PhaseMax: Convex Phase Retrieval via Basis Pursuit

Journal Article

Author(s):
Goldstein, Tom; Studer, Christoph (iD)

PhaseMax: Convex Phase Retrieval via Basis Pursuit

Tom Goldstein and Christoph Studer

Abstract—We consider the recovery of a (real- or complex-valued) signal from magnitude-only measurements, known as phase retrieval. We formulate phase retrieval as a convex optimization problem, which we call PhaseMax. Unlike other convex methods that use semidefinite relaxation and lift the phase retrieval problem to a higher dimension, PhaseMax is a "non-lifting" relaxation that operates in the original signal dimension. We show that the dual problem to PhaseMax is Basis Pursuit, which implies that phase retrieval can be performed using algorithms initially designed for sparse signal recovery. We develop sharp lower bounds on the success probability of PhaseMax for a broad range of random measurement ensembles, and we analyze the impact of measurement noise on the solution accuracy. We use numerical results to demonstrate the accuracy of our recovery guarantees, and we showcase the efficacy and limits of PhaseMax in practice.

I. INTRODUCTION

Phase retrieval is concerned with the recovery of an n-dimensional signal $\mathbf{x}^0 \in \mathcal{H}^n$, with $\mathcal{H}$ either $\mathbb{R}$ or $\mathbb{C}$, from $m \geq n$ squared-magnitude, noisy measurements [2]

$$b_i^2 = |\langle \mathbf{a}_i, \mathbf{x}^0 \rangle|^2 + \eta_i, \quad i = 1, 2, \ldots, m, \tag{1}$$

where $\mathbf{a}_i \in \mathcal{H}^n$, $i = 1, 2, \ldots, m$, are the (known) measurement vectors and $\eta_i \in \mathbb{R}$, $i = 1, 2, \ldots, m$, models measurement noise. Let $\hat{\mathbf{x}} \in \mathcal{H}^n$ be an approximation vector[1] to the true signal $\mathbf{x}^0$. We recover the signal $\mathbf{x}^0$ by solving the following convex problem we call *PhaseMax*:

$$(\text{PM}) \quad \begin{cases} \underset{\mathbf{x} \in \mathcal{H}^n}{\text{maximize}} & \langle \mathbf{x}, \hat{\mathbf{x}} \rangle_\Re \\ \text{subject to} & |\langle \mathbf{a}_i, \mathbf{x} \rangle| \leq b_i, \quad i = 1, 2, \ldots, m. \end{cases}$$

Here, $\langle \mathbf{x}, \hat{\mathbf{x}} \rangle_\Re$ denotes the real-part of the inner product between the vectors $\mathbf{x}$ and $\hat{\mathbf{x}}$. The main idea behind PhaseMax is to find the vector $\mathbf{x}$ that is most aligned with the approximation vector $\hat{\mathbf{x}}$ and satisfies a convex relaxation of the measurement constraints in (1).

Our main goal is to develop sharp lower bounds on the probability with which PhaseMax succeeds in recovering the true signal $\mathbf{x}^0$, up to an arbitrary phase ambiguity that does not affect the measurement constraints in (1). By assuming noiseless measurements, one of our main results is as follows.

T. Goldstein is with the Department of Computer Science, University of Maryland, College Park, MD (e-mail: tomg@cs.umd.edu).

C. Studer is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY (e-mail: studer@cornell.edu).

[1]Approximation vectors can be obtained via a variety of algorithms, or can even be chosen at random. See Section VI.

Theorem 1. *Consider the case of recovering a complex-valued signal $\mathbf{x} \in \mathbb{C}^n$ from m noiseless measurements of the form (1) with measurement vectors $\mathbf{a}_i$, $i = 1, 2, \ldots, m$, sampled independently and uniformly from the unit sphere. Let*

$$\text{angle}(\mathbf{x}^0, \hat{\mathbf{x}}) = \arccos\left(\frac{\langle \mathbf{x}^0, \hat{\mathbf{x}} \rangle_\Re}{\|\mathbf{x}^0\|_2 \|\hat{\mathbf{x}}\|_2} \right)$$

be the angle between the true vector $\mathbf{x}^0$ and the approximation $\hat{\mathbf{x}}$, and define the constant

$$\alpha = 1 - \tfrac{2}{\pi} \text{angle}(\mathbf{x}^0, \hat{\mathbf{x}})$$

that measures the approximation accuracy. Then, the probability that PhaseMax recovers the true signal $\mathbf{x}^0$, denoted by $p_\mathbb{C}(m, n)$, is bounded from below as follows:

$$p_\mathbb{C}(m, n) \geq 1 - \exp\left(-\frac{(\alpha m - 4n)^2}{2m} \right) \tag{2}$$

whenever $\alpha m > 4n$.

In words, if $m > 4n/\alpha$ and $\alpha > 0$, then PhaseMax will succeed with non-zero probability. Furthermore, for a fixed signal dimension n and an arbitrary approximation vector $\hat{\mathbf{x}}$ that satisfies $\text{angle}(\mathbf{x}^0, \hat{\mathbf{x}}) < \frac{\pi}{2}$, i.e., one that is not orthogonal to the vector $\mathbf{x}^0$, we can make the success probability of PhaseMax arbitrarily close to one by increasing the number of measurements m. As we shall see, our recovery guarantees are sharp and accurately predict the performance of PhaseMax in practice.

We emphasize that the convex formulation (PM) has been studied by several other authors, including Bahmani and Romberg [3], whose work appeared shortly before our own. Related work will be discussed in detail in Section I-C.

A. Convex Phase Retrieval via Basis Pursuit

It is quite intriguing that the following Basis Pursuit problem [4], [5]

$$(\text{BP}) \quad \begin{cases} \underset{\mathbf{z} \in \mathcal{H}^m}{\text{minimize}} & \|\mathbf{z}\|_1 \\ \text{subject to} & \hat{\mathbf{x}} = \mathbf{A}\mathbf{B}^{-1}\mathbf{z}, \end{cases}$$

with $\mathbf{B} = \text{diag}(b_1, b_2, \ldots, b_m)$ and $\mathbf{A} = [\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_m]$ is the dual problem to (PM); see, e.g., [6, Lem. 1]. As a consequence, if PhaseMax succeeds, then the phases of the solution vector $\mathbf{z} \in \mathcal{H}^m$ to (BP) are exactly the phases that were lost in the measurement process in (1), i.e., we have

$$y_i = \text{phase}(z_i)b_i = \langle \mathbf{a}_i, \mathbf{x}^0 \rangle, \quad i = 1, 2, \ldots, m,$$

with $\text{phase}(z) = z/|z|$ for $z \neq 0$ and $\text{phase}(0) = 1$. This observation not only reveals a fundamental connection between phase retrieval and sparse signal recovery, but also implies that

Basis Pursuit solvers can be used to recover the signal from the phase-less measurements in (1).

B. A Brief History of Phase Retrieval

Phase retrieval is a well-studied problem with a long history [7], [8] and enjoys widespread use in applications such as X-ray crystallography [9]–[11], microscopy [12], [13], imaging [14], and many more [15]–[18]. Early algorithms, such as the Gerchberg-Saxton [7] or Fienup [8] algorithms, rely on alternating projection to recover complex-valued signals from magnitude-only measurements. The papers [2], [19], [20] sparked new interest in the phase retrieval problem by showing that it can be relaxed to a semidefinite program. Prominent convex relaxations include PhaseLift [2] and PhaseCut [21]. These methods come with recovery guarantees but require the problem to be lifted to a higher dimensional space, which prevents their use for large-scale problems. Several authors recently addressed this problem by presenting efficient methods that can solve large-scale semidefinite programs without the complexity of representing the full-scale (or lifted) matrix of unknowns. These approaches include gauge duality methods [22], [23] and sketching methods [24].

More recently, a number of *non-convex* algorithms have been proposed (see e.g., [25]–[31]) that directly operate in the original signal dimension and exhibit excellent empirical performance. The algorithms in [25], [27]–[29], [32] come with recovery guarantees that mainly rely on accurate initializers, such as the (truncated) spectral initializer [25], [28], the Null initializer [33], the orthogonality-promoting method [29], or more recent initializers that guarantee optimal performance [34], [35] (see Section VI for additional details). These initializers enable non-convex phase retrieval algorithms to succeed, given a sufficiently large number of measurements; see [36] for more details on the geometry of such non-convex problems.

C. Related Work on Non-Lifting Convex Methods

Because of the extreme computational burden of traditional convex relaxations (which require "lifting" to a higher dimension), a number of authors have recently been interested in non-lifting convex relaxations for phase retrieval. Shortly before the appearance of this work, the formulation (PM) was studied by Bahmani and Romberg [3]. Using methods from machine learning theory, the authors derived bounds on the recovery of signals with and without noise. The results in [3] are stronger than those presented here in that they are uniform with respect to the approximation vector $\hat{\mathbf{x}}$, but weaker in the sense that they require significantly more measurements.

Several authors have studied PhaseMax after the initial appearance of this work. An alternative proof of accurate signal recovery was derived using measure concentration bounds in [37]. Compressed sensing methods for sparse signals [38], and corruption-robust methods for noisy signals [39] have also been presented. The closely related non-lifting relaxation BranchHull has also been proposed for recovering signals from entry-wise product measurements [40].

Finally, we note that recent works have proved tight asymptotic bounds for PhaseMax in the asymptotic limit, i.e.,

where $\beta = m/n$ is constant and $m \to \infty$. The authors of [41] derive an exact asymptotic bound of the performance of PhaseMax, and a related non-rigorous analysis was given in [42]. It was also shown that better signal recovery guarantees can be obtained by iteratively applying PhaseMax. The resulting method, called PhaseLamp, was analyzed in [41].

D. Contributions and Paper Outline

In contrast to algorithms relying on semidefinite relaxation or non-convex problem formulations, we propose *PhaseMax*, a novel, convex method for phase retrieval that directly operates in the original signal dimension. In Section II, we establish a deterministic condition that guarantees uniqueness of the solution to the (PM) problem. We borrow methods from geometric probability to derive sharp lower bounds on the success probability for real- and complex-valued systems in Section III. Section V generalizes our results to a broader range of random measurement ensembles and to systems with measurement noise. We show in Section VI that randomly chosen approximation vectors are sufficient to ensure faithful recovery, given a sufficiently large number of measurements. We numerically demonstrate the sharpness of our recovery guarantees and showcase the practical limits of PhaseMax in Section VII. We conclude in Section VIII.

E. Notation

Lowercase and uppercase boldface letters stand for column vectors and matrices, respectively. For a complex-valued matrix $\mathbf{A}$, we denote its transpose and Hermitian transpose by $\mathbf{A}^T$ and $\mathbf{A}^*$, respectively; the real and imaginary parts are $\mathbf{A}_\Re$ and $\mathbf{A}_\Im$. The ith column of the matrix $\mathbf{A}$ is denoted by $\mathbf{a}_i$ and the kth entry of the ith vector $\mathbf{a}_i$ is $[\mathbf{a}_i]_k$; for a vector $\mathbf{a}$ without index, we simply denote the kth entry by a_k. We define the inner product between two complex-valued vectors $\mathbf{a}$ and $\mathbf{b}$ as $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a}^* \mathbf{b}$. We use j to denote the imaginary unit. The ℓ_2-norm and ℓ_1-norm of the vector $\mathbf{a}$ are $\|\mathbf{a}\|_2$ and $\|\mathbf{a}\|_1$, respectively.

II. UNIQUENESS CONDITION

The measurement constraints in (1) do not uniquely define a vector. If $\mathbf{x}$ is a vector that satisfies (1), then any vector $\mathbf{x}' = e^{j\phi}\mathbf{x}$ for $\phi \in [0, 2\pi)$ also satisfies the constraints. In contrast, if $\mathbf{x}$ is a solution to (PM), then $e^{j\phi}\mathbf{x}$ with $\phi \neq 0$ will *not* be another solution. In fact, consider any vector $\mathbf{x}$ in the feasible set of (PM) with $\langle \mathbf{x}, \hat{\mathbf{x}} \rangle_\Im \neq 0$. By choosing $\omega = \text{phase}(\langle \mathbf{x}, \hat{\mathbf{x}} \rangle)$, we have

$$\langle \omega \mathbf{x}, \hat{\mathbf{x}} \rangle_\Re = |\langle \mathbf{x}, \hat{\mathbf{x}} \rangle| > \langle \mathbf{x}, \hat{\mathbf{x}} \rangle_\Re,$$

which implies that given such a vector $\mathbf{x}$, one can always increase the objective function of (PM) simply by *aligning* $\mathbf{x}$ with the approximation $\hat{\mathbf{x}}$ (i.e., modifying its phase so that $\langle \mathbf{x}, \hat{\mathbf{x}} \rangle$ is real valued). The following definition makes this observation rigorous.

Definition 1. *A vector* $\mathbf{x}$ *is said to be* aligned *with another vector* $\hat{\mathbf{x}}$, *if the inner product* $\langle \mathbf{x}, \hat{\mathbf{x}} \rangle$ *is real-valued and non-negative.*

From all the vectors that satisfy the measurement constraints in (1), there is only one that is a candidate solution to the convex problem (PM), which is also the solution that is aligned with $\hat{\mathbf{x}}$. For this reason, we adopt the following important convention throughout the rest of this paper.

> The true vector $\mathbf{x}^0$ denotes a solution to (1) that is aligned with the approximation vector $\hat{\mathbf{x}}$.

Remark 1. *There is an interesting relation between the convex formulation of PhaseMax and the semidefinite relaxation method PhaseLift [2], [19], [20]. Recall that the set of solutions to any convex problem is always convex. However, the solution set of the measurement constraints (1) is invariant under phase rotations, and thus non-convex. It is therefore impossible to design a convex problem that yields this set of solutions. PhaseMax and PhaseLift differ in how they remove the phase ambiguity from the problem to enable a convex formulation. Rather than trying to identify the true vector $\mathbf{x}^0$, PhaseLift reformulates the problem in terms of the quantity $\mathbf{x}^0(\mathbf{x}^0)^H$, which is unaffected by phase rotations in $\mathbf{x}^0$. Hence, PhaseLift removes the rotation symmetry from the solution set, thus yielding a problem with a convex set of solutions. PhaseMax does something much simpler: it* pins down *the phase of the solution to an arbitrary quantity, thus removing the phase ambiguity and restoring convexity to the solution set. This arbitrary phase choice is made when selecting the phase of the approximation $\hat{\mathbf{x}}$.*

We are now ready to state a deterministic condition under which PhaseMax succeeds in recovering the true vector $\mathbf{x}^0$. The result applies to the noiseless case, i.e., $\eta_i = 0$, $i = 1, 2, \ldots, m$. In this case, all inequality constraints in (PM) are active at $\mathbf{x}^0$. The noisy case will be discussed in Section V-B.

Theorem 2. *The true vector $\mathbf{x}^0$ is the unique maximizer of (PM) if, for any unit vector $\boldsymbol{\delta} \in \mathcal{H}^n$ that is aligned with the approximation $\hat{\mathbf{x}}$,*

$$\exists i, [\langle \mathbf{a}_i, \mathbf{x}^0 \rangle^* \langle \mathbf{a}_i, \boldsymbol{\delta} \rangle]_\Re > 0.$$

Proof. Suppose the conditions of this theorem hold, and consider some candidate solution $\mathbf{x}'$ in the feasible set for (PM) with $\langle \mathbf{x}', \hat{\mathbf{x}} \rangle \geq \langle \mathbf{x}^0, \hat{\mathbf{x}} \rangle$. Without loss of generality, we assume $\mathbf{x}'$ to be aligned with $\hat{\mathbf{x}}$. Then, the vector $\boldsymbol{\Delta} = \mathbf{x}' - \mathbf{x}^0$ is also aligned with $\hat{\mathbf{x}}$, and satisfies

$$\langle \boldsymbol{\Delta}, \hat{\mathbf{x}} \rangle = \langle \mathbf{x}', \hat{\mathbf{x}} \rangle - \langle \mathbf{x}^0, \hat{\mathbf{x}} \rangle \geq 0.$$

Since $\mathbf{x}'$ is a feasible solution for (PM), we have

$$|\langle \mathbf{a}_i, \mathbf{x}^0 + \boldsymbol{\Delta} \rangle|^2 =$$
$$|\langle \mathbf{a}_i, \mathbf{x}^0 \rangle|^2 + 2[\langle \mathbf{a}_i, \mathbf{x}^0 \rangle^* \langle \mathbf{a}_i, \boldsymbol{\Delta} \rangle]_\Re + |\langle \mathbf{a}_i, \boldsymbol{\Delta} \rangle|^2 \leq b_i^2, \ \forall i.$$

But $|\mathbf{a}_i^T \mathbf{x}^0|^2 = b_i^2$, and so

$$[\langle \mathbf{a}_i, \mathbf{x}^0 \rangle^* \langle \mathbf{a}_i, \boldsymbol{\Delta} \rangle]_\Re \leq -\tfrac{1}{2}|\langle \mathbf{a}_i, \boldsymbol{\Delta} \rangle|^2 \leq 0, \ \forall i.$$

Now, if $\|\boldsymbol{\Delta}\|_2 > 0$, then the unit-length vector $\boldsymbol{\delta} = \boldsymbol{\Delta}/\|\boldsymbol{\Delta}\|_2$ satisfies $[\langle \mathbf{a}_i, \mathbf{x}^0 \rangle^* \langle \mathbf{a}_i, \boldsymbol{\delta} \rangle]_\Re \leq 0$ for all i, which contradicts the hypothesis of the theorem. It follows that $\|\boldsymbol{\Delta}\|_2 = 0$ and

$\mathbf{x}' = \mathbf{x}^0$. $\qquad \square$

Theorem 2 has an intuitive geometrical interpretation. If $\mathbf{x}^0$ is an optimal point and $\boldsymbol{\delta}$ is an ascent direction, then one cannot move in the direction of $\boldsymbol{\delta}$ starting at $\mathbf{x}^0$ without leaving the feasible set. This condition is met if there is an $\mathbf{a}_i$ such that $\mathbf{x}^0$ and $\boldsymbol{\delta}$ both lie on the same side of the plane through the origin orthogonal to the measurement vector $\mathbf{a}_i$.

III. PRELIMINARIES: CLASSICAL SPHERE COVERING PROBLEMS AND GEOMETRIC PROBABILITY

To derive sharp conditions on the success probability of PhaseMax, we require a set of tools from geometric probability. Many classical problems in geometric probability involve calculating the likelihood of a sphere being covered by random "caps," or semi-spheres, which we define below.

Definition 2. *Consider the set $\mathcal{S}_\mathcal{H}^{n-1} = \{\mathbf{x} \in \mathcal{H}^n \mid \|\mathbf{x}\|_2 = 1\}$, the unit sphere embedded in $\mathcal{H}^n$. Given a vector $\mathbf{a} \in \mathcal{H}^n$, the cap centered at $\mathbf{a}$ with central angle θ is defined as*

$$\mathcal{C}_\mathcal{H}(\mathbf{a}, \theta) = \{\boldsymbol{\delta} \in \mathcal{S}_\mathcal{H}^{n-1} \mid \langle \mathbf{a}, \boldsymbol{\delta} \rangle_\Re > \cos(\theta)\}. \quad (3)$$

This cap contains all vectors that form an angle with $\mathbf{a}$ of less than θ radians. When $\theta = \pi/2$, we have a semisphere centered at $\mathbf{a}$, which is simply denoted by

$$\mathcal{C}_\mathcal{H}(\mathbf{a}) = \mathcal{C}_\mathcal{H}(\mathbf{a}, \pi/2) = \{\boldsymbol{\delta} \in \mathcal{S}_\mathcal{H}^{n-1} \mid \langle \mathbf{a}, \boldsymbol{\delta} \rangle_\Re > 0\}. \quad (4)$$

We say that a collection of caps *covers* the entire sphere if the sphere is contained in the union of the caps. Before we can say anything useful about when a collection of caps covers the sphere, we will need the following classical result, which is often attributed to Schläfli [43]. Proofs that use simple induction methods can be found in [44]–[46]. For completeness, we briefly include a short proof in Appendix A.

Lemma 1. *Consider a sphere $\mathcal{S}_\mathbb{R}^{n-1} \subset \mathbb{R}^n$. Suppose we slice the sphere with k planes through the origin. Suppose also that every subset of n planes have linearly independent normal vectors. These planes divide the sphere into*

$$r(n, k) = 2 \sum_{i=0}^{n-1} \binom{k-1}{i} \quad (5)$$

regions.

Classical results in geometric probability study the likelihood of a sphere being covered by random caps with centers chosen independently and uniformly from the sphere's surface. For our purposes, we need to study the more specific case in which caps are only chosen from a subset of the sphere. While calculating this probability is hard in general, it is quite simple when the set obeys the following symmetry condition.

Definition 3. *We say that the set $\mathcal{A}$ is* symmetric *if, for all $\mathbf{x} \in \mathcal{A}$, we also have $-\mathbf{x} \in \mathcal{A}$.*

We are now ready to prove a general result that states when the sphere is covered by random caps.

Lemma 2. *Consider some symmetric set $\mathcal{A} \subset \mathcal{S}_\mathbb{R}^{n-1}$ of positive ($n-1$ dimensional) measure. Choose some set of*

$m_\mathcal{A}$ measurements $\{\mathbf{a}_i\}_{i=1}^{m_\mathcal{A}}$ *uniformly from* $\mathcal{A}$. *Then, the caps* $\{\mathcal{C}_\mathbb{R}(\mathbf{a}_i)\}$ *cover the sphere* $\mathcal{S}_\mathbb{R}^{n-1}$ *with probability*

$$p_{cover}(m_\mathcal{A}, n) = 1 - \frac{1}{2^{m_\mathcal{A}-1}} \sum_{k=0}^{n-1} \binom{m_\mathcal{A}-1}{k}.$$

This is the probability of turning up n or more heads when flipping $m_\mathcal{A} - 1$ fair coins.

Proof. Consider the following two-step process for constructing the set $\{\mathbf{a}_i\}$. First, we sample $m_\mathcal{A}$ vectors $\{\mathbf{a}_i'\}$ independently and uniformly from $\mathcal{A}$. Note that, because $\mathcal{A}$ has positive measure, it holds with probability 1 that any subset of n vectors will be linearly independent.

Second, we define $\mathbf{a}_i = c_i \mathbf{a}_i'$, where $\{c_i\}$ are i.i.d. Bernoulli variables that take value $+1$ or -1 with probability $1/2$. We can think of this second step as randomly "flipping" a subset of uniform random vectors. Since $\mathcal{A}$ is symmetric and $\{\mathbf{a}_i'\}$ is sampled independently and uniformly, the random vectors $\{\mathbf{a}_i\}$ also have an independent and uniform distribution over $\mathcal{A}$. This construction may seem superfluous since both $\{\mathbf{a}_i\}$ and $\{\mathbf{a}_i'\}$ have the same distribution, but we will see below that this becomes useful.

Given a particular set of coin flips $\{c_i\}$, we can write the set of points that are *not* covered by the caps $\{\mathcal{C}_\mathbb{R}(\mathbf{a}_i)\}$ as

$$\bigcap_i \mathcal{C}_\mathbb{R}(-\mathbf{a}_i) = \bigcap_i \mathcal{C}_\mathbb{R}(-c_i \mathbf{a}_i'). \tag{6}$$

Note that there are $2^{m_\mathcal{A}}$ such intersections that can be formed, one for each choice of the sequence $\{c_i\}$. The caps $\{\mathcal{C}_\mathbb{R}(\mathbf{a}_i)\}$ cover the sphere whenever the intersection (6) is empty. Consider the set of planes $\{\{\mathbf{x} \,|\, \langle \mathbf{a}_i, \mathbf{x} \rangle = 0\}\}$. From Lemma 1, we know that $m_\mathcal{A}$ planes with a common intersection point divide the sphere into $r(n, m_\mathcal{A})$ non-empty regions. Each of these regions corresponds to the intersection (6) for one possible choice of $\{c_i\}$. Therefore, of the $2^{m_\mathcal{A}}$ possible intersections, only $r(n, m_\mathcal{A})$ of them are non-empty. Since the sequence $\{c_i\}$ is random, each intersection is equally likely to be chosen, and so the probability of covering the sphere is

$$p_{\text{cover}}(m_\mathcal{A}, n) = 1 - \frac{r(n, m_\mathcal{A})}{2^{m_\mathcal{A}}}.$$

$\square$

Remark 2. *Several papers have studied the probability of covering the sphere using points independently and uniformly chosen over the entire sphere. The only aspect that is unusual about Lemma 2 is the observation that this probability remains the same if we restrict our choices to the set $\mathcal{A}$, provided $\mathcal{A}$ is symmetric. We note that this result was observed by Gilbert [45] in the case $n = 3$, and we generalize it to any $n > 1$ using a similar argument.*

We now present a somewhat more complicated covering theorem. The next result considers the case where the measurement vectors are drawn only from a semisphere. We consider the question of whether these vectors cover enough area to contain not only their home semisphere, but another nearby semisphere as well.
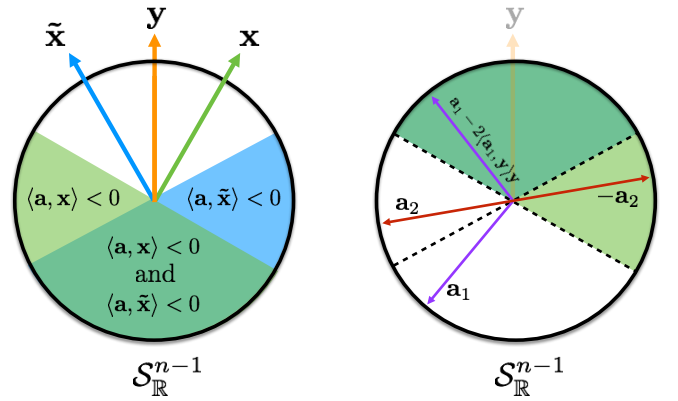

Fig. 1. (left) A diagram showing the construction of $\mathbf{x}$, $\mathbf{y}$, and $\tilde{\mathbf{x}}$ in the proof of Lemma 3. (right) The reflections defined in (7-9) map vectors $\mathbf{a}_i$ lying in the half-sphere $\mathbf{a}_i^T \mathbf{x} < 0$ onto the half-sphere $\mathbf{a}^T \mathbf{x} > 0$.

Lemma 3. *Consider two vectors $\mathbf{x}, \mathbf{y} \subset \mathcal{S}_\mathbb{R}^{n-1}$, and the caps $\mathcal{C}_\mathbb{R}(\mathbf{x})$ and $\mathcal{C}_\mathbb{R}(\mathbf{y})$. Let $\alpha = 1 - \frac{2}{\pi} \text{angle}(\mathbf{x}, \mathbf{y})$ be a measure of the similarity between the vectors $\mathbf{x}$ and $\mathbf{y}$. Draw some collection $\{\mathbf{a}_i \in \mathcal{C}_\mathbb{R}(\mathbf{x})\}_{i=1}^m$ of m vectors uniformly from $\mathcal{C}_\mathbb{R}(\mathbf{x})$ so that $m > 2n/\alpha$. Then*

$$\mathcal{C}_\mathbb{R}(\mathbf{y}) \subset \bigcup_i \mathcal{C}_\mathbb{R}(\mathbf{a}_i)$$

with probability at least

$$p_{cover}(m, n; \mathbf{x}, \mathbf{y}) \geq 1 - \exp\left(-\frac{(\alpha m - 2n)^2}{2m}\right).$$

Proof. Due to rotational symmetry, we assume $\mathbf{y} = [1, 0, \ldots, 0]^T$ without loss of generality. Consider the point $\tilde{\mathbf{x}} = [x_1, -x_2, \ldots, -x_n]^T$. This is the reflection of $\mathbf{x}$ over $\mathbf{y}$ (see Figure 1). Suppose we have some collection $\{\mathbf{a}_i\}$ independently and uniformly distributed on the entire sphere. Consider the collection of vectors

$$\mathbf{a}_i' = \begin{cases} \mathbf{a}_i, & \text{if } \langle \mathbf{a}_i, \mathbf{x} \rangle \geq 0 & (7) \\ \mathbf{a}_i - 2\langle \mathbf{a}_i, \mathbf{y} \rangle \mathbf{y} & \text{if } \langle \mathbf{a}_i, \mathbf{x} \rangle < 0, \langle \mathbf{a}_i, \tilde{\mathbf{x}} \rangle < 0 & (8) \\ -\mathbf{a}_i & \text{if } \langle \mathbf{a}_i, \mathbf{x} \rangle < 0, \langle \mathbf{a}_i, \tilde{\mathbf{x}} \rangle \geq 0. & (9) \end{cases}$$

The mapping $\mathbf{a}_i \to \mathbf{a}_i'$ maps the lower half sphere $\{\mathbf{a} \,|\, \langle \mathbf{a}, \mathbf{x} \rangle < 0\}$ onto the upper half sphere $\{\mathbf{a} \,|\, \langle \mathbf{a}, \mathbf{x} \rangle > 0\}$ using a combination of reflections and translations (depicted in Figure 1). Indeed, for all i we have $\langle \mathbf{a}_i', \mathbf{x} \rangle \geq 0$. This is clearly true in case (7) and (9). In case (8), we use the definition of $\mathbf{y}$ and $\tilde{\mathbf{x}}$ to write

$$\langle \mathbf{a}_i', \mathbf{x} \rangle = \langle \mathbf{a}_i, \mathbf{x} \rangle - 2\langle \mathbf{a}_i, \mathbf{y} \rangle \langle \mathbf{y}, \mathbf{x} \rangle$$
$$= \langle \mathbf{a}_i, \mathbf{x} \rangle - 2[\mathbf{a}_i]_1 x_1 = -\langle \mathbf{a}_i, \tilde{\mathbf{x}} \rangle \geq 0.$$

Because the mapping $\mathbf{a}_i \to \mathbf{a}_i'$ is onto and (piecewise) isometric, $\{\mathbf{a}_i'\}$ will be uniformly distributed over the half sphere $\{\mathbf{a} \,|\, \langle \mathbf{a}, \mathbf{x}^0 \rangle > 0\}$ whenever $\{\mathbf{a}_i\}$ are independently and uniformly distributed over the entire sphere.

Consider the "hourglass" shaped, symmetric set

$$\mathcal{A} = \{\mathbf{a} \,|\, \langle \mathbf{a}, \mathbf{x} \rangle > 0, \langle \mathbf{a}, \tilde{\mathbf{x}} \rangle > 0\} \cup \{\mathbf{a} \,|\, \langle \mathbf{a}, \mathbf{x} \rangle < 0, \langle \mathbf{a}, \tilde{\mathbf{x}} \rangle < 0\}.$$

We now make the following claim: $\mathcal{C}_\mathbb{R}(\mathbf{y}) \subset \bigcup_i \mathcal{C}_\mathbb{R}(\mathbf{a}_i')$

whenever

$$\mathcal{S}_{\mathbb{R}}^{n-1} \subset \bigcup_{\mathbf{a}_i \in \mathcal{A}} \mathcal{C}_{\mathbb{R}}(\mathbf{a}_i). \tag{10}$$

In words, if the caps defined by the subset of $\{\mathbf{a}_i\}$ in $\mathcal{A}$ cover the entire sphere, then the caps $\{\mathcal{C}_{\mathbb{R}}(\mathbf{a}_i')\}$ (which have centers in $\mathcal{C}_{\mathbb{R}}(\mathbf{x})$) not only cover $\mathcal{C}_{\mathbb{R}}(\mathbf{x})$, but also cover its neighbor cap $\mathcal{C}_{\mathbb{R}}(\mathbf{y})$. To justify this claim, suppose that (10) holds. Choose some $\boldsymbol{\delta} \in \mathcal{C}_{\mathbb{R}}(\mathbf{y})$. This point is covered by some cap $\mathcal{C}_{\mathbb{R}}(\mathbf{a}_i)$ with $\mathbf{a}_i \in \mathcal{A}$. If $\langle \mathbf{a}_i, \mathbf{x}\rangle > 0$ and $\langle \mathbf{a}_i, \tilde{\mathbf{x}}\rangle > 0$, then $\mathbf{a}_i = \mathbf{a}_i'$ and $\boldsymbol{\delta}$ is covered by $\mathcal{C}_{\mathbb{R}}(\mathbf{a}_i')$. Otherwise, we have $\langle \mathbf{a}_i, \mathbf{x}\rangle < 0$ and $\langle \mathbf{a}_i, \tilde{\mathbf{x}}\rangle < 0$, then

$$\langle \boldsymbol{\delta}, \mathbf{a}_i'\rangle = \langle \boldsymbol{\delta}, \mathbf{a}_i - 2\langle \mathbf{a}_i, \mathbf{y}\rangle \mathbf{y}\rangle$$
$$= \langle \boldsymbol{\delta}, \mathbf{a}_i\rangle - 2\langle \mathbf{a}_i, \mathbf{y}\rangle\langle \boldsymbol{\delta}, \mathbf{y}\rangle \geq \langle \boldsymbol{\delta}, \mathbf{a}_i\rangle \geq 0.$$

Note we have used the fact that $\langle \boldsymbol{\delta}, \mathbf{y}\rangle$ is real and non-negative because $\boldsymbol{\delta} \in \mathcal{C}_{\mathbb{R}}(y)$. We have also used $\langle \mathbf{a}_i, \mathbf{y}\rangle = [\mathbf{a}_i]_1 = \frac{1}{2}(\langle \mathbf{a}_i, \mathbf{x}\rangle + \langle \mathbf{a}_i, \tilde{\mathbf{x}}\rangle) < 0$, which follows from the definition of $\tilde{\mathbf{x}}$ and the definition of $\mathcal{A}$. In either case, $\langle \boldsymbol{\delta}, \mathbf{a}_i'\rangle > 0$, and we have $\boldsymbol{\delta} \in \mathcal{C}_{\mathbb{R}}(\mathbf{a}_i')$, which proves our claim.

We can now see that the probability that $\mathcal{C}_{\mathbb{R}}(\mathbf{y}) \subset \bigcup_i \mathcal{C}_{\mathbb{R}}(\mathbf{a}_i')$ is at least as high as the probability that (10) holds. Let $p_{\text{cover}}(m, n; \mathbf{x}, \mathbf{y} \,|\, m_{\mathcal{A}})$ denote the probability of covering $\mathcal{C}(\mathbf{y})$ conditioned on the number $m_{\mathcal{A}}$ of points lying in $\mathcal{A}$. From Lemma 2, we know that

$$p_{\text{cover}}(m, n; \mathbf{x}, \mathbf{y} \,|\, m_{\mathcal{A}}) \geq p_{\text{cover}}(m_{\mathcal{A}}, n) > p_{\text{cover}}(m_{\mathcal{A}}+1, n+1).$$

As noted in Lemma 2, the expression on the right is the chance of turning up more than n heads when flipping $m_{\mathcal{A}}$ fair coins.

The probability $p_{\text{cover}}(m, n; \mathbf{x}, \mathbf{y})$ is then given by

$$p_{\text{cover}}(m, n; \mathbf{x}, \mathbf{y}) = \mathbb{E}_{m_{\mathcal{A}}}[p_{\text{cover}}(m, n; \mathbf{x}, \mathbf{y} \,|\, m_{\mathcal{A}})]$$
$$\geq \mathbb{E}_{m_{\mathcal{A}}}[p_{\text{cover}}(m_{\mathcal{A}}, n)]$$
$$> \mathbb{E}_{m_{\mathcal{A}}}[p_{\text{cover}}(m_{\mathcal{A}} + 1, n + 1)].$$

The expression on the right hand side is the probability of getting more than n heads when one fair coin is flipped for every measurement $\mathbf{a}_i$ that lies in $\mathcal{A}$.

Let's evaluate how often this coin-flipping event occurs. The region $\mathcal{A}$ is defined by two planes that intersect at an angle of $\beta = \text{angle}(\mathbf{x}, \tilde{\mathbf{x}}) = 2\,\text{angle}(\mathbf{x}, \mathbf{y})$. The probability of a random point $\mathbf{a}_i$ lying in $\mathcal{A}$ is given by $\alpha = \frac{2\pi - 2\beta}{2\pi} = 1 - \frac{\beta}{\pi}$, which is the fraction of the unit sphere that lies either above or below both planes. The probability of a measurement $\mathbf{a}_i$ contributing to the heads count is half the probability of it lying in $\mathcal{A}$, or $\frac{1}{2}\alpha$. The probability of turning up more than n heads is therefore given by

$$1 - \sum_{k=0}^{n} \left(\frac{1}{2}\alpha\right)^k \left(1 - \frac{1}{2}\alpha\right)^{m-k} \binom{m}{k}.$$

Using Hoeffding's inequality, we obtain the following lower bound

$$p_{\text{cover}}(m, n; \mathbf{x}, \mathbf{y}) \geq 1 - \sum_{k=0}^{n} \left(\frac{1}{2}\alpha\right)^k \left(1 - \frac{1}{2}\alpha\right)^{m-k} \binom{m}{k}$$
$$\geq 1 - \exp\left(\frac{-(\alpha m - 2n)^2}{2m}\right),$$

which is only valid for $\alpha m > 2n$. $\qquad\square$

Remark 3. *In the proof of Lemma 3, we obtained a bound on $\mathbb{E}_{m_{\mathcal{A}}}[p_{cover}(m_{\mathcal{A}} + 1, n + 1)]$ using an intuitive argument about coin flipping probabilities. This expectation could have been obtained more rigorously (but with considerably more pain) using the method of probability generating functions.*

Lemma 3 contains most of the machinery needed for the proofs that follow. In the sequel, we prove a number of exact reconstruction theorems for (PM). Most of the results rely on short arguments followed by the invocation of Lemma 3.

We finally state a result that bounds from below the probability of covering the sphere with caps of small central angle. The following Lemma is a direct corollary of the results of Burgisser, Cucker, and Lotz in [47]. A derivation that uses their results is given in Appendix B.

Lemma 4. *Let $n \geq 9$, and $m > 2n$. Then the probability of covering the sphere $S_{\mathbb{R}}^{n-1}$ with independent uniformly sampled caps of central angle $\phi \leq \pi/2$ is lower bounded by*

$$p_{cover}(m, n, \phi) \geq$$
$$1 - \frac{(em)^n \sqrt{n-1}}{(2n)^{n-1}} \exp\left(-\frac{\sin^{n-1}(\phi)(m-n)}{\sqrt{8n}}\right)\cos(\phi)$$
$$- \exp\left(-\frac{(m - 2n + 1)^2}{2m - 2}\right).$$

IV. RECOVERY GUARANTEES

Using the uniqueness condition provided by Theorem 2 and the tools derived in Section III, we now develop sharp lower bounds on the success probability of PhaseMax for noiseless real- and complex-valued systems. The noisy case will be discussed in Section V-B.

A. The Real Case

We now study problem (PM) in the case that the unknown signal and measurement vectors are real valued. Consider some collection of measurement vectors $\{\mathbf{a}_i\}$ drawn independently and uniformly from $\mathcal{S}_{\mathbb{R}}^{n-1}$. For simplicity, we also consider the collection $\{\tilde{\mathbf{a}}_i\} = \{\text{phase}(\langle \mathbf{a}_i, \mathbf{x}^0\rangle)\mathbf{a}_i\}$ of aligned vectors that satisfy $\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0\rangle \geq 0$ for all i. Using this notation, Theorem 2 can be rephrased as a simple geometric condition.

Corollary 1. *Consider the set $\{\tilde{\mathbf{a}}_i\} = \{\text{phase}(\langle \mathbf{a}_i, \mathbf{x}^0\rangle)\mathbf{a}_i\}$ of aligned measurement vectors. Define the half sphere of aligned ascent directions*

$$\mathcal{D}_{\mathbb{R}} = \mathcal{C}_{\mathbb{R}}(\hat{\mathbf{x}}) = \{\boldsymbol{\delta} \in \mathcal{S}_{\mathbb{R}}^{n-1} \,|\, \langle \boldsymbol{\delta}, \hat{\mathbf{x}}\rangle \in \mathbb{R} \geq 0\}.$$

The true vector $\mathbf{x}^0$ will be the unique maximizer of (PM) if

$$\mathcal{D}_{\mathbb{R}} \subset \bigcup_i \mathcal{C}_{\mathbb{R}}(\tilde{\mathbf{a}}_i).$$

Proof. Choose some ascent direction $\boldsymbol{\delta} \in \mathcal{D}_{\mathbb{R}}$. If the assumptions of this Corollary hold, then there is some i with $\boldsymbol{\delta} \in \mathcal{C}_{\mathbb{R}}(\tilde{\mathbf{a}}_i)$, and so $\langle \tilde{\mathbf{a}}_i, \boldsymbol{\delta}\rangle \geq 0$. Since this is true for any $\boldsymbol{\delta} \in \mathcal{D}_{\mathbb{R}}$, the conditions of Theorem 2 are satisfied and exact reconstruction holds. $\qquad\square$

Using this observation, we can develop the following lower bound on the success probability of PhaseMax for real-valued systems.

Theorem 3. *Consider the case of recovering a real-valued signal $\mathbf{x}^0 \in \mathbb{R}^n$ from m noiseless measurements of the form (1) with measurement vectors $\mathbf{a}_i$, $i = 1, 2, \ldots, m$, sampled independently and uniformly from the unit sphere $\mathcal{S}_{\mathbb{R}}^{n-1}$. Then, the probability that PhaseMax recovers the true signal $\mathbf{x}^0$, denoted by $p_{\mathbb{R}}(m, n)$, is bounded from below as follows:*

$$p_{\mathbb{R}}(m, n) \geq 1 - \exp\left(\frac{-(\alpha m - 2n)^2}{2m}\right),$$

where $\alpha = 1 - \frac{2}{\pi} \operatorname{angle}(\mathbf{x}^0, \hat{\mathbf{x}})$ and $m > 2n/\alpha$.

Proof. Consider the set of m independent and uniformly sampled measurements $\{\mathbf{a}_i \in \mathcal{S}_{\mathbb{R}}^{n-1}\}_{i=1}^m$. The aligned vectors $\{\tilde{\mathbf{a}}_i = \operatorname{phase}(\langle \mathbf{a}_i, \mathbf{x}^0 \rangle)\mathbf{a}_i\}$ are uniformly distributed over the half sphere $\mathcal{C}_{\mathbb{R}}(\mathbf{x}^0)$. Exact reconstruction happens when the condition in Corollary 1 holds. To obtain a lower bound on the probability of this occurrence, we can simply invoke Lemma 3 with $\mathbf{x} = \mathbf{x}^0$ and $\mathbf{y} = \hat{\mathbf{x}}$. $\qquad\square$

B. The Complex Case

We now prove Theorem 1 given in Section I, which characterizes the success probability of PhaseMax for phase retrieval in complex-valued systems. For clarity, we restate our result in shorter form.

Theorem 1. *Consider the case of recovering a complex-valued signal $\mathbf{x}^0 \in \mathbb{C}^n$ from m noiseless measurements of the form (1), with $\{\mathbf{a}_i\}_{i=1}^m$ sampled independently and uniformly from the unit sphere $\mathcal{S}_{\mathbb{C}}^{n-1}$. Then, the probability that PhaseMax recovers the true signal $\mathbf{x}^0$ is bounded from below as follows:*

$$p_{\mathbb{C}}(m, n) \geq 1 - \exp\left(-\frac{(\alpha m - 4n)^2}{2m}\right),$$

where $\alpha = 1 - \frac{2}{\pi} \operatorname{angle}(\mathbf{x}^0, \hat{\mathbf{x}})$ and $m > 4n/\alpha$.

Proof. Consider the set $\{\tilde{\mathbf{a}}_i\} = \{\operatorname{phase}(\langle \mathbf{a}_i, \mathbf{x}^0 \rangle)\mathbf{a}_i\}$ of aligned measurement vectors. Define the half sphere of aligned ascent directions

$$\mathcal{D}_{\mathbb{C}} = \{\boldsymbol{\delta} \in \mathcal{S}_{\mathbb{C}}^{n-1} \mid \langle \boldsymbol{\delta}, \hat{\mathbf{x}} \rangle_{\Re} \in \mathbb{R}_0^+\}.$$

By Theorem 2, the true signal $\mathbf{x}^0$ will be the unique maximizer of (PM) if

$$\mathcal{D}_{\mathbb{C}} \subset \bigcup_i \mathcal{C}_{\mathbb{C}}(\tilde{\mathbf{a}}_i). \qquad (11)$$

Let us bound the probability of this event. Consider the set $\mathcal{B} = \{\boldsymbol{\delta} \mid \langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle_{\Im} = 0\}$. We now claim that (11) holds whenever

$$\mathcal{C}_{\mathbb{C}}(\hat{\mathbf{x}}) \cap \mathcal{B} \subset \bigcup_i \mathcal{C}_{\mathbb{C}}(\tilde{\mathbf{a}}_i). \qquad (12)$$

To prove this claim, consider some $\boldsymbol{\delta} \in \mathcal{D}_{\mathbb{C}}$. To keep notation light, we will assume without loss of generality that $\|\mathbf{x}^0\|_2 = 1$.

Form the vector $\boldsymbol{\delta}' = \boldsymbol{\delta} + j\langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle_{\Im} \mathbf{x}^0$, which is the projection of $\boldsymbol{\delta}$ onto $\mathcal{B}$. We can verify that $\boldsymbol{\delta}' \in \mathcal{B}$ by writing

$$\begin{aligned}
\langle \boldsymbol{\delta}', \mathbf{x}^0 \rangle &= \langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle + \langle j\langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle_{\Im} \mathbf{x}^0, \mathbf{x}^0 \rangle \\
&= \langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle - j\langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle_{\Im} \langle \mathbf{x}^0, \mathbf{x}^0 \rangle \\
&= \langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle - j\langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle_{\Im} = \langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle_{\Re},
\end{aligned}$$

which is real valued. Furthermore, $\boldsymbol{\delta}' \in \mathcal{C}_{\mathbb{C}}(\hat{\mathbf{x}})$ because

$$\langle \boldsymbol{\delta}', \hat{\mathbf{x}} \rangle = \langle \boldsymbol{\delta}, \hat{\mathbf{x}} \rangle + \langle j\langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle_{\Im} \mathbf{x}^0, \hat{\mathbf{x}} \rangle = \langle \boldsymbol{\delta}, \hat{\mathbf{x}} \rangle - j\langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle_{\Im} \langle \mathbf{x}^0, \hat{\mathbf{x}} \rangle.$$

The first term on the right is real-valued and non-negative (because $\boldsymbol{\delta} \in \mathcal{D}_{\mathbb{C}}$), and the second term is complex valued (because $\mathbf{x}^0$ is assumed to be aligned with $\hat{\mathbf{x}}$). It follows that $\langle \boldsymbol{\delta}', \hat{\mathbf{x}} \rangle_{\Re} \geq 0$ and $\boldsymbol{\delta}' \in \mathcal{C}_{\mathbb{C}}(\hat{\mathbf{x}})$. Since we already showed that $\boldsymbol{\delta}' \in \mathcal{B}$, we have $\boldsymbol{\delta}' \in \mathcal{C}_{\mathbb{C}}(\hat{\mathbf{x}}) \cap \mathcal{B}$. Suppose now that (12) holds. The claim will be proved if we can show that $\boldsymbol{\delta} \in \mathcal{D}$ is covered by one of the $\mathcal{C}_{\mathbb{C}}(\tilde{\mathbf{a}}_i)$. Since $\boldsymbol{\delta}' \in \mathcal{C}_{\mathbb{C}}(\hat{\mathbf{x}}) \cap \mathcal{B}$, there is some i with $\boldsymbol{\delta}' \in \mathcal{C}_{\mathbb{C}}(\tilde{\mathbf{a}}_i)$. But then

$$0 \leq \langle \boldsymbol{\delta}', \tilde{\mathbf{a}}_i \rangle_{\Re} = \langle \boldsymbol{\delta}, \tilde{\mathbf{a}}_i \rangle_{\Re} + \langle j\langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle_{\Im} \mathbf{x}^0, \tilde{\mathbf{a}}_i \rangle_{\Re} = \langle \boldsymbol{\delta}, \tilde{\mathbf{a}}_i \rangle_{\Re}. \tag{13}$$

We see that $\boldsymbol{\delta} \in \mathcal{C}_{\mathbb{C}}(\tilde{\mathbf{a}}_i)$, and the claim is proved.

We now know that exact reconstruction happens whenever condition (12) holds. We can put a bound on the frequency of this using Lemma 3. Note that the sphere $S_{\mathbb{C}}^{n-1}$ is isomorphic to $S_{\mathbb{R}}^{2n-1}$, and the set $\mathcal{B}$ is isomorphic to the sphere $S_{\mathbb{R}}^{2n-2}$. The aligned vectors $\{\tilde{\mathbf{a}}_i\}$ are uniformly distributed over a half sphere in $\mathcal{C}_{\mathbb{C}}(\mathbf{x}^0) \cap \mathcal{B}$, which is isomorphic to the upper half sphere in $S_{\mathbb{R}}^{2n-2}$. The probability of these vectors covering the cap $\mathcal{C}_{\mathbb{C}}(\hat{\mathbf{x}}) \cap \mathcal{B}$ is thus given by $p_{\text{cover}}(m, 2n - 1; \mathbf{x}^0, \hat{\mathbf{x}})$ from Lemma 3. We instead use the bound for $p_{\text{cover}}(m, 2n; \mathbf{x}^0, \hat{\mathbf{x}})$, which is slightly weaker. $\qquad\square$

Remark 4. *Theorems 1 and 3 guarantee exact recovery for a sufficiently large number of measurements m provided that $\operatorname{angle}(\mathbf{x}^0, \hat{\mathbf{x}}) < \frac{\pi}{2}$. In the case $\operatorname{angle}(\mathbf{x}^0, \hat{\mathbf{x}}) > \frac{\pi}{2}$, our theorems guarantee convergence to $-\mathbf{x}^0$ (which is also a valid solution) for sufficiently large m. Our theorems only fail for large m if $\arccos(\mathbf{x}^0, \hat{\mathbf{x}}) = \pi/2$, which happens with probability zero when the approximation vector $\hat{\mathbf{x}}$ is generated at random. See Section VI for more details.*

V. GENERALIZATIONS

Our theory thus far addressed the idealistic case in which the measurement vectors are independently and uniformly sampled from a unit sphere and for noiseless measurements. We now extend our results to more general random measurement ensembles and to noisy measurements.

A. Generalized Measurement Ensembles

The theorems of Section IV require the measurement vectors $\{\mathbf{a}_i\}$ to be drawn independently and uniformly from the surface of the unit sphere. This condition can easily be generalized to other sampling ensembles. In particular, our results still hold for all *rotationally symmetric distributions*. A distribution D is rotationally symmetric if the distribution of $\mathbf{a}/\|\mathbf{a}\|_2$ is uniform over the sphere when $\mathbf{a} \sim D$. For such a distribution, one can make the change of variables $\mathbf{a} \leftarrow \mathbf{a}/\|\mathbf{a}\|_2$, and then

apply Theorems 1 and 3 to the resulting problem. Note that this change of variables does not change the feasible set for (PM), and thus does not change the solution. Consequently, the same recovery guarantees apply to the original problem without explicitly implementing this change of variables. We thus have the following simple corollary.

Corollary 2. *The results of Theorem 1 and Theorem 3 still hold if the samples $\{\mathbf{a}_i\}$ are drawn from a multivariate Gaussian distribution with independent and identically distributed (i.i.d.) entries.*

Proof. A multivariate Gaussian distribution with i.i.d. entries is rotationally symmetric, and thus the change of variables $\mathbf{a} \leftarrow \mathbf{a}/\|\mathbf{a}\|_2$ yields an equivalent problem with measurements sampled uniformly from the unit sphere. $\square$

What happens when the distribution is not spherically symmetric? In this case, we can still guarantee recovery, but we require a larger number of measurements. The following result is, analogous to Theorem 1, for the noiseless complex case.

Theorem 4. *Suppose that m_D measurement vectors $\{\mathbf{a}_i\}_{i=1}^{m_D}$ are drawn from the unit sphere with (possibly non-uniform) probability density function $D : S_\mathbb{C}^{n-1} \to \mathbb{R}$. Let $\ell_D \leq \inf_{\mathbf{x} \in S_\mathbb{C}^{n-1}} D(\mathbf{x})$ be a lower bound on D over the unit sphere and let $\alpha = 1 - \frac{2}{\pi} \operatorname{angle}(\mathbf{x}^0, \hat{\mathbf{x}})$ as above. We use $s_n = \frac{2\pi^n}{\Gamma(n)}$ to denote the "surface area" of the complex sphere $\mathcal{S}_\mathbb{C}^{n-1}$, and set $m_U = \lfloor m_D s_n \ell_D \rfloor$. Then, exact reconstruction is guaranteed with probability at least*

$$1 - \exp\left(-\frac{(\alpha m_U - 4n)^2}{2m_U}\right)$$

whenever $\alpha m_U > 4n$ and $\ell_D > 0$. In other words, exact recovery with m_D non-uniform measurements happens at least as often as with m_U uniform measurements.

Proof. We compare two measurement models, a uniform measurement model in which m_U measurements are drawn uniformly from a unit sphere, and a non-uniform measurement model in which m_D measurements are drawn from the distribution D. Note that the sphere $\mathcal{S}_\mathbb{C}^{n-1}$ has surface area $s_n = \frac{2\pi^n}{\Gamma(n)}$, and the uniform density function U on this sphere has constant value s_n^{-1}. Consider some collection of measurements $\{\mathbf{a}_i^U\}_{i=1}^{m_U}$ drawn from the uniform model. This ensemble of measurements lies in $\mathbb{R}^{n \times m_U}$, and the probability density of sampling this measurement ensemble (in the space $\mathbb{R}^{n \times m_U}$) is

$$m_U! s_n^{-m_U}. \tag{14}$$

Now consider a random ensemble $\{\mathbf{a}_i^D\}_{i=1}^{m_D}$ drawn with density D. Given $\{\mathbf{a}_i^U\}$, the event that $\{\mathbf{a}_i^U\} \subset \{\mathbf{a}_i^D\}$ has density (in the space $\mathbb{R}^{n \times m_U}$)

$$\frac{m_D!}{(m_D - m_U)!} \prod_{i=1}^{m_U} D(\mathbf{a}_i). \tag{15}$$

The ratio of the non-uniform density (15) to the uniform density

(14) is

$$\binom{m_D}{m_U} \prod_{i=1}^{m_U} s_n D(\mathbf{a}_i) \geq \binom{m_D}{m_U} (s_n \ell_D)^{m_U} \geq \left(\frac{m_D s_n \ell_D}{m_U}\right)^{m_U},$$
$$\tag{16}$$

where we have used the bound $\binom{m_D}{m_U} > (m_D/m_U)^k$ to obtain the estimate on the right hand side. The probability of exact reconstruction using the non-uniform model will always be at least as large as the probability under the uniform model, provided the ratio (16) is one or higher. This holds whenever $m_U \leq m_D s_n \ell_D$. It follows that the probability of exact recovery using the non-uniform measurements is at least the probability of exact recovery from a uniform model with $m_U = \lfloor m_D s_n \ell_D \rfloor$ measurements. This probability is what is given by Theorem 1. $\square$

B. Noisy Measurements

We now analyze the sensitivity of PhaseLift to the measurement noise $\{\eta_i\}$. For brevity, we focus only on the case of complex-valued signals. To analyze the impact of noise, we re-write the problem (PM) in the following equivalent form:

$$\begin{cases} \underset{\mathbf{x} \in \mathcal{H}^n}{\text{maximize}} & \langle \mathbf{x}, \hat{\mathbf{x}} \rangle_\Re \\ \text{subject to} & |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2 \leq \hat{b}_i^2 + \eta_i, \ i = 1, 2, \ldots, m. \end{cases} \tag{17}$$

Here, $\hat{b}_i^2 = |\langle \mathbf{a}_i, \mathbf{x}^0 \rangle|^2$ is the (unknown) true magnitude measurement and $b_i^2 = \hat{b}_i^2 + \eta_i$. We are interested in bounding the impact that these measurement errors have on the solution to (PM). Note that the severity of a noise perturbation of size η_i depends on the (arbitrary) magnitude of the measurement vector $\mathbf{a}_i$. For this reason, we assume the vectors $\{\mathbf{a}_i\}$ have unit norm throughout this section.

We will begin by proving results only for the case of non-negative noise. We will then generalize our analysis to the case of arbitrary bounded noise. The following result gives a geometric characterization of the reconstruction error.

Theorem 5. *Suppose the vectors $\{\mathbf{a}_i \in \mathbb{C}^n\}$ in (17) are normalized to have unit length, and the noise vector $\boldsymbol{\eta}$ is non-negative. Let r be the maximum relative noise, defined by*

$$r = \max_{i=1,2,\ldots,m} \left\{\frac{\eta_i}{\hat{b}_i}\right\}, \tag{18}$$

and let $\mathcal{D}_\mathbb{C} = \{\boldsymbol{\delta} \in \mathcal{S}_\mathbb{C}^{n-1} | \langle \boldsymbol{\delta}, \hat{\mathbf{x}} \rangle_\Re \geq 0, \langle \boldsymbol{\delta}, \hat{\mathbf{x}} \rangle_\Im = 0\}$ be the set of aligned descent directions. Choose some error bound $\varepsilon > r/2$, and define the angle $\theta = \arccos(r/2\varepsilon)$. If the caps $\{\mathcal{C}_\mathbb{C}(\tilde{\mathbf{a}}_i, \theta)\}$ cover $\mathcal{D}_\mathbb{C}$, then the solution $\mathbf{x}^\star$ of (PM), and equivalently of the problem in (17), satisfies the bound

$$\|\mathbf{x}^\star - \mathbf{x}^0\|_2 \leq \varepsilon.$$

Proof. We first reformulate the problem (17) as

$$\begin{cases} \underset{\boldsymbol{\Delta} \in \mathbb{C}^n}{\text{maximize}} & \langle \mathbf{x}^0 + \boldsymbol{\Delta}, \hat{\mathbf{x}} \rangle_\Re \\ \text{subject to} & |\langle \tilde{\mathbf{a}}_i, (\mathbf{x}^0 + \boldsymbol{\Delta}) \rangle|^2 \leq \hat{b}_i^2 + \eta_i, \\ & i = 1, 2, \ldots, m, \end{cases} \tag{19}$$

where $\boldsymbol{\Delta} = \mathbf{x}^\star - \mathbf{x}^0$ is the recovery error vector and $\{\tilde{\mathbf{a}}_i\} = \{\mathrm{phase}(\langle \mathbf{a}_i, \mathbf{x}^0 \rangle)\mathbf{a}_i\}$ are aligned measurement vectors. In this form, the recovery error vector $\boldsymbol{\Delta}$ appears explicitly. Because we assume the errors $\{\eta_i\}$ to be non-negative, the true signal $\mathbf{x}^0$ is feasible for (17). It follows that the optimal objective of the perturbed problem (17) must be at least as large as the optimal value achieved by $\mathbf{x}^0$, i.e., $\langle \boldsymbol{\Delta}, \hat{\mathbf{x}} \rangle_\Re \geq 0$. Furthermore, the solution $\mathbf{x}^0 + \boldsymbol{\Delta}$ must be aligned with $\hat{\mathbf{x}}$, as is the true signal $\mathbf{x}^0$, and so $\langle \boldsymbol{\Delta}, \hat{\mathbf{x}} \rangle \in \mathbb{R}$. For the reasons just described, we know that the unit vector $\boldsymbol{\delta} = \boldsymbol{\Delta}/\|\boldsymbol{\Delta}\|_2 \in \mathcal{D}_\mathbb{C}$.

Our goal is to put a bound on the magnitude of the recovery error $\boldsymbol{\Delta}$. We start by reformulating the constraints in (19) to get

$$|\langle \tilde{\mathbf{a}}_i, (\mathbf{x}^0 + \boldsymbol{\Delta}) \rangle|^2 =$$
$$|\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle|^2 + 2[\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle^* \langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle]_\Re + |\langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle|^2 \leq \hat{b}_i^2 + \eta_i.$$

Subtracting $|\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle|^2 = \hat{b}_i^2$ from both sides yields

$$2[\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle^* \langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle]_\Re + |\langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle|^2 \leq \eta_i.$$

Since η_i is non-negative, we have

$$\begin{aligned}
\eta_i &\geq 2[\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle^* \langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle]_\Re + |\langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle|^2 \\
&= 2\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle \langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle_\Re + |\langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle|^2 \\
&\geq 2\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle \langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle_\Re + |\langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle_\Re|^2. \quad (20)
\end{aligned}$$

This inequality can hold only if $\langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle_\Re$ is sufficiently small. In particular we know

$$\begin{aligned}
\langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta} \rangle_\Re &\leq -\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle + \sqrt{(\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle)^2 + \eta_i} \\
&\leq \frac{\eta_i}{2\langle \tilde{\mathbf{a}}_i, \mathbf{x}^0 \rangle} = \frac{\eta_i}{2\hat{b}_i} \leq \frac{r}{2}. \quad (21)
\end{aligned}$$

Now suppose that $\|\boldsymbol{\Delta}\|_2 > \varepsilon$. From (21) we have

$$\langle \tilde{\mathbf{a}}_i, \boldsymbol{\delta} \rangle_\Re \leq \langle \tilde{\mathbf{a}}_i, \boldsymbol{\Delta}/\|\boldsymbol{\Delta}\|_2 \rangle_\Re < \frac{r}{2\varepsilon},$$

Therefore $\boldsymbol{\delta} \notin \mathcal{C}_\mathbb{C}(\tilde{\mathbf{a}}_i, \theta)$ where $\theta = \arccos(r/2\varepsilon)$. This is a contradiction because the caps $\{C_\mathbb{C}(\tilde{\mathbf{a}}_i, \theta)\}$ cover $\mathcal{D}_\mathbb{C}$. It follows that $\|\boldsymbol{\Delta}\|_2 \leq \varepsilon$. □

Using this result, we can bound the reconstruction error in the noisy case. For brevity, we present results only for the complex-valued case.

Theorem 6. *Suppose the vectors $\{\mathbf{a}_i\}$ in (17) are independently and uniformly distributed in $\mathcal{S}_\mathbb{C}^{n-1}$, and the noise vector $\boldsymbol{\eta}$ is non-negative. Let r be the maximum relative error defined in (18). Choose some error bound $\varepsilon > r/2$, and define the angle $\phi = \arccos(r/2\varepsilon) - \mathrm{angle}(\mathbf{x}^0, \hat{\mathbf{x}})$. Then, the solution $\mathbf{x}^\star$ to (PM) satisfies*

$$\|\mathbf{x}^\star - \mathbf{x}^0\| \leq \varepsilon.$$

with probability at least

$$p_{cover}(m, 2n, \phi) \geq$$
$$1 - \frac{(em)^{2n}\sqrt{2n-1}}{(4n)^{2n-1}} \exp\left(-\frac{\sin^{2n-1}(\phi)(m-n)}{\sqrt{16n}}\right)\cos(\phi)$$
$$- \exp\left(-\frac{(m-4n+1)^2}{2m-2}\right)$$

when $n \geq 5$ and $m > 4n$.

Proof. Define the following two sets:

$$\begin{aligned}
\mathcal{D} &= \{\boldsymbol{\delta} \in \mathcal{S}_\mathbb{C}^{n-1} \,|\, \langle \boldsymbol{\delta}, \hat{\mathbf{x}} \rangle \in \mathbb{R}_0^+\} \\
\mathcal{D}^0 &= \{\boldsymbol{\delta} \in \mathcal{S}_\mathbb{C}^{n-1} \,|\, \langle \boldsymbol{\delta}, \mathbf{x}^0 \rangle \in \mathbb{R}_0^+\}.
\end{aligned}$$

We now claim that the conditions of Theorem 5 hold whenever

$$\mathcal{D}^0 \subset \bigcup_i \mathcal{C}_\mathbb{C}(\tilde{\mathbf{a}}_i, \phi) \quad (22)$$

where $\{\tilde{\mathbf{a}}_i = \mathrm{phase}(\langle \mathbf{a}_i, \mathbf{x}^0 \rangle)\mathbf{a}_i\}$ is the set of aligned measurement vectors. To prove this claim, choose some $\boldsymbol{\delta} \in \mathcal{D}$ and assume that (22) holds. Since the half-sphere $\mathcal{D}^0$ can be obtained by rotating $\mathcal{D}$ by a principal angle of $\mathrm{angle}(\hat{\mathbf{x}}, \mathbf{x}^0)$, there is some point $\boldsymbol{\delta}^0 \in \mathcal{D}^0$ with $\mathrm{angle}(\boldsymbol{\delta}, \boldsymbol{\delta}^0) \leq \mathrm{angle}(\hat{\mathbf{x}}, \mathbf{x}^0)$. By property (22), there is some cap $\mathcal{C}_\mathbb{C}(\tilde{\mathbf{a}}_i, \phi)$ that contains $\boldsymbol{\delta}_0$. By the triangle inequality for spherical geometry it follows that:

$$\begin{aligned}
\mathrm{angle}(\boldsymbol{\delta}, \tilde{\mathbf{a}}_i) &\leq \mathrm{angle}(\boldsymbol{\delta}, \boldsymbol{\delta}^0) + \mathrm{angle}(\boldsymbol{\delta}^0, \tilde{\mathbf{a}}_i) \\
&\leq \mathrm{angle}(\mathbf{x}^0, \hat{\mathbf{x}}) + \phi \leq \theta.
\end{aligned}$$

Therefore, $\boldsymbol{\delta} \in \mathcal{C}_\mathbb{C}(\tilde{\mathbf{a}}_i, \theta)$, and the claim is proved.

It only remains to put a bound on the probability that (22) occurs. Note that the aligned vectors $\{\tilde{\mathbf{a}}_i\}$ are uniformly distributed in $\mathcal{D}^0$, which is isomorphic to a half-sphere in $S_\mathbb{R}^{2n-2}$. The probability of covering the half sphere $S_\mathbb{R}^{2n-2}$ with uniformly distributed caps drawn from that half sphere is at least as great as the probability of covering the whole sphere $S_\mathbb{R}^{2n-2}$ with caps drawn uniformly from the entire sphere. This probability is given by Lemma 4 as $p_{\mathrm{cover}}(m, 2n-1, \phi)$, and is lower bounded by $p_{\mathrm{cover}}(m, 2n, \phi)$. □

We now consider the case of non-negative noise. In this case, we bound the error by converting the problem into an equivalent problem with non-negative noise, and then apply Theorem 6.

Theorem 7. *Suppose the vectors $\{\mathbf{a}_i\}$ in (17) are normalized to have unit length, and that $\eta_i > -\hat{b}_i^2$ for all i.[2] Define the following measures of the noise*

$$s^2 = \min_{i=1,2,\ldots,m}\left\{\frac{\hat{b}_i^2 + \eta_i}{\hat{b}_i^2}\right\} = \min_{i=1,2,\ldots,m}\left\{\frac{b_i^2}{\hat{b}_i^2}\right\}$$

and

$$r = \frac{1}{s}\max_{i=1,2,\ldots,m}\left\{\hat{b}_i^2 - s^2\hat{b}_i + \frac{\eta_i}{\hat{b}_i}\right\}.$$

Choose some error bound $\varepsilon > r/2$, and define the angle $\phi = \arccos(r/2\varepsilon) - \mathrm{angle}(\mathbf{x}^0, \hat{\mathbf{x}})$. Then, we have the bound

$$\|\mathbf{x}^\star - \mathbf{x}^0\|_2 \leq \varepsilon + (1-s)\|\mathbf{x}^0\|_2.$$

with probability at least

$$p_{cover}(m, 2n, \phi) \geq$$
$$1 - \frac{(em)^{2n}\sqrt{2n-1}}{(4n)^{2n-1}} \exp\left(-\frac{\sin^{2n}(\phi)(m-n)}{\sqrt{16n}}\right)\cos(\phi)$$
$$- \exp\left(-\frac{(m-4n+1)^2}{2m-2}\right)$$

[2]This assumption is required so that $b_i^2 = \hat{b}_i^2 + \eta_i > 0$.

when $n \geq 5$ and $m > 4n$.

Proof. Consider the "shrunk" version of problem (17)

$$\begin{cases} \underset{\mathbf{x} \in \mathcal{H}^n}{\text{maximize}} & \langle \mathbf{x}, \hat{\mathbf{x}} \rangle_{\Re} \\ \text{subject to} & |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2 \leq s^2 \hat{b}_i^2 + \zeta_i, \quad i = 1, 2, \ldots, m. \end{cases}$$
(23)

for some real-valued "shrink factor" $s > 0$. Clearly, if $\mathbf{x}^0$ is aligned with $\hat{\mathbf{x}}$ and satisfies $|\langle \mathbf{a}_i, \mathbf{x}^0 \rangle| = b_i$ for all i, then $s\mathbf{x}^0$ is aligned with $\hat{\mathbf{x}}$ and satisfies $|\langle \mathbf{a}_i, s\mathbf{x}^0 \rangle| = sb_i$. We can now transform the noisy problem (17) into an equivalent problem with non-negative noise by choosing

$$s^2 = \min_{i=1,2,\ldots,m} \left\{ \frac{\hat{b}_i^2 + \eta_i}{\hat{b}_i^2} \right\} \quad \text{and} \quad \zeta_i = \hat{b}_i^2 - s^2 \hat{b}_i^2 + \eta_i \geq 0.$$

We then have $(sb_i)^2 + \zeta_i = b_i^2 + \eta_i^2$, and so problem (23) is equivalent to problem (17). However, the noise ζ_i in problem (23) is non-negative, and thus we can apply Theorem 6. This theorem requires the constant r for the shrunken problem, which is now

$$\begin{aligned} r_{\text{shrunk}} &= \max_{i=1,2,\ldots,m} \left\{ \frac{\zeta_i}{s\hat{b}_i} \right\} \\ &= \frac{1}{s} \max_{i=1,2,\ldots,m} \left\{ \hat{b}_i^2 - s^2 \hat{b}_i + \eta_i / \hat{b}_i \right\}. \end{aligned}$$

The solution to the shrunk problem (23) satisfies $\|\mathbf{x}^\star - s\mathbf{x}^0\|_2 \leq \epsilon$, with probability $p_{\text{cover}}(m, 2n, \phi)$, where $\phi = \arccos(r_{\text{shrunk}}/2\epsilon) - \text{angle}(\mathbf{x}^0, \hat{\mathbf{x}})$. If this condition is fulfilled, then we have

$$\begin{aligned} \|\mathbf{x}^\star - \mathbf{x}^0\|_2 &\leq \|\mathbf{x}^\star - s\mathbf{x}^0 + s\mathbf{x}^0 - \mathbf{x}^0\|_2 \\ &\leq \|\mathbf{x}^\star - s\mathbf{x}^0\|_2 + \|s\mathbf{x}^0 - \mathbf{x}^0\|_2 \\ &\leq \epsilon + (1-s)\|\mathbf{x}^0\|_2, \end{aligned}$$

which concludes the proof. $\square$

Remark 5. *Theorem 7 requires the noise to be sufficiently small so that the measurements are non-negative, and the PhaseMax formulation is feasible. A natural extension that avoids this caveat is to enforce constraints with a hinge penalty rather than a hard constraint. In addition, it is possible to achieve better noise robustness in this case. This direction was studied in [39].*

VI. How to Compute Approximation Vectors?

There exist a variety of algorithms that compute approximation vectors[3], such as the (truncated) spectral initializer [25], [28] or corresponding optimized variants [34], [35], the Null initializer [33], the orthogonality-promoting method [29], or least-squares methods [48]. We now show that even randomly generated approximation vectors guarantee the success of PhaseMax with high probability given a sufficiently large number of measurements. We then show that more sophisticated methods guarantee success with high probability if the number of measurements depends linearly on n.

[3] Approximation vectors are also known as initialization or anchor vectors.

A. Random Initialization

Consider the use of approximation vectors $\hat{\mathbf{x}}$ drawn randomly from the unit sphere $S_\mathbb{R}^{n-1}$. Do we expect such approximation vectors to be accurate enough to recover the unknown signal? To find out, we analyze the inner product between two real-valued random vectors on the unit sphere. Note that we only care about the *magnitude* of this inner product. If the inner product is negative, then PhaseMax simply recovers $-\mathbf{x}^0$ rather than $\mathbf{x}^0$. Our analysis will make use of the following result.

Lemma 5. *Consider the angle $\beta = \text{angle}(\mathbf{x}, \mathbf{y})$ between two random vectors $\mathbf{x}, \mathbf{y} \in \mathcal{S}_\mathcal{H}^{n-1}$ sampled independently and uniformly from the unit sphere. Then, the expected magnitude of the cosine distance between the two random vectors satisfies*

$$\sqrt{\frac{2}{\pi n}} \leq \mathbb{E}[|\cos(\beta)|] \leq \sqrt{\frac{2}{\pi(n - \frac{1}{2})}}, \text{ for } \mathcal{H} = \mathbb{R} \quad (24)$$

$$\sqrt{\frac{1}{\pi n}} \leq \mathbb{E}[|\cos(\beta)|] \leq \sqrt{\frac{4}{\pi(4n-1)}}, \text{ for } \mathcal{H} = \mathbb{C}. \quad (25)$$

Proof. We first consider the real case. The quantity $\cos(\beta) = \langle \mathbf{x}, \mathbf{y} \rangle / (\|\mathbf{x}\|_2 \|\mathbf{y}\|_2)$ is simply the sample correlation between two random vectors, whose distribution function is given by [49]

$$f(z) = \frac{(1 - z^2)^{\frac{n-3}{2}}}{2^{n-2} B\left(\frac{n-1}{2}, \frac{n-1}{2}\right)}, \quad (26)$$

where B is the beta function and $z \in [-1, +1]$. Hence, the expectation of the magnitude of the inner product is given by

$$\mathbb{E}[|\cos(\beta)|] = 2 \frac{\int_0^1 z(1 - z^2)^{\frac{n-3}{2}} \, dz}{2^{n-2} B\left(\frac{n-1}{2}, \frac{n-1}{2}\right)}.$$

The integral in the numerator was studied in [50, Eq. 31] and evaluates to $\frac{1}{2} B(1, \frac{n-1}{2})$. Plugging this expression into (26), and using the identity $\Gamma(a)/\Gamma(a/2) = 2^{a-1} \Gamma(\frac{a+1}{2})/\sqrt{\pi}$, we get

$$\mathbb{E}[|\cos(\beta)|] = \frac{\Gamma(\frac{n}{2})}{\sqrt{\pi} \, \Gamma(\frac{n+1}{2})}.$$

Finally, by using bounds on ratios of Gamma functions [51], we obtain the bounds in (24) for real-valued vectors. The bounds in (25) for complex-valued vectors are obtained by noting that $\mathcal{S}_\mathbb{C}^{n-1}$ is isomorphic to $\mathcal{S}_\mathbb{R}^{2n-1}$ and by simply replacing $n \leftarrow 2n$ in the bounds for the real-valued case. $\square$

B. Spectral Initializers

Spectral initializers were first used in phase retrieval by [25], and enable the computation of an approximation vector $\hat{\mathbf{x}}$ that exhibits strong theoretical properties [28]. This class of initializers was analyzed in detail in [34], and [35] developed an optimal spectral initializer for a class of problems including phase retrieval.

Fix some $\delta > 1$, and consider a measurement process with $m = \delta n$ measurement vectors sampled from a Gaussian distribution. Then, the spectral initializers in [34], [35] are known to deliver an initializer $\hat{\mathbf{x}}$ with $\lim_{n \to \infty} \text{angle}(\hat{\mathbf{x}}, \mathbf{x}^0) > \epsilon$ for some positive ϵ. Equivalently, there is some accuracy parameter

$\alpha_\star = \lim_{n\to\infty} 1 - \frac{2}{\pi}\operatorname{angle}(\hat{\mathbf{x}}, \mathbf{x}^0) > 0$. By combining this result with Theorem 1, we see that spectral initializers enable PhaseMax to succeed for large n with high probability provided that $m > n/\alpha_*$ measurements are used for signal recovery[4]; a number of measurements that is linear in n. Note that similar non-asymptotic guarantees can be achieved for finite n using other results on spectral initializers (e.g., Prop. 8 in [28]), although with less sharp bounds.

In Section VII-B, we investigate the numerical value of α_* and make quantitative statements about how many measurements are required when using a spectral initializer. As we will see there, PhaseMax requires roughly $m = 5.5n$ noiseless Gaussian measurements in theory[5], and somewhat fewer measurements in practice.

VII. Discussion

This section briefly compares our theoretical results to that of existing algorithms. We furthermore demonstrate the sharpness of our recovery guarantees and show the practical limits of PhaseMax.

A. Comparison with Existing Recovery Guarantees

Table I compares our noiseless recovery guarantees in a complex system to that of PhaseLift [19], truncated Wirtinger flow (TWF) [19], and truncated amplitude flow (TAW) [29]. We also compare to the recovery guarantee provided for PhaseMax using classical machine learning methods in [3]. [6] We see that PhaseMax requires the same sample complexity (number of required measurements) as compared to PhaseLift, TWF, and TAW, when used together with the truncated spectral initializer [19]. While the constants c_0, c_1, and c_2 in the recovery guarantees for all of the other methods are generally very large, our recovery guarantees contain no unspecified constants, explicitly depend on the approximation factor α, and are surprisingly sharp. We next demonstrate the accuracy of our results via numerical simulations.

B. Tightness of Recovery Guarantees

We now investigate the tightness of the recovery guarantees of PhaseMax using both experiments and theory. First, we compare the empirical success probability of PhaseMax in a noiseless and complex-valued scenario with measurement vectors taken independently and uniformly from the unit sphere. All experiments were obtained by using the implementations provided in the software library PhasePack [52]. This software library provides efficient implementations of phase retrieval methods using fast adaptive gradient solvers [53]. We declare signal recovery a success whenever the relative reconstruction error satisfies

$$RRE = \frac{\|\mathbf{x}^0 - \mathbf{x}\|_2^2}{\|\mathbf{x}^0\|_2^2} < 10^{-5}. \tag{27}$$

We compare empirical rates of success to the theoretical lower bound in Theorem 1. Figure 2 shows results for $n = 100$ and $n = 500$ measurements, where we artificially generate an approximation $\hat{\mathbf{x}}$ for different angles $\beta = \operatorname{angle}(\hat{\mathbf{x}}, \mathbf{x}^0)$ measured in degrees. Clearly, our theoretical lower bound accurately predicts the real-world performance of PhaseMax. For large n and large β, the gap between theory and practice becomes extremely tight. We furthermore observe a sharp phase transition between failure and success, with the transition getting progressively sharper for larger dimensions n.

Next, we investigate the tightness of the recovery guarantees when an initial vector is chosen using the spectral initialization method [34], [35]. Using analytical formulas, the authors of [35] calculate the accuracy of their proposed spectral initializer for real-valued signal estimation from Gaussian measurements with large n. They find that, when $2n$ measurements are used for spectral estimation, the initializer and signal have a squared cosine similarity of over 0.6. This corresponds to an accuracy parameter of $\alpha > 0.55$, and PhaseMax can recover the signal exactly with $m = n/\alpha \approx 3.5n$ measurements. This is within a factor of 2 of the information-theoretic lower bound for real-valued signal recovery, which is $2n - 2$ measurements.

Note that our analysis of PhaseMax assumes that the measurements used for recovery are statistically independent from those used by the initializer. For this reason, our theory does not allow us to perform signal recovery by "recycling" the measurements used for initialization. While we do not empirically observe any change in behavior of the method when the initialization measurements are used for recovery, the above reconstruction bounds formally require $5.5n$ measurements ($2n$ for initialization, plus $3.5n$ for recovery). The bounds in [3] are uniform with respect to the initializer, and thus enable measurement recycling (although this does not result in tighter bounds for the overall number of measurements because the required constants are quite large).

Finally, we would like to mention that asymptotically *exact* performance bounds for PhaseMax have recently been derived in [41], and tighter bounds for non-lifting phase retrieval have been shown for the method PhaseLamp, which repeatedly uses PhaseMax within an iterative process.

C. Experimental Results

We compare PhaseMax to other algorithms using synthetic and empirical datasets. The implementations used here are publicly available as part of the PhasePack software library [52]. This software library contains scripts for comparing different phase retrieval methods, including scripts for reproducing the experiments shown here.

1) Comparisons Using Synthetic Data: We briefly compare PhaseMax to a select set of phase retrieval algorithms in terms of relative reconstruction error with Gaussian measurements. We emphasize that this comparison is by no means intended to be exhaustive and serves the sole purpose of demonstrating the efficacy and limits of PhaseMax (see, e.g., [17], [21] for more extensive phase retrieval algorithm comparisons). We compare the Gerchberg-Saxton algorithm [7], the Fienup algorithm [8], truncated Wirtinger flow [28], and PhaseMax—all of these

---

[4]Our results assume the initial vector is independent of the measurement vectors. The total number of measurements needed is thus $m > n/\alpha_* + \delta n$ for initialization and recovery combined.

[5]The constants that appear in the bounds [34], [35] were approximated using stochastic numerical methods, and were not obtained exactly using analysis.

[6]Since AltMinPhase [25] requires an online measurement model that differs significantly from the other algorithms considered here, we omit a comparison.

TABLE I
COMPARISON OF THEORETICAL RECOVERY GUARANTEES FOR NOISELESS PHASE RETRIEVAL

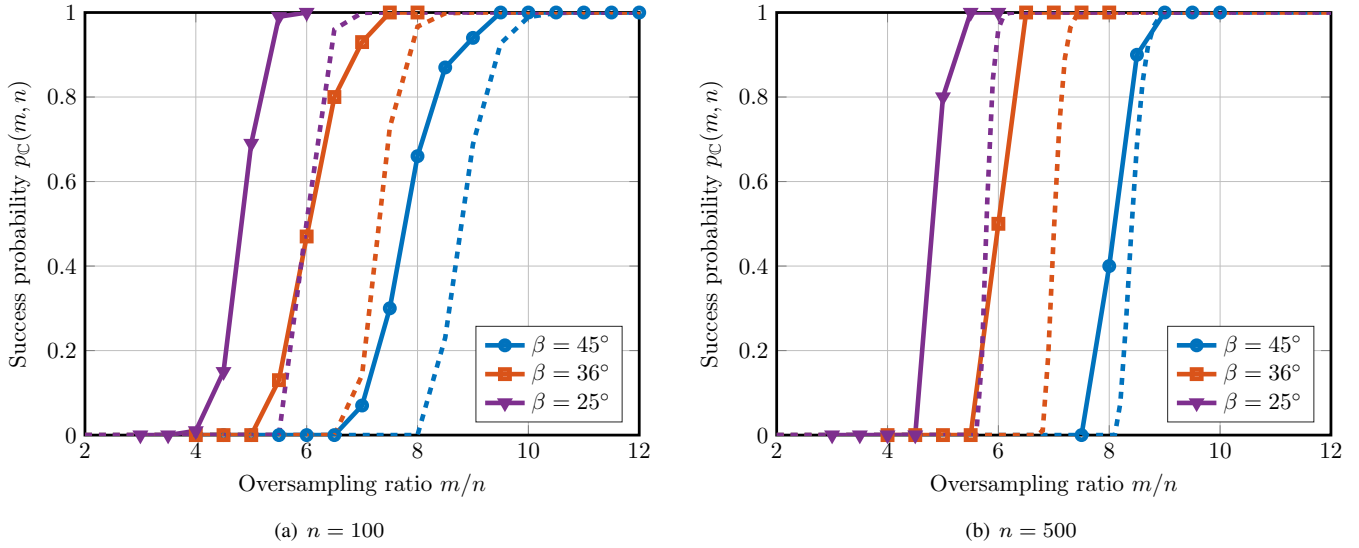| Algorithm	Sample complexity	Lower bound on $p_{\mathbb{C}}(m,n)$
PhaseMax	$m > 4n/\alpha$	$1 - e^{-(\alpha m - 4n)^2/(2m)}$
PhaseLift [19]	$m \geq c_0 n$	$1 - c_1 e^{-c_2 m}$
TWF [19]	$m \geq c_0 n$	$1 - c_1 e^{-c_2 m}$
TAF [29]	$m \geq c_0 n$	$1 - (m+5)e^{-n/2} - c_1 e^{-c_2 m} - 1/n^2$
Bahmani and Romberg [3]	$m > \frac{32}{\sin^4(\alpha)} \log\left(\frac{8e}{\sin^4(\alpha)}\right) n$	$1 - 8e^{-\sin^4(\alpha)\left(M - \frac{32}{\sin^4(\alpha)}\log\left(\frac{8e}{\sin^4(\alpha)}\right)N\right)/16}$

Fig. 2. Comparison between the empirical success probability (solid lines) and our theoretical lower bound (dashed lines) for varying angles β between the true signal and the approximation vector. Our theoretical results accurately characterize the empirical success probability of PhaseMax. Furthermore, PhaseMax exhibits a sharp phase transition for larger dimensions.
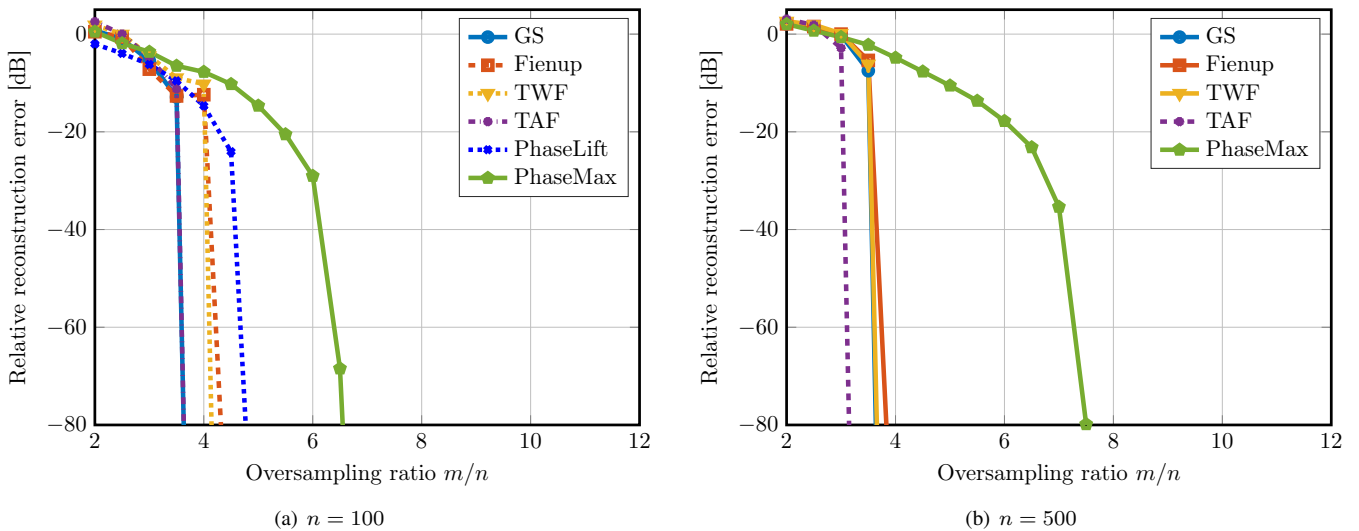

Fig. 3. Comparison of the relative reconstruction error. We use the truncated spectral initializer for Gerchberg-Saxton (GS), Fienup, truncated Wirtinger flow (TWF), truncated amplitude flow (TAF), and PhaseMax. PhaseMax does not achieve exact recovery for the lowest number of measurements among the considered methods, but is convex, operates in the original dimension, and comes with sharp performance guarantees. PhaseLift only terminates in reasonable computation time for $n = 100$.
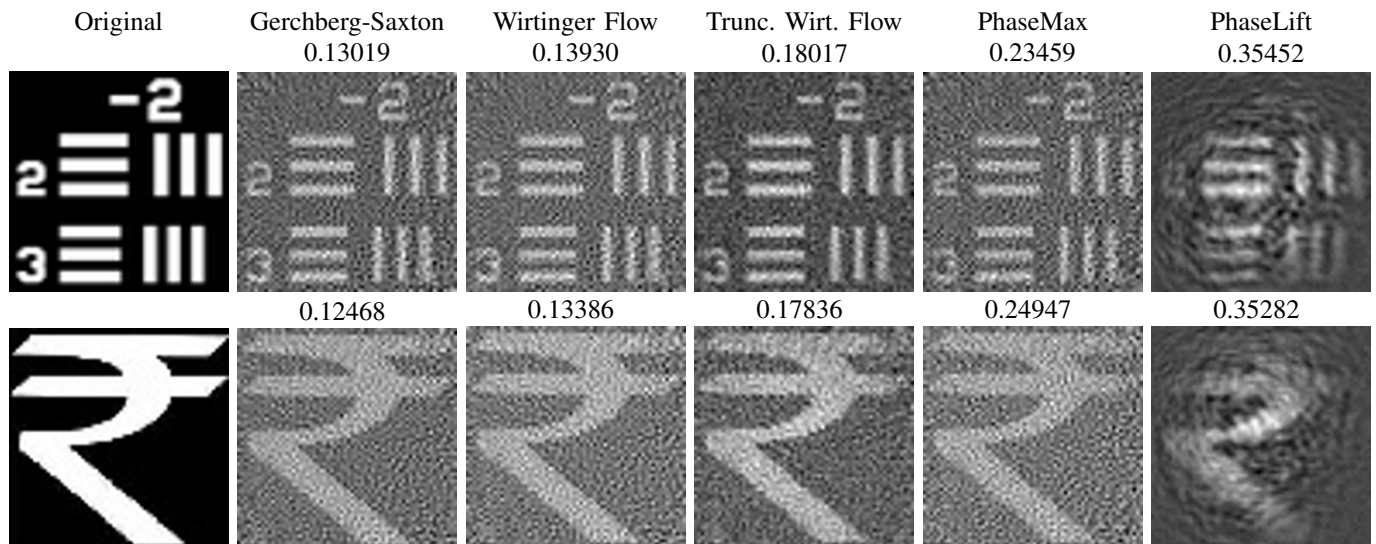
Fig. 4. Reconstruction of two 64×64 masks from empirical phaseless measurements obtained through a diffusive medium. The numbers on top of each of the recovered images denote the relative measurement error $\|| \mathbf{Ax}| - \mathbf{b}\|_2 / \|\mathbf{b}\|_2$ achieved by each method.

methods use the truncated spectral initializer [28]. We also run simulations using the semidefinite relaxation (SDR)-based method PhaseLift [2] implemented via FASTA [53]; this is, together with PhaseCut [21], the only convex alternative to PhaseMax, but lifts the problem to a higher dimension.

Figure 3 reveals that PhaseMax requires larger oversampling ratios m/n to enable faithful signal recovery compared to non-convex phase-retrieval algorithms that operate in the original signal dimension. This is because the truncated spectral initializer requires oversampling ratios of about six or higher to yield sufficiently accurate approximation vectors $\hat{\mathbf{x}}$ that enable PhaseMax to succeed.

2) Comparisons Using Empirical Data: The PhasePack library contains scripts for comparing phase retrieval algorithms using synthetic datasets as well as publicly available real measurements. The datasets contain measurements obtained using the experimental setup described in [54], in which a binary mask is imaged through a diffusive medium.

Figure 4 shows reconstructions from measurements obtained using two different test images. Both images were acquired using phaseless measurements from the same measurement operator $\mathbf{A}$. Reconstructions are shown for the Gerchberg-Saxton [7], Wirtinger flow [27], and truncated Wirtinger flow [28] methods. We also show results of the two convex methods PhaseLift (which lifts the problem dimension) [2] and PhaseMax (the proposed non-lifting relaxation). For each algorithm, Figure 4 shows the recovered images and reports the relative measurement error $\|| \mathbf{Ax}| - \mathbf{b}\|_2 / \|\mathbf{b}\|_2$. All phase retrieval algorithms were initialized using a variant of the spectral method proposed in [35].

We find that the Gerchberg-Saxton algorithm outperforms all other considered methods by a small margin (in terms of measurement error) followed by the Wirtinger flow. On this dataset, measurements seem to be quite valuable, and the truncated Wirtinger flow method (with default truncation parameters as described in [28]) performs slightly worse than the original Wirtinger flow, which uses the full dataset.

PhaseMax produces results that are visually comparable to that of the Wirtinger flow, but with slightly higher relative measurement error.

D. Advantages of PhaseMax

While PhaseMax does not achieve exact reconstruction with the lowest number of measurements, it is convex, operates in the original signal dimension, can be implemented via efficient solvers for Basis Pursuit, and comes with *sharp* performance guarantees that do not sweep constants under the rug (cf. Figure 2). The convexity of PhaseMax enables a natural extension to sparse phase retrieval [55], [56] or other signal priors (e.g., total variation or bounded infinity norm) that can be formulated with convex functions. Such non-differentiable priors cannot be efficiently minimized using simple gradient descent methods (which form the basis of Wirtinger or amplitude flow, and many other methods), but can potentially be solved using standard convex solvers when combined with the PhaseMax formulation.

VIII. CONCLUSIONS

We have proposed a novel, convex phase retrieval algorithm, which we call *PhaseMax*. We have provided accurate bounds on the success probability that depend on the signal dimension, the number of measurements, and the angle between the approximation vector and the true vector. Our analysis covers a broad range of random measurement ensembles and characterizes the impact of general measurement noise on the solution accuracy. We have demonstrated the sharpness of our recovery guarantees and studied the practical limits of PhaseMax via simulations.

There are many avenues for future work. Developing a computationally efficient algorithm for solving PhaseMax (or the related PhaseLamp procedure) accurately and for large problems is a challenging problem. An accurate analysis of more general noise models is left for future work.

APPENDIX A
PROOF OF LEMMA 1

The proof is by induction. As a base case, we note that $r(n, 1) = 2$ for $n \geq 1$ and $r(2, k) = 2k$ for $k \geq 1$. Now suppose we have a sphere $S_{\mathbb{R}}^{n-1}$ in n dimensions sliced by $k - 1$ planes into $r(n, k - 1)$ "original" regions. Consider the effect of adding a kth plane, $\mathcal{P}_k$. Every original region that is intersected by $\mathcal{P}_k$ will split in two, and increase the number of total regions by 1. The increase in the number of regions is then the number of original regions that are intersected by $\mathcal{P}_k$. Equivalently, this is the number of regions formed inside $\mathcal{P}_k$ by the original $k - 1$ planes. Any subset of $k - 1$ planes will have normal vectors that remain linearly independent when projected into $\mathcal{P}_k$. By the induction hypothesis, the number of regions formed inside $\mathcal{P}_k$ is then given by $r(n - 1, k - 1)$. Adding this to the number of original regions yields

$$r(n, k) = r(n, k - 1) + r(n - 1, k - 1).$$

We leave it to the reader to verify that (5) satisfies this recurrence relation and base case.

APPENDIX B
PROOF OF LEMMA 4

In this section, we prove Lemma 4. This Lemma is a direct corollary of the following result of Burgisser, Cucker, and Lotz [47]. For a complete proof of this result, see Theorem 1.1 of [47], and the upper bound on the constant "C" given in Proposition 5.5.

Theorem 8. *Let $m > n \geq 2$. Then the probability of covering the sphere $S_{\mathbb{R}}^{n-1}$ with independent and uniform random caps of central angle $\phi \leq \pi/2$ is bounded by*

$$p_{cover}(m, n, \phi) \geq$$
$$1 - \binom{m}{n} C \int_0^\epsilon (1 - t^2)^{(n^2 - 2n - 1)/2} (1 - \lambda(t))^{m-n} \, dt$$
$$- \frac{1}{2^{m-1}} \sum_{k=0}^{n-1} \binom{m-1}{k}$$

where $\lambda(t) = \frac{V_{n-1}}{V_n} \int_0^{\arccos(t)} \sin^{n-2}(\phi) \, d\phi$, $V_n = Vol(S_{\mathbb{R}}^{n-1}) = \frac{2\pi^{n/2}}{\Gamma(n/2)}$, $C = \frac{n\sqrt{n-1}}{2^{n-1}}$, *and* $\epsilon = \cos(\phi)$.

While Theorem 8 provides a bound on $p_{\text{cover}}(m, n, \phi)$, the formulation of this bound does not provide any intuition of the scaling of $p_{\text{cover}}(m, n, \phi)$ or its dependence on m and n. For this reason, we derive Lemma 4, which is a weaker but more intuitive result. We restate Lemma 4 here for clarity.

Lemma 4. Let $n \geq 9$, and $m > 2n$. Then, the probability of covering the sphere $S_{\mathbb{R}}^{n-1}$ with caps of central angle $\phi \leq \pi/2$ is lower bounded by

$$p_{\text{cover}}(m, n, \phi) \geq$$
$$1 - \frac{(em)^n \sqrt{n-1}}{(2n)^{n-1}} \exp\left(-\frac{\sin^{(n-1)}(\phi)(m-n)}{\sqrt{8n}}\right) \cos(\phi)$$
$$- \exp\left(-\frac{(m - 2n + 1)^2}{2m - 2}\right).$$

Proof. Let us simplify the result of Theorem 8. If we assume $m > 2n$, then Hoeffding's inequality yields

$$\frac{1}{2^{m-1}} \sum_{k=0}^{n-1} \binom{m-1}{k} \leq \exp\left(-\frac{(m - 2n + 1)^2}{2m - 2}\right).$$

Next, we derive a lower bound as follows:

$$\lambda(t) = \frac{\Gamma(n/2)}{\Gamma((n-1)/2)\sqrt{\pi}} \int_0^{\arccos(t)} \sin^{n-2}(\phi) \, d\phi$$
$$\geq \sqrt{(n/2 - 1)/\pi} \int_0^{\arccos(t)} \sin^{n-2}(\phi) \cos(\phi) \, d\phi$$
$$= \sqrt{(n/2 - 1)/\pi} \frac{1}{n - 1} \sin^{n-1} \arccos(t)$$
$$\geq \frac{1}{\sqrt{8n}} (1 - t^2)^{(n-1)/2}.$$

We have used the fact that $\sqrt{(n/2 - 1)/\pi} \frac{1}{n-1} > \frac{1}{\sqrt{8n}}$ for $n \geq 4$, and also the "Wallis ratio" bound $\frac{\Gamma(n/2)}{\Gamma((n-1)/2)} \geq \sqrt{n/2 - 1}$ [57], [58]. Finally, we plug in the inequality $\binom{m}{n} \leq \frac{(em)^n}{n^n}$. We now have

$$\binom{m}{n} C \int_0^\epsilon (1 - t^2)^{(n^2 - 2n - 1)/2} (1 - \lambda(t))^{m-n} \, dt$$
$$\leq \frac{(em)^n \sqrt{n-1}}{(2n)^{n-1}} \int_0^\epsilon (1 - t^2)^{(n^2 - 2n - 1)/2}$$
$$\times \left(1 - \frac{1}{\sqrt{8n}} (1 - t^2)^{(n-1)/2}\right)^{m-n} \, dt.$$

Now we simplify the integral. Using the identity $(1 - x)^a < e^{-ax}$, which holds for $x \leq 1$, we can convert each term in the integrand into an exponential. We do this first with $x = t^2$ and then with $x = \frac{1}{\sqrt{8n}}(1 - t^2)^{(n-1)/2}$ to obtain

$$(1 - t^2)^{(n^2 - 2n - 1)/2} \left(1 - \frac{1}{\sqrt{8n}}(1 - t^2)^{(n-1)/2}\right)^{m-n}$$
$$\leq \exp\left(-\frac{t^2(n^2 - 2n - 1)}{2} - \frac{(1 - t^2)^{(n-1)/2}(m - n)}{\sqrt{8n}}\right). \tag{28}$$

We then apply the Cauchy-Schwarz inequality to get

$$\int_0^\epsilon \exp\left(-\frac{t^2(n^2 - 2n - 1)}{2} - \frac{(1 - t^2)^{(n-1)/2}(m - n)}{\sqrt{8n}}\right) dt$$
$$\leq \left[\int_0^\epsilon \exp\left(-t^2(n^2 - 2n - 1)\right) dt\right]^{1/2}$$
$$\times \left[\int_0^\epsilon \exp\left(-\frac{(1 - t^2)^{(n-1)/2}(m - n)}{\sqrt{2n}}\right) dt\right]^{1/2}$$
$$\leq [\epsilon]^{1/2} \left[\epsilon \exp\left(-\frac{(1 - \epsilon^2)^{(n-1)/2}(m - n)}{\sqrt{2n}}\right)\right]^{1/2}$$
$$= \epsilon \exp\left(-\frac{(1 - \epsilon^2)^{(n-1)/2}(m - n)}{\sqrt{8n}}\right).$$

Replacing the integral with this bound and using the definition $\epsilon = \cos(\phi)$ yields the result. $\square$

REFERENCES

[1] T. Goldstein and C. Studer, "Convex phase retrieval without lifting via PhaseMax," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. International Convention Centre, Sydney, Australia: PMLR, 06–11 Aug 2017, pp. 1273–1281. [Online]. Available: http://proceedings.mlr.press/v70/goldstein17a.html

[2] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Commun. Pure Appl. Math.*, vol. 66, no. 8, pp. 1241–1274, 2013.

[3] S. Bahmani and J. Romberg, "Phase retrieval meets statistical learning theory: A flexible convex relaxation," in *Intl. Conf. on Artificial Intelligence and Statistics (AISTATS)*, May 2017, pp. 252–260.

[4] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Rev.*, vol. 43, no. 1, pp. 129–159, Jul. 2001.

[5] S. Chen and D. Donoho, "Basis pursuit," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, vol. 1, Oct. 1994, pp. 41–44.

[6] C. Studer, W. Yin, and R. G. Baraniuk, "Signal representations with minimum ℓ_∞-norm," in *Proc. Allerton Conf. Commun., Contr., Comput.*, Oct. 2012, pp. 1270–1277.

[7] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik*, vol. 35, pp. 237–246, Aug. 1972.

[8] J. R. Fienup, "Phase retrieval algorithms: a comparison," *Appl. Opt.*, vol. 21, no. 15, pp. 2758–2769, Aug. 1982.

[9] R. W. Harrison, "Phase problem in crystallography," *J. Opt. Soc. Am. A*, vol. 10, no. 5, pp. 1046–1055, May 1993.

[10] J. Miao, T. Ishikawa, Q. Shen, and T. Earnest, "Extending X-ray crystallography to allow the imaging of noncrystalline materials, cells, and single protein complexes," *Ann. Rev. Phys. Chem.*, vol. 59, pp. 387–410, Nov. 2008.

[11] F. Pfeiffer, T. Weitkamp, O. Bunk, and C. David, "Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources," *Nat. Phys.*, vol. 2, no. 4, pp. 258–261, Apr. 2006.

[12] S. S. Kou, L. Waller, G. Barbastathis, and C. J. Sheppard, "Transport-of-intensity approach to differential interference contrast (TI-DIC) microscopy for quantitative phase imaging," *Opt. Lett.*, vol. 35, no. 3, pp. 447–449, Feb. 2010.

[13] H. Faulkner and J. Rodenburg, "Movable aperture lensless transmission microscopy: a novel phase retrieval algorithm," *Phys. Rev. Lett.*, vol. 93, no. 2, Jul. 2004.

[14] J. Holloway, M. S. Asif, M. K. Sharma, N. Matsuda, R. Horstmeyer, O. Cossairt, and A. Veeraraghavan, "Toward long-distance subdiffraction imaging using coherent camera arrays," *IEEE Trans. Comput. Imag.*, vol. 2, no. 3, pp. 251–265, Sept. 2016.

[15] F. Fogel, I. Waldspurger, and A. d'Aspremont, "Phase retrieval for imaging problems," *Math. Prog. Comp.*, vol. 8, no. 3, pp. 311–335, Sept. 2016.

[16] E. J. Candès, E.es, X. Li, and M. Soltanolkotabi, "Phase retrieval from coded diffraction patterns," *Appl. Comput. Harm.. Anal.*, vol. 39, no. 2, pp. 277–299, Sept. 2015.

[17] K. Jaganathan, Y. C. Eldar, and B. Hassibi, "Phase retrieval: An overview of recent developments," *arXiv:1510.07713*, Oct. 2015.

[18] L. Tian and L. Waller, "3D intensity and phase imaging from light field measurements in an LED array microscope," *Optica*, vol. 2, no. 2, pp. 104–111, Feb. 2015.

[19] E. J. Candès and X. Li, "Solving quadratic equations via phaselift when there are about as many equations as unknowns," *Found. Comput. Math.*, vol. 14, no. 5, pp. 1017–1026, Oct. 2014.

[20] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM Rev.*, vol. 57, no. 2, pp. 225–251, Nov 2015.

[21] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," *Math. Prog.*, vol. 149, no. 1-2, pp. 47–81, Feb. 2015.

[22] M. P. Friedlander, I. Macedo, and T. K. Pong, "Gauge optimization and duality," *SIAM Journal on Optimization*, vol. 24, no. 4, pp. 1999–2022, 2014.

[23] A. Y. Aravkin, J. V. Burke, D. Drusvyatskiy, M. P. Friedlander, and K. MacPhee, "Foundations of gauge and perspective duality," *arXiv preprint arXiv:1702.08649*, 2017.

[24] A. Yurtsever, M. Udell, J. A. Tropp, and V. Cevher, "Sketchy decisions: Convex low-rank matrix optimization with optimal storage," in *Intl. Conf. on Artificial Intelligence and Statistics (AISTATS)*, May 2017.

[25] P. Netrapalli, P. Jain, and S. Sanghavi, "Phase retrieval using alternating minimization," in *Adv. Neural Inf. Process. Syst.*, 2013, pp. 2796–2804.

[26] P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate message passing," *IEEE Trans. Sig. Process.*, vol. 63, no. 4, pp. 1043–1055, Feb. 2015.

[27] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval via Wirtinger flow: Theory and algorithms," *IEEE Trans. Inf. Theory*, vol. 61, no. 4, pp. 1985–2007, Feb. 2015.

[28] Y. Chen and E. Candès, "Solving random quadratic systems of equations is nearly as easy as solving linear systems," in *Adv. Neural Inf. Process. Syst.*, 2015, pp. 739–747.

[29] G. Wang, G. B. Giannakis, and Y. C. Eldar, "Solving systems of random quadratic equations via truncated amplitude flow," *arXiv: 1605.08285*, Jul. 2016.

[30] K. Wei, "Solving systems of phaseless equations via kaczmarz methods: A proof of concept study," *Inverse Problems*, vol. 31, no. 12, p. 125008, 2015.

[31] W.-J. Zeng and H. So, "Coordinate descent algorithms for phase retrieval," *arXiv preprint arXiv:1706.03474*, 2017.

[32] Z. Yuan and H. Wang, "Phase retrieval via reweighted Wirtinger flow," *Applied optics*, vol. 56, no. 9, p. 2418, 2017.

[33] P. Chen, A. Fannjiang, and G.-R. Liu, "Phase retrieval with one or two diffraction patterns by alternating projections of the null vector," *arXiv:1510.07379*, Apr. 2015.

[34] Y. M. Lu and G. Li, "Phase transitions of spectral initialization for high-dimensional nonconvex estimation," *arXiv preprint arXiv:1702.06435*, 2017.

[35] M. Mondelli and A. Montanari, "Fundamental limits of weak recovery with applications to phase retrieval," *arXiv preprint: 1708.05932*, 2017.

[36] J. Sun, Q. Qu, and J. Wright, "A geometric analysis of phase retrieval," *arXiv:1602.06664*, Mar. 2016.

[37] P. Hand and V. Voroninski, "An elementary proof of convex phase retrieval in the natural parameter space via the linear program phasemax," *arXiv preprint: 1611.03935*, 2016.

[38] ——, "Compressed sensing from phaseless Gaussian measurements via linear programming in the natural parameter space," *arXiv preprint: 1611.05985*, 2016.

[39] ——, "Corruption robust phase retrieval via linear programming," *arXiv preprint: 1612.03547*, 2016.

[40] A. Aghasi, A. Ahmed, and P. Hand, "BranchHull: convex bilinear inversion from the entrywise product of signals with known signs," *arXiv preprint: 1702.04342*, 2017.

[41] O. Dhifallah, C. Thrampoulidis, and Y. M. Lu, "Phase retrieval via linear programming: Fundamental limits and algorithmic improvements," *arXiv preprint: 1710.05234*, 2017.

[42] O. Dhifallah and Y. M. Lu, "Fundamental limits of phasemax for phase retrieval: A replica analysis," *arXiv preprint arXiv:1708.03355*, 2017.

[43] L. Schläfli, *Gesammelte Mathematische Abhandlungen I*. Springer Basel, 1953.

[44] J. G. Wendel, "A problem in geometric probability," *Math. Scand.*, vol. 11, pp. 109–111, 1962.

[45] E. Gilbert, "The probability of covering a sphere with n circular caps," *Biometrika*, vol. 52, no. 3/4, pp. 323–330, Dec. 1965.

[46] Z. Füredi, "Random polytopes in the d-dimensional cube," *Disc. Comput. Geom.*, vol. 1, no. 4, pp. 315–319, Dec. 1986.

[47] P. Bürgisser, F. Cucker, and M. Lotz, "Coverage processes on spheres and condition numbers for linear programming," *Ann. Probab.*, vol. 38, no. 2, pp. 570–604, 2010.

[48] T. Bendory and Y. C. Eldar, "Non-convex phase retrieval from STFT measurements," *IEEE Trans. Inf. Theory*, Aug. 2017.

[49] J. F. Kenney and E. Keeping, *Mathematics of Statistics, Part 2*. D. Van Nostrand, 1951.

[50] L. Jacques, "A quantized Johnson–Lindenstrauss lemma: The finding of Buffon's needle," *IEEE Trans. Inf. Theory*, vol. 61, no. 9, pp. 5012–5027, Sept. 2015.

[51] F. Qi and Q.-M. Luo, "Bounds for the ratio of two gamma functions—from Wendel's and related inequalities to logarithmically completely monotonic functions," *Banach J. Math. Anal*, vol. 6, no. 2, pp. 132–158, May. 2012.

[52] R. Chandra, Z. Zhong, J. Hontz, V. McCulloch, C. Studer, and T. Goldstein, "PhasePack: A phase retrieval library," *arXiv preprint*, 2017.

[53] T. Goldstein, C. Studer, and R. Baraniuk, "A field guide to forward-backward splitting with a FASTA implementation," *arXiv:1411.3406*, Feb. 2014.

[54] C. A. Metzler, M. K. Sharma, S. Nagesh, R. G. Baraniuk, O. Cossairt, and A. Veeraraghavan, "Coherent inverse scattering via transmission matrices: Efficient phase retrieval algorithms and a public dataset," in *Computational Photography (ICCP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 1–16.

[55] K. Jaganathan, S. Oymak, and B. Hassibi, "Sparse phase retrieval: Convex algorithms and limitations," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2013, pp. 1022–1026.

[56] Y. Shechtman, A. Beck, and Y. C. Eldar, "GESPAR: efficient phase retrieval of sparse signals," *IEEE Trans. Sig. Process.*, vol. 62, no. 4, pp. 928–938, Jan. 2014.

[57] C. Mortici, "New approximation formulas for evaluating the ratio of gamma functions," *Math. Comput. Model.*, vol. 52, no. 1, pp. 425–433, Jul. 2010.

[58] W. Gautschi, "Some elementary inequalities relating to the gamma and incomplete gamma function," *J. Math. Phys.*, vol. 38, no. 1, pp. 77–81, Apr. 1959.