness), Theorem 10 in [6] guarantees that the primal solution X is unique. Since X is nonzero and $\hat{S}X = 0$, it must be that $\text{rank}(X) = 1$. \square

Certainly, if (SDP) admits a rank-one solution, it has to be of the form $X = xx^*$, with x an optimal solution of the original problem (QP). Based on this consideration, our proof of Lemma 3.2 goes as follows. We let x denote a global optimizer of (QP) and we consider $X = xx^*$ as a candidate solution for (SDP). Using the optimality of x and assumptions on the noise, we then construct and verify a dual certificate matrix \hat{S} as required per Lemma 4.3. In such proofs, one of the nontrivial parts is to guess an analytical form for \hat{S} given a candidate solution X . We achieve this by inspecting the first-order optimality conditions of (QP) (which x necessarily satisfies). The main difficulty is then to show the suitability of the candidate S , as it depends nonlinearly on the global optimum x , which itself is a complicated function of the noise W . We show feasibility of S via a program of inequalities, relying on z -discordance of the noise W (see Definition 3.1).

4.3 Construction of the dual certificate S

Every global optimizer of the combinatorial problem (QP) must, a fortiori, satisfy first-order necessary optimality conditions. We derive those now.

Endow the complex plane \mathbb{C} with the Euclidean metric

$$\langle y_1, y_2 \rangle = \Re\{y_1^* y_2\}. \tag{4.6}$$

This is equivalent to viewing \mathbb{C} as \mathbb{R}^2 with the canonical inner product, using the real and imaginary parts of a complex number as its first and second coordinates. Denote the complex circle by

$$S = \{y \in \mathbb{C} : y^* y = 1\}.$$

The circle can be seen as a submanifold of \mathbb{C} , with tangent space at each y given by (simply differentiating the constraint):

$$T_y S = \{\dot{y} \in \mathbb{C} : \dot{y}^* y + y^* \dot{y} = 0\} = \{\dot{y} \in \mathbb{C} : \langle y, \dot{y} \rangle = 0\}.$$

Restricting the Euclidean inner product to each tangent space equips S with a Riemannian submanifold geometry. The search space of (QP) is exactly S^n , itself a Riemannian submanifold of \mathbb{C}^n with the product geometry. Thus, problem (QP) consists in *minimizing* a smooth function

$$g(x) = -x^* C x$$

over the smooth Riemannian manifold S^n . Therefore, the first-order necessary optimality conditions for (QP) (i.e., the KKT conditions) can be stated simply as $\text{grad } g(x) = 0$, where $\text{grad } g(x)$ is the Riemannian gradient of g at $x \in S^n$ [3].

This gradient is given by the orthogonal projection of the Euclidean (the classical) gradient of g onto the tangent space of \mathcal{S}^n at x [3, Eq. (3.37)],

$$\begin{aligned} T_x \mathcal{S}^n &= T_{x_1} \mathcal{S} \times \cdots \times T_{x_n} \mathcal{S} \\ &= \{ \dot{x} \in \mathbb{C}^n : \Re\{\text{diag}(\dot{x}x^*)\} = 0 \}. \end{aligned} \quad (4.7)$$

The projector and the Euclidean gradient are given respectively by:

$$\begin{aligned} \text{Proj}_x : \mathbb{C}^n &\rightarrow T_x \mathcal{S}^n : \dot{x} \mapsto \text{Proj}_x \dot{x} = \dot{x} - \Re\{\text{ddiag}(\dot{x}x^*)\}x, \\ \nabla g(x) &= -2Cx, \end{aligned}$$

where $\text{ddiag} : \mathbb{C}^{n \times n} \rightarrow \mathbb{C}^{n \times n}$ sets all off-diagonal entries of a matrix to zero. Thus, for x a global optimizer of (QP), it holds that

$$0 = \text{grad } g(x) = \text{Proj}_x \nabla g(x) = 2(\Re\{\text{ddiag}(Cxx^*)\} - C)x. \quad (4.8)$$

This suggests the following definitions:

$$X = xx^*, \quad S = \Re\{\text{ddiag}(Cxx^*)\} - C. \quad (4.9)$$

Note that S is Hermitian and $Sx = 0$. Referring to the KKT conditions in Lemma 4.3, it follows immediately that X is feasible for (SDP) (conditions 1 and 2); that $SX = (Sx)x^* = 0$ (condition 3); and that $S + C$ is a diagonal matrix (condition 4). It thus only remains to show that S is also positive semidefinite and has rank $n - 1$. If such is the case, then X is the unique global optimizer of (SDP). Note the special role of the first-order necessary optimality conditions: they guarantee complementary slackness, without requiring further work.

We will also use second-order necessary optimality conditions, namely, that the (Riemannian) Hessian of g at an optimizer x is positive semidefinite on the tangent space (4.7). The action of the Riemannian Hessian is obtained by projecting the directional derivatives of the Riemannian gradient [3, Eq. (5.15)]. Explicitly, for any tangent vector $\dot{x} \in T_x \mathcal{S}^n$, one can compute that (with $Dh(x)[\dot{x}]$ denoting the directional derivative of h at x along \dot{x})

$$\text{Hess } g(x)[\dot{x}] = \text{Proj}_x D \text{grad } g(x)[\dot{x}] = \text{Proj}_x 2S\dot{x}.$$

Thus, if x is a (local) optimizer, then, since Proj_x is self-adjoint,

$$\langle \dot{x}, \text{Hess } g(x)[\dot{x}] \rangle = 2 \langle \dot{x}, S\dot{x} \rangle \geq 0$$

for all $\dot{x} \in T_x \mathcal{S}^n$: a necessary but insufficient step towards making S positive semidefinite. We will later use this condition along selected directions.

The following lemma further shows that S is the right candidate dual certificate. More precisely, for x a critical point of (QP), it is necessary and sufficient for S to be positive semidefinite in order for $X = xx^*$ to be optimal for (SDP). This confirms that, by studying S , nothing is lost with respect to the original question. See also [17].

Lemma 4.4 *A feasible X (of any rank) for (SDP) is optimal if and only if $S = \Re\{\text{ddiag}(CX)\} - C$ (4.9) is positive semidefinite. There exists no other certificate.*

Proof The *if* part follows from Lemma 4.3: set $\hat{S} = S$ and observe that, since $\text{Tr}(SX) = 0$ by construction, and since $S, X \succeq 0$, it follows that $SX = 0$. We now show the *only if* part. Assume X is optimal. Then, by Lemma 4.3, there exists $\hat{S} \succeq 0$ which satisfies $\hat{S}X = 0$ and $\hat{S} + C = \hat{D}$, where \hat{D} is diagonal. Thus, $CX = (\hat{D} - \hat{S})X = \hat{D}X$ and $\Re\{\text{ddiag}(CX)\} = \hat{D}$. Consequently, $S = \hat{D} - C = \hat{S}$. \square

4.4 A sufficient condition for rank recovery

Knowing which certificate S (4.9) to verify for a given x , it remains to characterize the point x . Of course, x is the global optimizer of (QP), but this is not a convenient property to exploit. Instead, we let x be a *second-order critical point*⁵ for (QP) which outperforms the planted signal z (certainly, the MLE is one such point). In effect, we prove the following, which immediately yields Lemma 3.2.

Proposition 4.5 *Let $z \in \mathbb{C}^n$ have unit-modulus entries, W be z -discordant, $\sigma \leq \frac{1}{18}n^{1/4}$ and $C = zz^* + \sigma W$. If $x \in \mathbb{C}^n$ is a second-order critical point for (QP) such that $x^*Cx \geq z^*Cz$, then xx^* is the unique solution of (SDP) and x solves (QP).*

To this end, we prove that, at such points x , the certificate S is positive semidefinite of rank $n - 1$. As a first observation, note that x being critical ($Sx = 0$) implies that $\Re\{(Cx)_i \bar{x}_i\} x_i = (Cx)_i$ for all i . Thus, $(Cx)_i \bar{x}_i$ is real, and it follows that

$$S = \text{ddiag}(Cxx^*) - C.$$

Furthermore, since x is second-order critical, the Hessian of g at x is positive semidefinite, implying $\langle \dot{x}, S\dot{x} \rangle \geq 0$ for all $\dot{x} \in T_x \mathcal{S}^n$ (4.7). In particular, for all i , $\dot{x} = (jx_i)e_i$ is a tangent vector at x (where $j = \sqrt{-1}$ and e_i is the i th canonical basis vector), and it follows that

$$\langle (jx_i)e_i, (jx_i)Se_i \rangle = S_{ii} = (Cx)_i \bar{x}_i - C_{ii} \geq 0.$$

Since $C_{ii} = 1$ by construction,⁶ it follows in particular that $(Cx)_i \bar{x}_i$ is real and positive. Thus:

$$(Cx)_i \bar{x}_i = |(Cx)_i \bar{x}_i| = |(Cx)_i|.$$

We can now turn to lower-bounding the eigenvalues of S . Since $Sx = 0$, it is sufficient to study u^*Su for $u \in \mathbb{C}^n$ orthogonal to x , that is, $u^*x = 0$. Without loss

⁵ A second-order critical point satisfies first- and second-order necessary optimality conditions, namely, the gradient is zero and the Hessian is positive semidefinite (if minimizing) [26, Sect. 3.2.1].

⁶ As before, S is independent of $\text{diag}(C)$. Assuming $C_{ii} = 1$ involves no loss of generality.

of generality, fix the global phase of x such that $z^*x = |z^*x|$. Using once more that $C = zz^* + \sigma W$,

$$\begin{aligned} u^*Su &= u^* \text{ddiag}(Cxx^*)u - u^*Cu \\ &= \sum_{i=1}^n |u_i|^2 |(Cx)_i| - u^*Cu \\ &= \sum_{i=1}^n |u_i|^2 (z^*x)z_i + \sigma (Wx)_i - |u^*z|^2 - \sigma u^*Wu \\ &\geq |z^*x| \|u\|_2^2 - \sigma \sum_{i=1}^n |u_i|^2 |(Wx)_i| - |u^*(z-x)|^2 - \sigma \|W\|_{\text{op}} \|u\|_2^2 \\ &\geq \|u\|_2^2 \left(|z^*x| - \sigma \frac{\sum_{i=1}^n |u_i|^2 |(Wx)_i|}{\sum_{i=1}^n |u_i|^2} - \|z-x\|_2^2 - \sigma \|W\|_{\text{op}} \right). \end{aligned}$$

The fraction is a weighted mean of nonnegative numbers, $|(Wx)_i|$ for $i = 1 \dots n$, so that it is bounded by $\|Wx\|_\infty$. Now using z -discordance of W , invoke Lemma 4.1 to bound $|z^*x| \geq n - 72\sigma^2$, $\|z-x\|_2^2 \leq 144\sigma^2$, and $\|W\|_{\text{op}} \leq 3n^{1/2}$ to establish

$$u^*Su \geq \|u\|_2^2 (n - 216\sigma^2 - \sigma \|Wx\|_\infty - 3\sigma n^{1/2}).$$

As a result, a sufficient condition for S to be positive semidefinite with rank $n - 1$ is to have

$$n > \sigma \left(216\sigma + 3n^{1/2} + \|Wx\|_\infty \right). \tag{4.10}$$

At this point, notice that if $\|Wx\|_\infty$ were bounded by $\tilde{O}(\sqrt{n})$, we would indeed obtain the targeted rate for σ . Unfortunately, we could not obtain such a bound (see Remark 4.6). Instead, we carry on with the following suboptimal argument:

$$\begin{aligned} \|Wx\|_\infty &\leq \|W(x-z)\|_\infty + \|Wz\|_\infty \\ &\leq \|W(x-z)\|_2 + \|Wz\|_\infty \\ &\leq \|W\|_{\text{op}} \|x-z\|_2 + \|Wz\|_\infty. \end{aligned}$$

By z -discordance, $\|W\|_{\text{op}} \leq 3\sqrt{n}$ and $\|Wz\|_\infty \leq 3\sqrt{n \log n}$. By Lemma 4.1, $\|x-z\|_2 \leq 12\sigma$. Thus,

$$\|Wx\|_\infty \leq 36\sigma\sqrt{n} + 3\sqrt{n \log n}. \tag{4.11}$$

Combining with (4.10) yields the following sufficient condition:

$$\sqrt{n} > 3\sigma \left(72\sigma n^{-1/2} + 1 + 12\sigma + \sqrt{\log n} \right). \tag{4.12}$$

This condition involves only σ and n , and as such is the best result we prove. The bottleneck is the term in σ^2 , which leads to a bound of the type $\sigma \leq c \cdot n^{1/4}$. Further inspection of (4.12) shows $\sigma \leq \frac{1}{18}n^{1/4}$ is indeed a sufficient condition, for $n \geq 2$. This concludes the proof of Lemma 3.2.

Remark 4.6 (About $\|Wx\|_\infty$) Each entry $(Wx)_i$ is a sum of $n - 1$ random variables in the complex plane. If W_{ij} are standard (complex) Gaussians (so that their phases are uniformly distributed) and if x were independent from W , then the quantity $|(Wx)_i|$ would grow as $\tilde{O}(n^{1/2})$. The largest one, $\|Wx\|_\infty$, would grow similarly (as is the case for $\|Wz\|_\infty$, where we do have independence). Thus, we would recover the observed rate for σ in Fig. 2, namely, $\sigma \leq \frac{n^{1/2}}{\text{polylog}(n)}$ would suffice to ensure tightness. Likewise, the proof in Lemma 4.2 would lead to $\|x - z\|_\infty = \tilde{O}(\sigma/\sqrt{n})$. Unfortunately, W and x are not independent: the MLE depends on W in a complicated fashion.

One avenue to improve the current bound might be to expand $x = x(\sigma W)$ around 0, to first order. Some computations (omitted here) show that such a development would lead to $x = z + \frac{\sigma}{n} \mathfrak{S} \{ \text{ddiag}(Wzz^*) \} (iz) + \epsilon$. Then,

$$\|Wx\|_\infty \lesssim \|Wz\|_\infty + \frac{\sigma}{n} \|W \mathfrak{S} \{ \text{ddiag}(Wzz^*) \} z\|_\infty.$$

Still assuming Gaussian entries, the last norm appears to grow like $\tilde{O}(n)$, which would lead to $\|Wx\|_\infty = \tilde{O}(\sqrt{n} + \sigma)$, as desired. Unfortunately, controlling the error term ϵ and its effect on the above bound seems nontrivial, as this involves assessing smoothness of the argmax function $x(\sigma W)$.

5 Conclusions

Computing maximum likelihood estimators is important for statistical problems, in particular when exact signal recovery is out of reach. In this paper, we have made partial progress toward understanding when this computation can be executed in polynomial time for the angular synchronization problem. It still remains to improve the proposed result to understand how this MLE can be computed for noise levels as high as those witnessed empirically. As it stands, the suboptimal bound (4.11) is the bottleneck. Reducing it to $\tilde{O}(\sqrt{n})$ would yield the (tentatively) correct rate, as per Remark 4.6.

For many other problems of interest, it has likewise been observed that semidefinite relaxations tend to be tight. Notably, this is the case for synchronization of rotations: a problem similar to angular synchronization but with rotations in \mathbb{R}^d rather than \mathbb{R}^2 . Analyzing the latter might prove more difficult, because rotations in \mathbb{R}^d do not commute for $d > 2$. Even for $d = 2$, the case of an incomplete graph of (weighted) measurements is of interest, and not covered here. Tightness has also been observed for the Procrustes problem and the multi-reference alignment (MRA) problem [13]. The situation for MRA differs markedly from the present setting, as the noise originates at the vertices of a graph, not on the edges. Notwithstanding those differences, we hope that the present work will help future attempts to explain the tightness phenomenon in various settings.

Acknowledgements A. S. Bandeira was supported by AFOSR Grant No. FA9550-12-1-0317. Most of this work was done while he was with the Program for Applied and Computational Mathematics at Princeton University, and some while he was with the Department of Mathematics at the Massachusetts Institute of Technology. N. Boumal was supported by a Belgian F.R.S.-FNRS fellowship while working at the Université catholique de Louvain (Belgium), by a Research in Paris fellowship at Inria and ENS, the “Fonds Spéciaux de Recherche” (FSR UCLouvain), the Chaire Havas “Chaire Economie et gestion des nouvelles données” and the ERC Starting Grant SIPA. A. Singer was partially supported by Award Number R01GM090200 from the NIGMS, by Award Numbers FA9550-12-1-0317 and FA9550-13-1-0076 from AFOSR, by Award Number LTR DTD 06-05-2012 from the Simons Foundation, and by the Moore Foundation.

Appendix: Wigner matrices are discordant

This appendix is a proof for Proposition 3.3, namely, that for arbitrary $z \in \mathbb{C}^n$ such that $|z_1| = \dots = |z_n| = 1$, complex Wigner matrices are z -discordant (Definition 3.1) with high probability.

A matrix W is z -discordant if and only if $\text{diag}(z)^* W \text{diag}(z)$ is $\mathbb{1}$ -discordant. Since $\text{diag}(z)^* W \text{diag}(z)$ has the same distribution as W (owing to complex normal random variables having uniformly random phase), we may without loss of generality assume $z = \mathbb{1}$ in the remainder of the proof.

1. $\Pr \{ \|W\|_{\text{op}} > 3n^{1/2} \} \leq e^{-n/2}$.

Although tail bounds for the real version of this are well-known (see for example [15,56]) and they mostly hold verbatim in the complex case, for the sake of completeness we include a classical argument, based on Slepian’s comparison theorem and Gaussian concentration, for a tail bound in the complex valued case.

We will bound the largest eigenvalue of W . It is clear that a simple union bound argument will allow us to bound also the smallest, and thus bound the largest in magnitude. Let $\lambda_+ = \max_{v \in \mathbb{C}^n: \|v\|=1} v^* W v$ denote the largest eigenvalue of W . For any unit-norm $u, v \in \mathbb{C}^n$, the real valued Gaussian process $X_v = v^* W v$ satisfies:

$$\begin{aligned} \mathbb{E} (X_v - X_u)^2 &= \mathbb{E} \left(\sum_{i < j} W_{ij} (\bar{v}_i v_j - \bar{u}_i u_j) + W_{ji} (\bar{v}_j v_i - \bar{u}_j u_i) \right)^2 \\ &= \sum_{i < j} \mathbb{E} [W_{ij} (\bar{v}_i v_j - \bar{u}_i u_j) + W_{ji} (\bar{v}_j v_i - \bar{u}_j u_i)]^2. \end{aligned}$$

The variable W_{ij} has uniformly random phase, hence so does W_{ij}^2 , so that $\mathbb{E} W_{ij}^2 = 0$. As a result,

$$\begin{aligned} \mathbb{E} (X_v - X_u)^2 &= \sum_{i < j} 2\mathbb{E} |W_{ij}|^2 |\bar{v}_i v_j - \bar{u}_i u_j|^2 \\ &= 2 \sum_{i < j} |\bar{v}_i v_j - \bar{u}_i u_j|^2 \\ &\leq \sum_{i, j} |\bar{v}_i v_j - \bar{u}_i u_j|^2. \end{aligned}$$

Note that, since $\|u\| = \|v\| = 1$,

$$\begin{aligned} \sum_{i,j} |\overline{v_i} v_j - \overline{u_i} u_j|^2 &= \sum_{ij} \left[|v_i|^2 |v_j|^2 + |u_i|^2 |u_j|^2 - \overline{v_i} v_j u_i \overline{u_j} - v_i \overline{v_j} u_i \overline{u_j} \right] \\ &= 2 - 2 |v^* u|^2 \\ &\leq 2 (2 - 2 |v^* u|) \\ &\leq 4 (1 - \Re [v^* u]) \\ &= 2 \|v - u\|^2. \end{aligned}$$

This means that we can use Slepian’s comparison theorem (see for example [39, Cor. 3.12]) to get

$$\mathbb{E} \lambda_+ \leq \sqrt{2} \mathbb{E} \max_{\tilde{v} \in \mathbb{R}^{2n}: \|\tilde{v}\|=1} \tilde{v}^T g \leq 2\sqrt{n}, \tag{5.1}$$

where g is a standard Gaussian vector in \mathbb{R}^{2n} and \tilde{v} is a vector in \mathbb{R}^{2n} obtained from $v \in \mathbb{C}^n$ by stacking its real and imaginary parts. Since $|\|W_1\| - \|W_2\|| \leq \|W_1 - W_2\| \leq \|W_1 - W_2\|_F$, Gaussian concentration [39] gives

$$\Pr \{ \lambda_+ - \mathbb{E} \lambda_+ \geq t \} \leq e^{-t^2/2}. \tag{5.2}$$

Using (5.1) and (5.2) gives

$$\Pr \left\{ \|W\|_{\text{op}} > 3n^{1/2} \right\} \leq e^{-n/2}.$$

2. $\Pr \{ \|W\mathbb{1}\|_{\infty} > 3\sqrt{n \log n} \} \leq 2n^{-5/4}.$

The random vector given by $\frac{1}{(n-1)^{1/2}} W\mathbb{1}$ is jointly Gaussian where the marginal of each entry is a standard complex Gaussian. By a suboptimal union bound argument, the maximum absolute value among k standard complex Gaussian random variables (not necessarily independent) is larger than t with probability at most $2ke^{-t^2/4}$. Hence,

$$\begin{aligned} \Pr \left\{ \|W\mathbb{1}\|_{\infty} > 3\sqrt{n \log n} \right\} &\leq \Pr \left\{ \left\| \frac{1}{(n-1)^{1/2}} W\mathbb{1} \right\|_{\infty} > 3\sqrt{\log n} \right\} \\ &\leq 2ne^{-\frac{9}{4} \log n} = 2n^{-5/4}. \end{aligned}$$

To support the discussion following Proposition 3.3, we further argue that

$$\Pr \left\{ |\mathbb{1}^* W\mathbb{1}| > n^{3/2} \right\} \leq e^{-n/2}.$$

It is easy to see that $\mathbb{1}^* W\mathbb{1}$ is a real Gaussian random variable with zero mean and variance $2\frac{n(n-1)}{2} = n(n-1)$. This implies that:

$$\Pr \left\{ \|\mathbb{1}^* W \mathbb{1}\| > n^{3/2} \right\} \leq \Pr \left\{ \frac{1}{(n(n-1))^{1/2}} \|\mathbb{1}^* W \mathbb{1}\| > n^{1/2} \right\} \\ \leq \frac{1}{\sqrt{2\pi}} \frac{1}{n^{1/2}} e^{-\frac{n}{2}} \leq e^{-n/2}.$$

References

1. Abbe, E., Bandeira, A.S., Bracher, A., Singer, A.: Decoding binary node labels from censored edge measurements: phase transition and efficient recovery. *IEEE Trans. Network Sci. Eng.* **1**(1), 10–22 (2014)
2. Abbe, E., Bandeira, A.S., Hall, G.: Exact recovery in the stochastic block model. *IEEE Trans. Inf. Theory* **62**(1), 471–487 (2015)
3. Absil, P.-A., Mahony, R., Sepulchre, R.: *Optimization Algorithms on Matrix*. Princeton University Press, Princeton (2008)
4. Agrawal, A., Raskar, R., Chellappa, R.: What is the range of surface reconstructions from a gradient field? In: Leonardis, A., Bischof, H., Pinz, A. (eds.) *Computer Vision—ECCV 2006*. Lecture Notes in Computer Science, vol. 3951, pp. 578–591. Springer, Berlin Heidelberg (2006)
5. Alexeev, B., Bandeira, A.S., Fickus, M., Mixon, D.G.: Phase retrieval with polarization. *SIAM J. on Imaging Sci.* **7**(1), 35–66 (2013)
6. Alizadeh, F., Haeberly, J.-P., Overton, M.L.: Complementarity and nondegeneracy in semidefinite programming. *Math. Program.* **77**(1), 111–128 (1997)
7. Amelunxen, D., Lotz, M., McCoy, M.B., Tropp, J.A.: Living on the edge: phase transitions in convex programs with random data. *Inf. Inference* **3**(3), 224–294 (2014)
8. Ames, B.P.W.: Guaranteed clustering and biclustering via semidefinite programming. *Math. Program.* **147**, 429–465 (2014)
9. Bandeira, A.S., Chen, Y., Mixon, D.G.: Phase retrieval from power spectra of masked signals. *Inf. Inference J. IMA* **3**, 83–102 (2014)
10. Bandeira, A.S., Kennedy, C., Singer, A.: Approximating the little Grothendieck problem over the orthogonal group. *Math. Program. Ser. A*, 1–43 (2016). doi:[10.1007/s10107-016-0993-7](https://doi.org/10.1007/s10107-016-0993-7)
11. Bandeira, A.S., Singer, A., Spielman, D.A.: A Cheeger inequality for the graph connection Laplacian. *SIAM J. Matrix Anal. Appl.* **34**(4), 1611–1630 (2013)
12. Bandeira, A.S.: Random Laplacian matrices and convex relaxations. (2015). [arXiv:1504.03987](https://arxiv.org/abs/1504.03987)
13. Bandeira, A.S., Charikar, M., Singer, A., Zhu, A.: Multireference alignment using semidefinite programming. In: *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pp. 459–470. ACM, (2014)
14. Bandeira, A.S., Khoo, Y., and Singer, A.: Open problem: Tightness of maximum likelihood semidefinite relaxations. In: Maria Florina B., Vitaly F., Csaba S., (ed.), *Proceedings of the 27th Conference on Learning Theory*, vol. 35 of *JMLR W&CP*, pp. 1265–1267, (2014)
15. Bandeira, A.S., van Handel, R.: Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *Annal. Probab.*, (to appear)
16. Barvinok, A.I.: Problems of distance geometry and convex properties of quadratic maps. *Discrete Comput. Geom.* **13**(1), 189–202 (1995)
17. Boumal, N.: A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints. (2015). [arXiv:1506.00575](https://arxiv.org/abs/1506.00575)
18. Boumal, N., Singer, A., Absil, P.-A., Blondel, V.D.: Cramér-Rao bounds for synchronization of rotations. *Inf. Inference* **3**, 1–39 (2014)
19. Briët, J., de Oliveira Filho, F.M., Vallentin, F.: Grothendieck inequalities for semidefinite programs with rank constraint. *Theory Comput.* **10**(4), 77–105 (2014)
20. Candès, E.J., Recht, B.: Exact matrix completion via convex optimization. *Found. Comput. Math.* **9**(6), 717–772 (2009)
21. Candès, E.J., Romberg, J., Tao, T.: Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory* **52**, 489–509 (2006)
22. Candès, E.J., Romberg, J., Tao, T.: Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.* **59**, 1207–1223 (2006)

23. Candès, E.J., Strohmer, T., Vershynina, V.: Phaselift: exact and stable signal recovery from magnitude measurements via convex programming. *Commun. Pure Appl. Math.* **66**, 1241–1274 (2011)
24. Chandrasekaran, V., Parrilo, P.A., Willsky, A.S.: Latent variable graphical model selection via convex optimization. *Annal. Stat.* **40**, 1935–1967 (2012)
25. Chandrasekaran, V., Recht, B., Parrilo, P.A., Willsky, A.S.: The convex geometry of linear inverse problems. *Found. Comput. Math.* **12**(6), 805–849 (2012)
26. Conn, A.R., Gould, N.I.M., Toint, P.L.: Trust-region methods. MOS-SIAM Series on Optimization, Philadelphia. doi:10.1137/1.9780898719857 (2000)
27. Cucuringu, M.: Sync-Rank: Robust ranking, constrained ranking and rank aggregation via eigenvector and semidefinite programming synchronization. *IEEE Trans. Netw. Sci.* **3**(1), 58–79 (2015)
28. Demanet, L., Hand, P.: Stable optimizationless recovery from phaseless linear measurements. *J. Fourier Anal. Appl.* **20**(1), 199–221. doi:10.1007/s00041-013-9305-2
29. Demanet, L., Jugnon, V.: Convex recovery from interferometric measurements. (2013). [arXiv:1307.6864](https://arxiv.org/abs/1307.6864)
30. Ding, X., Jiang, T.: Spectral distributions of adjacency and Laplacian matrices of random graphs. *Annal. Appl. Probab.* **20**(6), 2086–2117 (2010)
31. Donoho, D.L.: Compressed sensing. *IEEE Trans. Inf. Theory* **52**, 1289–1306 (2006)
32. Giridhar, A., Kumar, P.R.: Distributed clock synchronization over wireless networks: Algorithms and analysis. In: *IEEE Conf. Decis. Control*, 2006 45th, pp. 4915–4920. IEEE, (2006)
33. Goemans, M.X., Williamson, D.P.: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM (JACM)* **42**(6), 1115–1145 (1995)
34. Goemans, M.X., Williamson, D.P.: Approximation algorithms for Max-3-Cut and other problems via complex semidefinite programming. *J. Comput. Syst. Sci.* **68**(2), 442–470 (2004)
35. Grothendieck, A.: Resume de la theorie metrique des produits tensoriels topologiques (french). Reprint of *Bol. Soc. Mat. Sao Paulo*, p. 179, (1996)
36. Hartley, R., Trunpf, J., Dai, Y., Li, H.: Rotation averaging. *Int. J. Comput. Vis.* **103**(3), 267–305 (2013)
37. Huang, Q.X., Guibas, L.: Consistent shape maps via semidefinite programming. In: *Computer Graphics Forum*. vol. 32, pp. 177–186. Wiley Online Library, (2013)
38. Journée, M., Bach, F., Absil, P.-A., Sepulchre, R.: Low-rank optimization on the cone of positive semidefinite matrices. *SIAM J. Optim.* **20**(5), 2327–2351 (2010)
39. Ledoux, M., Talagrand, M.: Probability in Banach Spaces: isoperimetry and processes, vol. 23, *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge / A Series of Modern Surveys in Mathematics*. (1991)
40. Low, S.H.: Convex relaxation of optimal power flow: a tutorial. In: *Bulk Power System Dynamics and Control-IX Optimization. Security and Control of the Emerging Power Grid (IREP)*, 2013 IREP Symposium, pp. 1–15. IEEE, (2013)
41. Luo, Z., Ma, W., So, A.M.C., Ye, Y., Zhang, S.: Semidefinite relaxation of quadratic optimization problems. *Sign. Process. Mag. IEEE* **27**(3), 20–34 (2010)
42. Martinec, D., Pajdla, T.: Robust rotation and translation estimation in multiview reconstruction. In: *IEEE Conference on Computer Vision and Pattern Recognition*, 2007. CVPR '07. pp. 1–8, June (2007)
43. Pataki, G.: On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Math. Operations Res.* **23**(2), 339–358 (1998)
44. Pisier, G.: Grothendieck's theorem, past and present. *Bull. Amer. Math. Soc.* **49**, 237–323 (2011)
45. Rubinstein, J., Wolansky, G.: Reconstruction of optical surfaces from ray data. *Opt. Rev.* **8**(4), 281–283 (2001)
46. Ruszczyński, A.P.: *Nonlinear Optimization*. Princeton University Press, Princeton (2006)
47. Sagnol, G.: A class of semidefinite programs with rank-one solutions. *Linear Algebra Appl.* **435**(6), 1446–1463 (2011)
48. Shapiro, A.: Rank-reducibility of a symmetric matrix and sampling theory of minimum trace factor analysis. *Psychometrika* **47**(2), 187–199 (1982)
49. Singer, A.: Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmonic Anal.* **30**(1), 20–36 (2011)
50. Singer, A., Shkolnisky, Y.: Three-dimensional structure determination from common lines in Cryo-EM by eigenvectors and semidefinite programming. *SIAM J. Imaging Sci.* **4**(2), 543–572 (2011)
51. So, A.M.-C.: Probabilistic analysis of the semidefinite relaxation detector in digital communications. *Proc. SODA* (2010)

52. So, A.M.-C., Zhang, J., Ye, Y.: On approximating complex quadratic optimization problems via semidefinite programming relaxations. *Math. Program.* **110**(1), 93–110 (2007)
53. Sojoudi, S., Lavaei, J.: Exactness of semidefinite relaxations for nonlinear optimization problems with underlying graph structure. *SIAM J. Optim.* **24**(4), 1746–1778 (2014)
54. Tropp, J.A.: Just relax: convex programming methods for identifying sparse signals in noise. *IEEE Trans Inf. Theory* **52**(3), 1030–1051 (2006)
55. Vanderberghe, L., Boyd, S.: Semidefinite programming. *SIAM Rev.* **38**, 49–95 (1996)
56. Vershynin, R.: Introduction to the non-asymptotic analysis of random matrices. In: Eldar, Y., Kutyniok, G. (eds.) Chapter 5 of: *Compressed Sensing, Theory and Applications*. Cambridge University Press, Cambridge (2012)
57. Wang, L., Singer, A.: Exact and stable recovery of rotations for robust synchronization. *Inf. Inference* **2**(2), 145–193 (2013)
58. Williamson, D.P., Shmoys, D.B.: *The Design of Approximation Algorithms*. Cambridge University Press, Cambridge (2011)
59. Zhang, S., Huang, Y.: Complex quadratic optimization and semidefinite programming. *SIAM J. Optim.* **16**(3), 871–890 (2006)
60. Zhang, T., and Singer, A.: Disentangling two orthogonal matrices. (2015). [arXiv:1506.02217](https://arxiv.org/abs/1506.02217)