

Phase Retrieval via Matrix Completion*

Emmanuel J. Candès[†]
Yonina C. Eldar[‡]
Thomas Strohmer[§]
Vladislav Voroninski[¶]

Abstract. This paper develops a novel framework for phase retrieval, a problem which arises in X-ray crystallography, diffraction imaging, astronomical imaging, and many other applications. Our approach, called PhaseLift, combines multiple structured illuminations together with ideas from convex programming to recover the phase from intensity measurements, typically from the modulus of the diffracted wave. We demonstrate empirically that a complex-valued object can be recovered from the knowledge of the magnitude of just a few diffracted patterns by solving a simple convex optimization problem inspired by the recent literature on matrix completion. More importantly, we also demonstrate that our noise-aware algorithms are stable in the sense that the reconstruction degrades gracefully as the signal-to-noise ratio decreases. Finally, we introduce some theory showing that one can design very simple structured illumination patterns such that three diffracted figures uniquely determine the phase of the object we wish to recover.

Key words. diffraction, Fourier transform, convex optimization, trace-norm minimization

AMS subject classifications. 49N30, 49N45, 94A20

DOI. 10.1137/151005099

Foreword. We owe SIAM a great debt of gratitude for recognizing our work with the Outstanding Paper Prize, and for having selected it as a SIGEST paper. From the number of citations this paper has received since publication, one can infer that it helped in triggering a wave of new and significant research that we were very far from imagining at the time of the original submission. We warmly thank the community for its interest in this work, and for bringing a renewed and broader understanding of the phase retrieval problem, and more generally of the problem of finding solutions to

*Published electronically May 8, 2015. This paper originally appeared in *SIAM Journal on Imaging Sciences*, Volume 6, Issue 1, 2013, pages 199–225.

<http://www.siam.org/journals/sirev/57-2/100509.html>

[†]Departments of Mathematics and of Statistics, Stanford University, Stanford, CA 94305 (candes@stanford.edu). This author's work was partially supported by NSF via grant CCF-0963835 and the 2006 Waterman Award.

[‡]Department of Electrical Engineering, Technion, Israel Institute of Technology, Haifa 32000, Israel (yonina@ee.technion.ac.il). This author's work was partially supported by the Israel Science Foundation under grant 170/10.

[§]Department of Mathematics, University of California at Davis, Davis, CA 95616 (strohmer@math.ucdavis.edu). This author's work was partially supported by the NSF via grants DMS 0811169 and DMS-1042939, and by DARPA via grant N66001-11-1-4090.

[¶]Department of Mathematics, University of California at Berkeley, Berkeley, CA 94720 (vladv@math.berkeley.edu). This author's work was supported by the Department of Defense (DoD) through the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program.

quadratic equations. Because of the sheer number of works, it is not the scope of this note to review them all; rather, we briefly list a few recent exciting developments.

- First and foremost, PhaseLift now comes with theoretical guarantees stating that the recovery is exact from noiseless data for some random models of sampling vectors. In some cases, the number of samples (of quadratic equations) scales just like the number of unknowns. Similar conclusions hold, via reduction, for another powerful semidefinite relaxation named PhaseCut introduced by d’Aspremont, Mallat, and Waldspurger; PhaseCut is modeled after the famous Goemans–Williamson semidefinite relaxation of the MAXCUT problem. Further, there also is theory explaining the graceful degradation of the PhaseLift reconstruction as the noise level increases.
- In the spirit of compressed sensing, we have also learned a lot about what is and is not possible with SDP relaxations, polarization, and empirical risk approaches when one imposes structure, such as sparsity, on the signal we wish to recover.
- Abstractly, the phase retrieval problem is equivalent to the pure-state quantum tomography problem, and recent theoretical progress on the former has re-ignited interest in the latter and lead to improvements in our basic understanding thereof, for instance, in relation to minimal informational completeness and Wright’s conjecture.
- New algorithmic approaches based to phase retrieval have also flourished. Some are based on combinatorial ideas such as expander graphs as suggested by Alexeev, Mixon, and Bandeira, while others rely on finding exact solutions to nonconvex problems; see, for instance, the work on Wirtinger flows.
- Finally, some of the new phase retrieval algorithms which have been proposed in the last three years or so have found their way in successful applications as varied as speech processing, short-time Fourier transform analysis, ptychography, and blind deconvolution problems.

We would be extremely pleased if our work played a small role in the development of this most impressive literature.

1. Introduction.

1.1. The Phase Retrieval Problem. This paper considers the fundamental problem of recovering a general signal, an image, for example, from the magnitude of its Fourier transform. This problem, also known as *phase retrieval*, arises in many applications and has challenged engineers, physicists, and mathematicians for decades. Its origin comes from the fact that detectors oftentimes can record only the squared modulus of the Fresnel or Fraunhofer diffraction pattern of the radiation that is scattered from an object. In such settings, one cannot measure the phase of the optical wave reaching the detector, and, therefore, much information about the scattered object or the optical field is lost since, as is well known, the phase encodes a lot of the structural content of the image we wish to form.

Historically, one of the first important applications of phase retrieval is X-ray crystallography [58, 33], and today this may very well still be the most important application. Over the last century or so, this field has developed a wide array of techniques for recovering Bragg peaks from missing-phase data. Of course, the phase retrieval problem permeates many other areas of imaging science, and other applications include diffraction imaging [12], optics [74], astronomical imaging [21], and microscopy [55], to name just a few. In particular, X-ray tomography has become an invaluable tool in biomedical imaging to generate quantitative three-dimensional den-

sity maps of extended specimens on the nanoscale [22]. Other subjects where phase retrieval plays an important role are quantum mechanics [65, 20] and even differential geometry [8]. We note that phase retrieval has seen a resurgence in activity in recent years, fueled on the one hand by the desire to image individual molecules and other nanoparticles, and on the other hand by new imaging capabilities: one such recent modality is the availability of new X-ray synchrotron sources that provide extraordinary X-ray fluxes; see, for example, [59, 68, 9, 55, 22]. References and various instances of the phase retrieval problem as well as some theoretical and numerical solutions can be found in [37, 49, 43].

There are many ways in which one can pose the phase retrieval problem, for instance, depending upon whether one assumes a continuous- or discrete-space model for the signal. In this paper, we consider finite length signals (one-dimensional (1D) or multidimensional) for simplicity and because numerical algorithms ultimately operate with digital data. To fix ideas, suppose we have a 1D signal $x = (x[0], x[1], \dots, x[n-1]) \in \mathbb{C}^n$ and write its Fourier transform as

$$(1.1) \quad \hat{x}[\omega] = \frac{1}{\sqrt{n}} \sum_{0 \leq t < n} x[t] e^{-i2\pi\omega t/n}, \quad \omega \in \Omega.$$

Here, Ω is a grid of sampled frequencies, and an important special case is $\Omega = \{0, 1, \dots, n-1\}$ so that the mapping is the classical unitary discrete Fourier transform (DFT).¹ The phase retrieval problem consists in finding x from the magnitude coefficients $|\hat{x}[\omega]|$, $\omega \in \Omega$. When Ω is the usual frequency grid as above and without further information about the unknown signal x , this problem is ill-posed since there are many different signals whose Fourier transforms have the same magnitude. Clearly, if x is a solution to the phase retrieval problem, then (1) cx for any scalar $c \in \mathbb{C}$ obeying $|c| = 1$ is also a solution, (2) the “mirror function” or time-reversed signal $\bar{x}[-t \bmod n]$, where $t = 0, 1, \dots, n-1$, is also a solution, and (3) the shifted signal $x[t-a \bmod n]$ is also a solution. From a physical viewpoint these “trivial associates” of x are acceptable ambiguities. But in general infinitely many solutions can be obtained from $\{|\hat{x}[\omega]| : \omega \in \Omega\}$ beyond these trivial associates [67].

We wish to make clear that our claims apply only to the discrete setting considered in this problem. It would, of course, be possible to develop a framework for continuous objects as well. This would involve specifying a sampling mechanism and a finite model suitable for finite computations. If, as empirically demonstrated in this paper, the discrete model is robust vis-à-vis small errors, then accurate recovery of continuous objects would be possible since small approximation/discretization errors would lead to small reconstruction errors.

1.2. Main Approaches to Phase Retrieval. Holographic techniques are among the more popular methods that have been proposed to measure the phase of the optical wave. While holographic techniques have been successfully applied in certain areas of optical imaging, they are generally difficult to implement in practice [23]. Hence, the development of algorithms for signal recovery from magnitude measurements is still a very active field of research. Existing methods for phase retrieval rely on all kinds of a priori information about the signal, such as positivity, atomicity, support constraints, and real-valuedness [27, 28, 51, 19]. Direct methods [34] are limited in their applicability to small-scale problems due to their large computational complexity.

¹For later reference, we denote the Fourier transform operator by F and the inverse Fourier transform by F^{-1} .

Oversampling in the Fourier domain has been proposed as a means of mitigating the nonuniqueness of the phase retrieval problem [54]. While oversampling offers no benefit for most 1D signals, the situation is more favorable for multidimensional signals, where it has been shown that twofold oversampling in each dimension almost always yields uniqueness for finitely supported, real-valued, and nonnegative signals [11, 35, 67]. In other words, a digital image of the form $x = \{x[t_1, t_2]\}$ with $0 \leq t_1 < n_1$ and $0 \leq t_2 < n_2$, whose Fourier transform is given by

$$(1.2) \quad \hat{x}[\omega_1, \omega_2] = \frac{1}{\sqrt{n_1 n_2}} \sum x[t_1, t_2] e^{-i2\pi(\omega_1 t_1/n_1 + \omega_2 t_2/n_2)},$$

is usually uniquely determined from the values of $|\hat{x}[\omega_1, \omega_2]|$ on the oversampled grid $\omega = (\omega_1, \omega_2) \in \Omega = \Omega_1 \times \Omega_2$ in which $\Omega_i = \{0, 1/2, 1, 3/2, \dots, n_i + 1/2\}$. (In other words, if we think of $(1/n_1, 1/n_2)$ as some Nyquist frequency, then we would need to sample at a rate that is at least twice this Nyquist frequency.) This holds, provided that x has proper spatial support and is real-valued and nonnegative.

As pointed out in [49], these uniqueness results do not say anything about how a signal can be recovered from its intensity measurements, or about the robustness and stability of commonly used reconstruction algorithms—a fact we shall make very clear in what follows. In fact, theoretical uniqueness conditions do not readily translate into numerical methods and, as a result, the algorithmic and practical aspects of the phase retrieval problem (from noisy data) still pose significant challenges.

By and large, the most popular methods for phase retrieval from oversampled data are alternating projection algorithms pioneered by Gerchberg and Saxton [30] and Fienup [27, 28]. These methods often require careful exploitation of signal constraints and delicate parameter selection to increase the likelihood of convergence to a correct solution [62, 51, 19, 50]. We describe the simplest realization of a widely used alternating projection approach [73], which assumes support constraints in the spatial domain and oversampled measurements in the frequency domain. With T being a *known* subset containing the support of the signal x ($\text{supp}(x) \subset T$) and Fourier magnitude measurements $\{y[\omega]\}_{\omega \in \Omega}$ with $y[\omega] = |\hat{x}[\omega]|$, the method works as follows:

1. *Initialization:* Choose an initial guess x_0 and set $z_0[\omega] = y[\omega] \frac{\hat{x}_0[\omega]}{|\hat{x}_0[\omega]|}$ for $\omega \in \Omega$.
2. *Loop:* For $k = 1, 2, \dots$, inductively define

$$(1) \quad x_k[t] = \begin{cases} (F^{-1} z_{k-1})[t] & \text{if } t \in T, \\ 0 & \text{else;} \end{cases}$$

$$(2) \quad z_k[\omega] = y[\omega] \frac{\hat{x}_k[\omega]}{|\hat{x}_k[\omega]|} \quad \text{for } \omega \in \Omega$$

until convergence.

While this algorithm is simple to implement and amenable to additional constraints such as the positivity of x , its convergence remains problematic. Projection algorithms onto convex sets are well understood [10, 32, 75, 3]. However, the set $\{z : |\hat{z}[\omega]| = |\hat{x}[\omega]|\}$ is not convex, and, therefore, the algorithm is not known to converge in general or even to give a reasonable solution [44, 3, 49]. Some progress toward understanding the convergence behavior of certain alternating projection methods has been made in [48]. Good numerical results have been reported in certain oversampling settings, but they nevertheless appear to be somewhat problematic in light of our numerical experiments from section 4. Reference [52] points out that oversampling is not always practically feasible as certain experimental geometries allow only for sub-Nyquist sampling; an example is the Bragg sampling from periodic crystalline

structures. Alternating projection algorithms may be more competitive when applied within the framework of multiple structured illuminations, as proposed in this paper, instead of oversampling. Another direction of investigation is to utilize sparsity of the signal; see [52, 47, 69]. Here, the signal is known to have only a few nonzero coefficients, but the locations of the nonzero coefficients (that is, the support of the signal) are not known a priori.

In a different direction, a frame-theoretic approach to phase retrieval has been proposed in [2, 1], where the authors derive various necessary and sufficient conditions for the uniqueness of the solution, as well as various numerical algorithms. While theoretically appealing, the practical applicability of these results is limited by the fact that very specific types of measurements are required, which cannot be realized in most applications of interest.

To summarize our discussion, we have seen many methods which all represent some important attempts to find efficient algorithms, and work well in certain situations. However, these techniques do not always provide a consistent and robust result.

1.3. PhaseLift—a Novel Methodology. This paper develops a novel methodology for phase retrieval based on a rigorous and flexible numerical framework. Whereas most of the existing methods seek to overcome nonuniqueness by imposing additional constraints on the signal, we pursue a different direction by assuming no constraints at all on the signal. There are two main components to our approach.

- *Multiple structured illuminations.* We suggest collecting several diffraction patterns providing “different views” of the sample or specimen. This can be accomplished in a number of ways: for instance, by modulating the light beam falling onto the sample or by placing a mask right after the sample; see section 2 for details. Taking multiple diffraction patterns usually yields uniqueness as discussed in section 3.

The concept of using multiple measurements as an attempt to resolve the phase ambiguity for diffraction imaging is, of course, not new, and was suggested in [60]. Since then, a variety of methods have been proposed to carry out these multiple measurements; depending on the particular application, these may include the use of various gratings and/or masks, the rotation of the axial position of the sample, and the use of defocusing implemented in a spatial light modulator; see [23] for details and references. Other approaches include *ptychography*, an exciting field of research, where one records several diffraction patterns from overlapping areas of the sample; see [66, 70] and references therein.

- *Formulation of phase recovery as a matrix completion problem.* We suggest (1) lifting up the problem of recovering a vector from quadratic constraints into that of a recovering of a rank-one matrix from affine constraints, and (2) relaxing the combinatorial problem into a convenient convex program. Since the lifting step is fundamental to our approach, we will refer to the proposed numerical framework as *PhaseLift*. The price we pay for trading the nonconvex quadratic constraints for convex constraints is that we must deal with a highly underdetermined problem. However, recent advances in the areas of compressive sensing and matrix completion have shown that such convex approximations are often exact.

Although our algorithmic framework appears to be novel for phase retrieval, the idea of solving problems involving nonconvex quadratic constraints by

semidefinite relaxations has a long history in optimization; see [7] and the references therein, and section 1.4 below for more discussion.

The goal of this paper is to demonstrate that, taken together, multiple coded illuminations and convex programming (trace-norm minimization) provide a powerful new approach to phase retrieval. Further, a significant aspect of our methodology is that our systematic optimization framework offers a principled way of dealing with noise and makes it easy to handle various statistical noise models. This is important because in practice, measurements are always noisy. In fact, our framework can be understood as an elaborate regularized maximum likelihood method. Finally, our framework can also include a priori knowledge about the signal that can be formulated or relaxed as convex constraints.

1.4. Precedents. At the abstract level, the phase retrieval problem is that of finding $x \in \mathbb{C}^n$ obeying quadratic equations of the form $|\langle a_k, x \rangle|^2 = b_k$. Casting such quadratic constraints as affine constraints about the matrix variable $X = xx^*$ has been widely used in the optimization literature for finding good bounds on a number of quadratically constrained quadratic problems (QCQP). Indeed, solving the general case of a QCQP is known to be an NP-hard problem since it includes the family of Boolean linear programs [7]. The approach usually consists in finding a relaxation of the QCQP using semidefinite programming (SDP), for instance, via Lagrangian duality. An important example of this strategy is Max Cut, an NP-hard problem in graph theory which can be formulated as a QCQP. In a celebrated paper, Goemans and Williamson introduced a relaxation [31] for this problem, which lifts or linearizes a nonlinear, nonconvex problem to the space of symmetric matrices. Although there are evident connections to our work, our relaxation is quite different from these now-standard techniques.

The idea of linearizing the phase retrieval problem by reformulating it as a problem of recovering a matrix from linear measurements can be found in [1]. While this reference also contains some intriguing numerical recovery algorithms, their practical relevance for most applications is limited by the fact that the proposed measurement matrices either require a very specific algebraic structure which does not seem to be compatible with the physical properties of diffraction, or the number of measurements is proportional to the square of the signal dimension, which is not feasible in most applications.

In terms of framework, the closest approach is the paper [18], in which the authors use a matrix completion approach for array imaging from intensity measurements. Although this paper executes a relaxation similar to ours, there are some differences. We present a “noise-aware” framework, which makes it possible to account for a variety of noise models in a systematic way. Moreover, our emphasis is on a novel combination of structured illuminations and convex programming, which seems to bear great potential.

2. Methodology.

2.1. Structured Illumination. Suppose $x = \{x[t]\}$ is the object of interest (t may be a 1D or multidimensional index). In this paper, we shall discuss illumination schemes collecting the diffraction pattern of the modulated object $w[t]x[t]$, where the waveforms or patterns $w[t]$ may be selected by the user. There are many ways in which this can be implemented in practice, and we discuss just a few.

- *Masking.* One possibility is to modify the phase front after the sample by inserting a mask or a phase plate; see [45], for example. A schematic layout

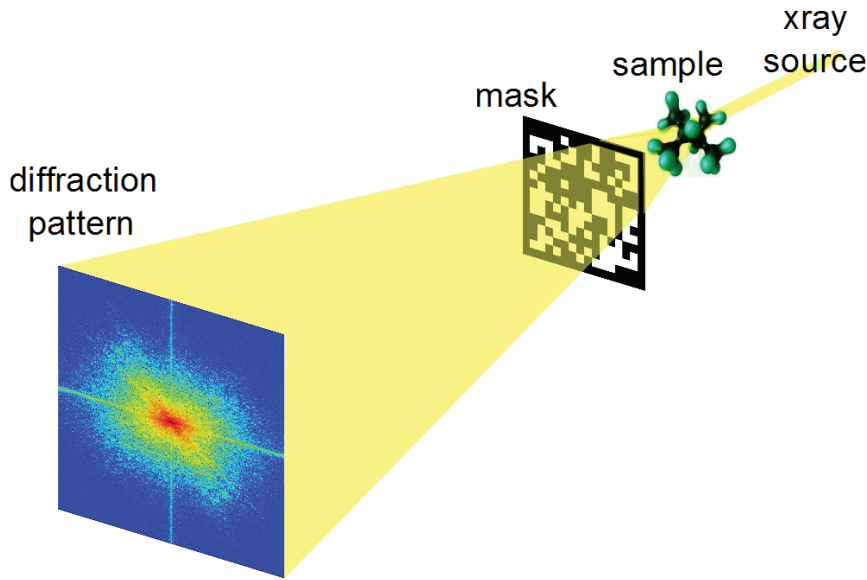


Fig. 1 A typical setup for structured illuminations in diffraction imaging using a phase mask.

is shown in Figure 1. In [40], the sample is scanned by shifting the phase plate as in ptychography (discussed below); the difference is that one scans the known phase plate rather than the object being imaged.

- *Optical grating.* Another standard approach would be to change the profile or modulate the illuminating beam, which can easily be accomplished by the use of optical gratings [46]. A simplified representation would look similar to the scheme depicted in Figure 1, with a grating instead of the mask (the grating could be placed before or after the sample).
- *Ptychography.* Here, one measures multiple diffraction patterns by scanning a finite illumination on an extended specimen [66, 70]. In this setup, it is common to maintain a substantial overlap between adjacent illumination positions.
- *Oblique illuminations.* One can use illuminating beams that hit the sample at a user specified angle [24]; see Figure 2 for a schematic illustration of this approach. One can also imagine having multiple simultaneous oblique illuminations.

As is clear, there is no shortage of options, and one might prefer solutions which require generating as few diffraction patterns as possible for stable recovery.

2.2. Lifting. Suppose we have $x_0 \in \mathbb{C}^n$ or $\mathbb{C}^{n_1 \times n_2}$ (or some higher-dimensional version) about which we have quadratic measurements of the form

$$(2.1) \quad \mathbb{A}(x_0) = \{|\langle a_k, x_0 \rangle|^2 : k = 1, 2, \dots, m\}.$$

In the setting where we would collect the diffraction pattern of $w[t]x_0[t]$ as discussed earlier, the waveform $a_k[t]$ can be written as

$$(2.2) \quad a_k[t] \propto w[t]e^{i2\pi \langle \omega_k, t \rangle};$$




Fig. 2 A typical setup for structured illuminations in diffraction imaging using oblique illuminations. The left image shows direct (on-axis) illumination, and the right image corresponds to oblique (off-axis) illumination.

here, ω_k is a frequency value so that $a_k[t]$ is a patterned complex sinusoid. One can assume for convenience that the normalizing constant is such that a_k is unit normed, i.e., $\|a_k\|_2^2 = \sum_t |a_k[t]|^2 = 1$. Phase retrieval is then the feasibility problem

$$(2.3) \quad \begin{array}{ll} \text{find} & x \\ \text{obeying} & \mathbb{A}(x) = \mathbb{A}(x_0) := b. \end{array}$$

As is well known, quadratic measurements can be lifted up and interpreted as linear measurements about the rank-one matrix $X = xx^*$. Indeed,

$$|\langle a_k, x \rangle|^2 = \text{Tr}(x^* a_k a_k^* x) = \text{Tr}(a_k a_k^* x x^*) := \text{Tr}(A_k X),$$

where A_k is the rank-one matrix $a_k a_k^*$. In what follows, we will let \mathcal{A} be the linear operator mapping positive semidefinite matrices into $\{\text{Tr}(A_k X) : k = 1, \dots, m\}$. Hence, the phase retrieval problem is equivalent to

$$(2.4) \quad \begin{array}{ll} \text{find} & X \\ \text{subject to} & \mathcal{A}(X) = b, \\ & X \succeq 0, \\ & \text{rank}(X) = 1 \end{array} \quad \Leftrightarrow \quad \begin{array}{ll} \text{minimize} & \text{rank}(X) \\ \text{subject to} & \mathcal{A}(X) = b, \\ & X \succeq 0. \end{array}$$

Upon solving the left-hand side of (2.4), we would factorize the rank-one solution X as xx^* , hence finding solutions to the phase retrieval problem. Note that the equivalence between the left- and right-hand sides of (2.4) is straightforward since by definition $b = \mathbb{A}(x_0) = \mathcal{A}(x_0 x_0^*)$ and there exists a rank-one solution. Therefore, our problem is a rank-minimization problem over an affine slice of the positive semidefinite cone. As such, it falls into the realm of low-rank *matrix completion* or *matrix recovery*, a class of optimization problems that has gained tremendous attention in recent years; see,

e.g., [64, 13, 15]. Just as in matrix completion, the linear system $\mathcal{A}(X) = b$, with the unknown in the positive semidefinite cone, is highly underdetermined. For instance, suppose our signal x_0 has n complex unknowns. Then we may imagine collecting six diffraction patterns with n measurements for each (no oversampling). Thus $m = 6n$, whereas the dimension of the space of $n \times n$ Hermitian matrices over the reals is n^2 , which is obviously much larger.

We are, of course, interested in low-rank solutions, and this makes the search feasible. This also raises an important question: what is the minimal number of diffraction patterns needed to recover x , whatever x may be? Since each pattern yields n real-valued coefficients and x has n complex-valued unknowns, the answer is at least two. Further, in the context of quantum state tomography, Theorem II in [29] shows that one needs at least $3n - 2$ intensity measurements to guarantee uniqueness, hence suggesting an absolute minimum of three diffraction patterns. Are three patterns sufficient? For some answers to this question, see section 3.

2.3. Recovery via Convex Programming. The rank-minimization problem (2.4) is NP-hard. We propose using the trace norm as a convex surrogate [5, 53] for the rank functional, giving the familiar SDP (and a crucial component of PhaseLift),

$$(2.5) \quad \begin{array}{ll} \text{minimize} & \text{Tr}(X) \\ \text{subject to} & \mathcal{A}(X) = b, \\ & X \succeq 0; \end{array}$$

here and below $X \succeq 0$ means that X is Hermitian positive semidefinite. This problem is convex, and there exists a wide array of numerical solvers including the popular Nesterov’s accelerated first-order method [61]. As far as the relationship between (2.4) and (2.5) is concerned, the matrix \mathcal{A} in most diffraction imaging applications is not known to obey any of the conditions derived in the literature [13, 15, 64] that would guarantee a formal equivalence between the two programs. Nevertheless, the formulation (2.5) enjoys great empirical performance as demonstrated in section 4. Furthermore, shortly after this work was submitted, it was shown in [14] that for measurement vectors a_k sampled independently and uniformly at random on the unit sphere, PhaseLift can recover x exactly (up to a global phase factor) with high probability, provided that the number of measurements is on the order of $n \log n$.

We mentioned earlier that measurements are typically noisy and that our formulation allows for a principled approach to dealing with this issue for a variety of noise models. Suppose the measurement vector $\{b_k\}$ is sampled from a probability distribution $p(\cdot; \mu)$, where $\mu = \mathbb{A}(x_0)$ is the vector of noiseless values, $\mu_k = |\langle a_k, x_0 \rangle|^2$. Then a classical fitting approach simply consists of maximizing the likelihood:

$$(2.6) \quad \begin{array}{ll} \text{maximize} & p(b; \mu) \\ \text{subject to} & \mu = \mathbb{A}(x) \end{array}$$

with optimization variables μ and x . (A more concise description is to find x such that $p(b; \mathbb{A}(x))$ is maximum.) Using the lifting technique and the monotonicity of the logarithm, an equivalent formulation is

$$\begin{array}{ll} \text{minimize} & -\log(p(b; \mu)) \\ \text{subject to} & \mu = \mathcal{A}(X), \\ & X \succeq 0, \text{rank}(X) = 1. \end{array}$$

This is, of course, not tractable and our convex formulation suggests solving instead

$$(2.7) \quad \begin{array}{ll} \text{minimize} & -\log p(b; \mu) + \lambda \text{Tr}(X) \\ \text{subject to} & \mu = \mathcal{A}(X), \\ & X \succeq 0 \end{array}$$

with optimization variables μ and X (i.e., find $X \succeq 0$ such that $-\log p(b; \mathcal{A}(X)) + \lambda \text{Tr}(X)$ is minimum). Above, λ is a positive scalar, and, hence, our approach is a penalized or regularized maximum likelihood method, which trades off between goodness and complexity of the fit. When the likelihood is log-concave, problem (2.7) is convex and solvable. We give two examples for concreteness:

- *Poisson data.* Suppose that $\{b_k\}$ is a sequence of independent samples from the Poisson distributions $\text{Poi}(\mu_k)$. The Poisson log-likelihood for independent samples has the form $\sum_k b_k \log \mu_k - \mu_k$ (up to an additive constant factor), and thus our problem becomes

$$\begin{array}{ll} \text{minimize} & \sum_k [\mu_k - b_k \log \mu_k] + \lambda \text{Tr}(X) \\ \text{subject to} & \mu = \mathcal{A}(X), \\ & X \succeq 0. \end{array}$$

- *Gaussian data.* Suppose that $\{b_k\}$ is a sequence of independent samples from the Gaussian distribution with mean μ_k and variance σ_k^2 (or is well approximated by Gaussian variables). Then our problem becomes

$$\begin{array}{ll} \text{minimize} & \sum_k \frac{1}{2\sigma_k^2} (b_k - \mu_k)^2 + \lambda \text{Tr}(X) \\ \text{subject to} & \mu = \mathcal{A}(X), \\ & X \succeq 0. \end{array}$$

If Σ is a diagonal matrix with diagonal elements σ_k^2 , this can be written as

$$\begin{array}{ll} \text{minimize} & \frac{1}{2} [b - \mathcal{A}(X)]^* \Sigma^{-1} [b - \mathcal{A}(X)] + \lambda \text{Tr}(X) \\ \text{subject to} & X \succeq 0. \end{array}$$

Both formulations are, of course, convex and in both cases one recovers the noiseless trace-minimization problem (2.5) as $\lambda \rightarrow 0^+$.

In addition, it is straightforward to include further constraints frequently discussed in the phase retrieval literature such as real-valuedness, positivity, and atomicity. Suppose the support of x is known to be included in a set T known a priori. Then we would add the linear constraint

$$X_{ij} = 0, \quad (i, j) \notin T \times T.$$

(Algorithmically, one would simply work with a reduced-size matrix.) Suppose we would like to enforce real-valuedness; then we simply assume that X is real-valued and positive semidefinite. Finally positivity can be expressed as linear inequalities

$$X_{ij} \geq 0.$$

Of course, many other types of constraints can be incorporated into this framework, which provides appreciable flexibility.

2.4. PhaseLift with Reweighting. The trace norm promotes low-rank solutions, and this is why it is often used as a convex proxy for the rank. However, it is possible to further promote low-rank solutions by solving a sequence of weighted trace-norm problems, a technique which has been shown to provide even more accurate solutions [26, 16]. The reweighting scheme works like this: choose $\varepsilon > 0$; start with $W_0 = I$ and for $k = 0, 1, \dots$, inductively define X_k as the optimal solution to

$$(2.8) \quad \begin{aligned} & \text{minimize} && \text{Tr}(W_k X) \\ & \text{subject to} && \mathcal{A}(X) = b, \\ & && X \succeq 0, \end{aligned}$$

and update the “weight matrix” as

$$W_{k+1} = (X_k + \varepsilon I)^{-1}.$$

The algorithm terminates on convergence or when the iteration count k attains a specified maximum number of iterations k_{\max} . One can see that the first step of this procedure is precisely (2.5); after this initial step, the algorithm proceeds to solve a sequence of trace-norm problems in which the matrix weights W_k are roughly the inverse of the current guess.

As explained in the literature [26, 25], this reweighting scheme can be viewed as attempting to solve

$$(2.9) \quad \begin{aligned} & \text{minimize} && f(X) = \log(\det(X + \varepsilon I)) \\ & \text{subject to} && \mathcal{A}(X) = b, \\ & && X \succeq 0, \end{aligned}$$

by minimizing the tangent approximation to f at each iterate; that is to say, at step k , (2.8) is equivalent to minimizing $f(X_{k-1}) + \langle \nabla f(X_{k-1}), X - X_{k-1} \rangle$ over the feasible set. (As for the trace, the function $\log \det(X + \varepsilon I)$ serves as a surrogate for the rank functional.) This can also be applied to noise-aware variants, where one would simply replace the objective functional in (2.7) with

$$-\log p(b; \mu) + \lambda \text{Tr}(W_k X)$$

at each step and update W_k in exactly the same way as before.

The log-det functional is closer to the rank functional than the trace norm. In fact, minimizing this functional solves the phase retrieval problem as incorporated into the following theorem.

THEOREM 2.1. *Suppose that \mathbb{A} is one to one and that the identity matrix I is in the span of the sensing matrices A_k . Then the unique solution of the phase retrieval problem (2.3) is also the unique minimizer to (2.9)—up to global phase.² This holds for all values of $\varepsilon > 0$. The same conclusion holds without the inclusion assumption, provided that one modifies the reweighting scheme and substitutes the objective function $f(X)$ in (2.9) with $f(RXR)$, where $R = (\sum_k A_k)^{1/2}$.*

Since the reweighting algorithm is a good heuristic for solving (2.9), we potentially have an interesting and tractable method for phase retrieval. It is not a perfect

²Quadratic measurements can, of course, never distinguish between x and cx in which $c \in \mathbb{C}$ has unit norm. When the solution is unique up to a multiplication with such a scalar, we say that unicity holds up to global phase. From now on, whenever we talk about unicity, it is implied up to global phase.

heuristic, however, as we cannot expect this procedure to always find the global minimum since the objective functional is concave.

The assumption that the identity matrix is in the span of the A_k 's holds whenever the modulus of the Fourier transform of the sample is measured. Indeed, if $|Fx|^2$ is observed, then, letting $\{f_k^*\}$ be the rows of F , we have $\sum_k f_k f_k^* = I$.

Proof. Our assumption implies that for any feasible X , $\text{Tr}(X) = \sum_k h_k \text{Tr}(A_k X) = \sum_k h_k b_k$ is fixed. Assume without loss of generality that feasible points obey $\text{Tr}(X) = 1$ (if $\text{Tr}(X) = 0$, then the unique solution is $X = 0$). If x_0 is the unique solution to phase retrieval (up to global phase), then $X_0 = x_0 x_0^*$ is the only rank-one feasible point. We thus need to show that any feasible X with $\text{rank}(X) > 1$ obeys $f(X) > f(X_0)$, a fact which follows from the strong concavity of f (of the logarithm). Let $X = \sum_j \lambda_j u_j u_j^*$ be any eigenvalue decomposition of a feasible point. Then

$$f(X) = \log(\det(\varepsilon I + X)) = \sum_j \log(\varepsilon + \lambda_j),$$

and it follows from the strict concavity of the log that

$$\sum_j \log(\varepsilon + \lambda_j) > \sum_j \lambda_j \log(\varepsilon + 1) + (1 - \lambda_j) \log \varepsilon = \log(\varepsilon + 1) + (n - 1) \log \varepsilon.$$

The first strict inequality holds unless X is rank one, in which case we have equality. The equality follows from $\sum_j \lambda_j = \text{Tr}(X) = 1$. Since the right-hand side is none other than $f(X_0)$, the theorem is established.

For the second part, set $B_k = R^{-1} A_k R^{-1}$ and consider a new data problem with constraints $\{X : \text{Tr}(B_k X) = b_k, X \succeq 0\}$. Now X is feasible for our problem if and only if RXR is feasible for this new problem. (This is because the mapping $X \mapsto RXR$ preserves the positive semidefinite cone.) Now suppose that x_0 is the solution to phase retrieval and set $X_0 = x_0 x_0^*$ as before. Since I is in the span of the sensing matrices B_k , we have just learned that for all $RXR \neq RX_0R$ and X feasible for our problem,

$$f(RXR) > f(RX_0R).$$

This concludes the proof. \square

3. Theory. Our PhaseLift framework poses two main theoretical questions:

1. When do multiple diffracted images imply unicity of the solution?
2. When does our convex heuristic succeed in recovering the unique solution to the phase retrieval problem?

Developing comprehensive answers to these questions constitutes a whole research program, which clearly is beyond the scope of this work. In this paper, we shall limit ourselves to introducing some theoretical results showing simple ways of designing diffraction patterns which give unicity. Our focus is on getting uniqueness from a very limited number of diffraction patterns. For example, we shall demonstrate that in some cases three diffraction images are sufficient for perfect recovery. Thus, we give below partial answers to the first question and will address the second in a later publication.

A frequently discussed approach to retrieving phase information uses a technique from holography. Roughly speaking, the idea is to let the signal of interest x interfere with a known reference beam y . One typically measures $|x + y|^2$ and $|x - iy|^2$, and precise knowledge of y allows us, in principle, to recover the amplitude and phase of x .

Holographic techniques are hard to implement [23] in practice. Instead, we propose using a modulated version of the signal itself as a reference beam, which in some cases may be easier to implement.

To discuss this idea, we need to introduce some notation. For a complex signal $z \in \mathbb{C}^n$, we let $|z|^2$ be the nonnegative real-valued n -dimensional vector containing the squared magnitudes of z . Suppose first that x is a 1D signal $(x[0], x[1], \dots, x[n-1])$ and that F_n is the $n \times n$ unitary DFT. In this section, we consider taking $3n$ real-valued measurements of the form

$$(3.1) \quad \mathbb{A}(x) = \{|F_n x|^2, |F_n(x + D^s x)|^2, |F_n(x - iD^s x)|^2\},$$

where D is the modulation,

$$D = \text{diag}(\{e^{i2\pi t/n}\}_{0 \leq t \leq n-1}),$$

and s is a nonnegative integer. These measurements can be obtained by illuminating the sample with the three light fields 1 , $1 + e^{i2\pi st/n}$, and $1 + e^{i2\pi(st/n-1/4)}$. We show below that these $3n$ measurements are generally sufficient for perfect recovery.

THEOREM 3.1. *Suppose that the DFT of $x \in \mathbb{C}^n$ does not vanish. Then x can be recovered up to global phase from the $3n$ real numbers $\mathbb{A}(x)$ (3.1) if and only if s is prime with n . In particular, assuming primality, if the trace-minimization program (2.5) or the iteratively reweighted algorithm returns a rank-one solution, then this solution is exact.*

Conversely, if the DFT vanishes at two frequency points k and k' obeying $k - k' \neq s \pmod n$, then recovery is not possible from the $3n$ real numbers (3.1).

The proof of this theorem is constructive, and we give a simple algorithm that achieves perfect reconstruction. Further, one can use masks to scramble the Fourier transform so as to make sure it does not vanish. Suppose, for instance, that we collect

$$\mathbb{A}(Wx), \quad W = \text{diag}(\{z[t]\}_{0 \leq t \leq n-1}),$$

where the $z[t]$'s are independent and identically distributed (i.i.d.) $\mathcal{N}(0, 1)$. Then since the Fourier transform of $z[t]x[t]$ does not vanish with probability one, we have the following corollary.

COROLLARY 3.2. *Assume that s is prime with n . Then with probability one, x can be recovered up to global phase from the $3n$ real numbers $\mathbb{A}(Wx)$, where W is the diagonal matrix with Gaussian entries above.*

Of course, one could derive similar results by scrambling the Fourier transform with the aid of other types of masks, e.g., binary masks. We do not pursue such calculations.

We now turn our attention to the situation in higher dimensions and will consider the two-dimensional (2D) case (higher dimensions are treated in the same way). Here, we have a discrete signal $x[t_1, t_2] \in \mathbb{C}^{n_1 \times n_2}$ about which we take the $3n_1 n_2$ measurements

$$(3.2) \quad \{|\mathcal{F}_{n_1 \times n_2} x|^2, |\mathcal{F}_{n_1 \times n_2}(x + \mathcal{D}^s x)|^2, |\mathcal{F}_{n_1 \times n_2}(x - i\mathcal{D}^s x)|^2\}, \quad s = (s_1, s_2);$$

$\mathcal{F}_{n_1 \times n_2}$ is the 2D unitary Fourier transform defined by (1.2) in which the frequencies belong to the 2D grid $\{0, 1, \dots, n_1 - 1\} \times \{0, 1, \dots, n_2 - 1\}$, s is a pair of nonnegative integers, and \mathcal{D}^s is the modulation

$$[\mathcal{D}^s x][t_1, t_2] = e^{i2\pi s_1 t_1/n_1} e^{i2\pi s_2 t_2/n_2} x[t_1, t_2].$$

With these definitions, we have the following result.

THEOREM 3.3. *Suppose that the DFT of $x \in C^{n_1 \times n_2}$ does not vanish. Then x can be recovered up to global phase from the $3n_1n_2$ real numbers (3.2) if and only if s_1 is prime with n_1 , s_2 is prime with n_2 , and n_1 is prime with n_2 . Under these assumptions, if the trace-minimization program (2.5) or the iteratively reweighted algorithm returns a rank-one solution, then this solution is exact.*

Again, one can apply a random mask to turn this statement into a probabilistic statement holding either with probability one or with very large probability depending upon the mask that is used.

One can always choose s_1 and s_2 such that they are prime with n_1 and n_2 , respectively. The last condition may be less friendly, but one can decide to pad one dimension with zeros to guarantee primality. This is equivalent to a slight oversampling of the DFT along one direction. An alternative is to take $5n_1n_2$ measurements in which we modulate the signal horizontally and then vertically; that is to say, we modulate with $s = (s_1, 0)$ and then with $s = (0, s_2)$. These $5n_1n_2$ measurements guarantee recovery if s_1 is prime with n_1 and s_2 is prime with n_2 for all sizes n_1 and n_2 ; see section 3.3 for details.

3.1. Proof of Theorem 3.1. Let $\hat{x} = (\hat{x}[0], \dots, \hat{x}[n-1])$ be the DFT of x . Then knowledge of $\mathbb{A}(x)$ is equivalent to knowledge of

$$|\hat{x}[k]|^2, \quad |\hat{x}[k] + \hat{x}[k-s]|^2, \quad \text{and} \quad |\hat{x}[k] - i\hat{x}[k-s]|^2$$

for all $k \in \{0, 1, \dots, n-1\}$ (above, $k-s$ is understood mod n). Write $\hat{x}[k] = |\hat{x}[k]|e^{i\phi[k]}$ so that $\phi[k]$ is the missing phase, and observe that

$$\begin{aligned} |\hat{x}[k] + \hat{x}[k-s]|^2 &= |\hat{x}[k]|^2 + |\hat{x}[k-s]|^2 + 2|\hat{x}[k]||\hat{x}[k-s]|\operatorname{Re}(e^{i(\phi[k-s]-\phi[k])}), \\ |\hat{x}[k] - i\hat{x}[k-s]|^2 &= |\hat{x}[k]|^2 + |\hat{x}[k-s]|^2 + 2|\hat{x}[k]||\hat{x}[k-s]|\operatorname{Im}(e^{i(\phi[k-s]-\phi[k])}). \end{aligned}$$

Hence, if $\hat{x}[k] \neq 0$ for all $k \in \{0, 1, \dots, n-1\}$, our data gives us knowledge of all phase shifts of the form

$$\phi[k-s] - \phi[k], \quad k = 0, 1, \dots, n-1.$$

We can, therefore, initialize $\phi[0]$ to be zero and then get the values of $\phi[-s]$, $\phi[-2s]$, and so on.

This process can be represented as a cycle in the group $\mathbb{Z}/n\mathbb{Z}$ as the sequence $(0, -s, -2s, \dots)$. We would like this cycle to contain n unique elements, which is true if and only if the cyclic subgroup $(0, s, 2s, \dots)$ has order n . This is equivalent to requiring $\gcd(s, n) = 1$. If this subgroup has a smaller order, then recovery is impossible since we finish the cycle before we have all the phases; the phases that we are able to recover do not enable us to determine any more phases without making further assumptions.

For the second part of the theorem, assume without loss of generality that $s = -1$ and that $(k, k') = (0, k_0)$ ($1 < k_0 < n-1$). For simplicity suppose these are the only zeros of the DFT. This creates two disjoint sets of frequency indices: those for which $0 < k < k_0$ and those for which $k_0 < k \leq n-1$. We are given no information about the phase difference between elements of these two subgroups, and hence recovery is not possible. This argument extends to situations where the DFT vanishes more often, in which case we have even more indeterminacy.

3.2. Proof of Theorem 3.3. Let $\hat{x} = \{\hat{x}[k_1, k_2]\}$, where $(k_1, k_2) \in \{0, 1, \dots, n_1-1\} \times \{0, 1, \dots, n_2-1\}$ is the DFT of x . Then we have knowledge of

$$|\hat{x}[k_1, k_2]|^2, \quad |\hat{x}[k_1, k_2] + \hat{x}[k_1-s_s, k_2-s_2]|^2, \quad \text{and} \quad |\hat{x}[k_1, k_2] - i\hat{x}[k_1-s_1, k_2-s_2]|^2$$

for all $(k_1 k_2)$. With the same notation as before, this gives us knowledge of all phase shifts of the form

$$\phi[k_1 - s_1, k_2 - s_2] - \phi[k_1, k_2], \quad 0 \leq k_1 \leq n_1, \quad 0 \leq k_2 \leq n_2 - 1.$$

Hence, we can initialize $\phi[0, 0]$ to be zero and then get the values of $\phi[-s_1, -s_2]$, $\phi[-2s_1, -2s_2]$ and so on. The argument is as before: we would like the cyclic subgroup $((0, 0), (s_1, s_2), (2s_1, 2s_2), \dots)$ in $\mathbb{Z}/n_1\mathbb{Z} \times \mathbb{Z}/n_2\mathbb{Z}$ to have order $n_1 n_2$. Now the order of an element $(s_1, s_2) \in \mathbb{Z}/n_1\mathbb{Z} \times \mathbb{Z}/n_2\mathbb{Z}$ is equal to

$$\text{lcm}(|s_1|, |s_2|) = \text{lcm}(n_1/\text{gcd}(n_1, s_1), n_2/\text{gcd}(n_2, s_2)),$$

where $|s_1|$ is the order of s_1 in $\mathbb{Z}/n_1\mathbb{Z}$ and likewise for $|s_2|$. Noting that $\text{lcm}(a, b) \leq ab$ and that equality is achieved if and only if $\text{gcd}(a, b) = 1$, we must simultaneously have

$$\text{gcd}(s_1, n_1) = 1, \quad \text{gcd}(s_2, n_2) = 1, \quad \text{and} \quad \text{gcd}(n_1, n_2) = 1$$

to have uniqueness.

3.3. Extensions. It is clear from our analysis that if we were to collect $|\mathcal{F}_{n_1 \times n_2} x|^2$ together with

$$\{|\mathcal{F}_{n_1 \times n_2}(x + \mathcal{D}^{s_k} x)|^2, |\mathcal{F}_{n_1 \times n_2}(x - i\mathcal{D}^{s_k} x)|^2\}, \quad k = 1, \dots, K,$$

so that one collects $(2K + 1)n_1 n_2$ measurements, then 2D recovery would be possible if and only if $\{s_1, \dots, s_K\}$ generates $\mathbb{Z}/n_1\mathbb{Z} \times \mathbb{Z}/n_2\mathbb{Z}$ (and the Fourier transform has no nonzero components). This can be understood by analyzing the generators of the group $\mathbb{Z}/n_1\mathbb{Z} \times \mathbb{Z}/n_2\mathbb{Z}$.

A simple instance consists in choosing one modulation pattern to be $(s_1, 0)$ and another to be $(0, s_2)$. If s_1 is prime with n_1 and s_2 is prime with n_2 , these two modulations generate the whole group regardless of the relationship between n_1 and n_2 . An algorithmic way to see this is as follows. Initialize $\phi(0, 0)$. Then by using horizontal modulations, one recovers all phases of the form $\phi(k_1, 0)$. Further, by using vertical modulations (starting with $\phi(k_1, 0)$), one can recover all phases of the form $\phi(k_1, k_2)$ by moving upward.

4. Numerical Experiments. This section introduces numerical simulations to illustrate and study the effectiveness of PhaseLift.

4.1. Numerical Solvers. All numerical algorithms were implemented in MATLAB using TFOCS [6] as well as modifications of TFOCS template files. TFOCS is a library of MATLAB files designed to facilitate the construction of first-order methods for a variety of convex optimization problems, including those we consider here.

In a nutshell, suppose we wish to solve the problem

$$(4.1) \quad \begin{aligned} & \text{minimize} && g(X) := -\ell(b; \mathcal{A}(X)) + \lambda \text{Tr}(X) \\ & \text{subject to} && X \succeq 0, \end{aligned}$$

in which $\ell(b; \mathcal{A}(X))$ is a smooth and concave (in X) log-likelihood. Then a projected gradient method would start with an initial guess X_0 and inductively define

$$(4.2) \quad X_k = \mathcal{P}(X_{k-1} - t_k \nabla g(X_{k-1})),$$

where $\{t_k\}$ is a sequence of stepsize rules and \mathcal{P} is the projection onto the positive semidefinite cone. (Various stepsize rules are typically considered, including fixed stepsizes, backtracking line search, and exact line search.)

TFOCS implements a variety of accelerated first-order methods pioneered by Nesterov; see [61] and the references therein. One variant [4] works as follows. Choose X_0 , set $Y_0 = X_0$ and $\theta_0 = 1$, and inductively define

$$(4.3) \quad \begin{aligned} X_k &= \mathcal{P}(Y_{k-1} - t_k \nabla g(Y_{k-1})), \\ \theta_k &= 2 \left[1 + \sqrt{1 + 4/\theta_{k-1}^2} \right]^{-1}, \\ \beta_k &= \theta_k (\theta_{k-1}^{-1} - 1), \\ Y_k &= X_k + \beta_k (X_k - X_{k-1}), \end{aligned}$$

where $\{t_k\}$ is a sequence of stepsize rules as before. The sequence $\{\theta_k\}$ is usually referred to as a sequence of accelerated parameters, and $\{Y_k\}$ is an auxiliary sequence at which the gradient is to be evaluated. The advantage of this approach is that the computational work per iteration is as in the projected gradient method, but the number of iterations needed to reach a certain accuracy is usually much lower [61]. TFOCS implements such iterations and others like it but with various improvements.

For large problems, e.g., images with a large number N of pixels, it is costly to hold the $N \times N$ optimization variable X in memory. To overcome this issue, our computational approach maintains a low-rank factorization of X . This is achieved by substituting the projection onto the semidefinite cone (the expensive step) with a proxy. Whereas \mathcal{P} dumps the negative eigenvalues as in

$$\mathcal{P}(X) = \sum_i \max(\lambda_i, 0) u_i u_i^*,$$

where $\sum_i \lambda_i u_i u_i^*$ ($\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N$) is any eigenvalue decomposition of X , our proxy keeps only the k largest eigenvalues in the expansion as in

$$(4.4) \quad \mathcal{P}_k(X) = \sum_{i \leq k} \max(\lambda_i, 0) u_i u_i^*.$$

For small values of k —we use k between 10 and 20—this can be efficiently computed since we need only compute the top eigenvectors of a low-rank matrix at each step. Although this approximation gives good empirical results, convergence is no longer guaranteed. For a method like (4.2) or (4.3), the main computational cost of a single iteration is dominated by computing (4.4), whose complexity is in turn governed by the costs of applying \mathcal{A} and \mathcal{A}^* . By maintaining a low-rank factorization of X or Y , these costs are on the order of $k \times M \times n \log n$ for $x \in \mathbb{C}^n$, where M is the number of illuminations. Roughly, each iteration costs on the order of $k \times M$ FFTs.

4.2. Error Measures. To measure performance, we will use the mean-square error (MSE). However, since a solution x_0 is unique only up to global phase, it makes no sense to compute the squared distance between x_0 and the recovery \hat{x}_0 . Rather, we compute the distance to the solution space; i.e., we are interested in the relative MSE defined as

$$\min_{c: |c|=1} \frac{\|cx_0 - \hat{x}_0\|_2^2}{\|x_0\|_2^2}.$$

This is the definition we will adopt throughout the paper;³ the signal-to-noise ratio (SNR) of the measured data is defined as $\text{SNR} = 10 \log_{10} \|b - \tilde{b}\|_2^2 / \|b\|_2^2$, where \tilde{b} denotes the noisy data.

³Alternatively, we could use $\|x_0 x_0^* - \hat{x}_0 \hat{x}_0^*\|_F / \|x_0 x_0^*\|_F$, which gives very similar values.

Although our algorithm favors low-rank solutions, it is not guaranteed to find a rank-one solution. Therefore, if our optimal solution \hat{X}_0 does not have exactly rank one, we extract the rank-one approximation $\hat{x}_0\hat{x}_0^*$, where \hat{x}_0 is an eigenvector associated with the largest eigenvalue. We use a scaling such that $\|\hat{x}_0\|_2^2 = \|x_0\|_2^2$. Note that the ℓ_2 norm of the true solution is generally known since by Parseval's theorem, the ℓ_2 norm of Fx_0 is equal to $\|x_0\|_2$. Hence, observing the diffraction pattern of the object x_0 reveals its squared ℓ_2 norm.

4.3. Alternating Projections. For comparison, we will also apply an alternating projection algorithm in some of the experiments. To describe this algorithm, put $Ax := \{(a_j, x)\}_{j=1}^m$, in which the a_j 's are as in (2.2) so that $\mathbb{A}(x) = |Ax|^2$. In the setting of multiple illuminations, the alternating projection algorithm consists of the following steps: (1) choose an initial guess x_0 ; (2) compute $b_0 = Ax_0$ and for $k = 0, 1, \dots$,

- (i) adjust the modulus of b_k so that it fits the measurements b ,

$$\tilde{b}_k[i] = b[i] \frac{b_k[i]}{|b_k[i]|}, \quad i = 1, \dots, m;$$

- (ii) reproject \tilde{b}_k onto the range of A ,

$$\begin{aligned} x_{k+1} &= \operatorname{argmin} \|Ax - \tilde{b}_k\|_2, \\ b_{k+1} &= Ax_{k+1}. \end{aligned}$$

Observe that we can incorporate appropriate additional information about x (such as positivity, for example) via a suitable modification of the projection step (ii).


4.4. ID Simulations. Phase retrieval for 1D signals arises in fiber optics [17, 39, 38], terahertz communications [42], and speech recognition [63], as well as in the determination of concentration profiles and the detection of planar disorder in diffraction imaging [12, 71]. We evaluate PhaseLift for noiseless and noisy data using a variety of different ‘‘illuminations’’ and test signals.

4.4.1. Noise-Free Measurements. In the first set of experiments we demonstrate the recovery of two very different signals from noiseless data. Both test signals are of length $n = 128$. The first signal, shown in Figure 3(a), is a linear combination of a few sinusoids and represents a typical transfer function one might encounter in optics. The second signal is a complex signal, with independent Gaussian complex entries (each entry is of the form $a + ib$, where a and b are independent $\mathcal{N}(0, 1)$ variables) so that the real and imaginary parts are independent white noise sequences; the real part of the signal is shown in Figure 3(b).

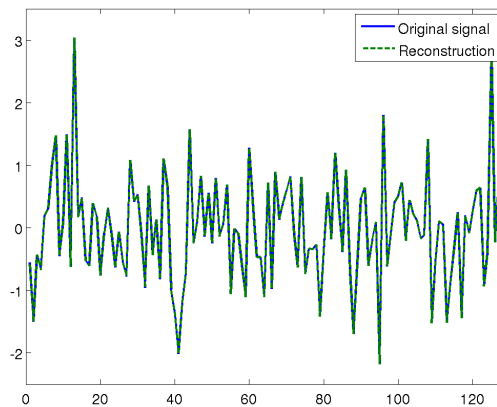
Four random binary masks are used to perform the structured illumination so that we measure $|Ax|^2$, in which

$$A = F \begin{bmatrix} W_1 \\ W_2 \\ W_3 \\ W_4 \end{bmatrix},$$

where each W_i is diagonal with either 0 or 1 on the diagonal, resulting in a total of 512 intensity measurements. We work with the objective functional $\frac{1}{2}\|b - \mathcal{A}(X)\|_2^2 + \lambda \operatorname{Tr}(X)$ and the constraint $X \succeq 0$ to recover the signal, in which we use a small



(a) Smooth signal and its reconstruction



(b) Random signal and its reconstruction (real part only)

Fig. 3 Two test signals and their reconstructions. The recovered signals are essentially indistinguishable from the originals.

value for λ such as 0.05 since we are dealing with noise-free data. We apply the reweighting scheme as discussed in section 2.4. (To achieve perfect reconstruction, one would have to let $\lambda \rightarrow 0$ as the iteration count increases.) The algorithm is terminated when the relative *residual error* is less than a fixed tolerance, namely, $\|\mathcal{A}(\hat{x}_0 \hat{x}_0^*) - b\|_2 \leq 10^{-6} \|b\|_2$, where \hat{x}_0 is the reconstructed signal just as before. The original and recovered signals are plotted in Figures 3(a) and 3(b). The MSE on a dB-scale (i.e., $10 \log_{10}(\text{MSE})$) is 87.3dB in the first case and 90.5dB in the second. The computation time for the recovery of such a 1D signal was in the order of a few minutes.

We have repeated these experiments with the same test signals and the same algorithm, but using Gaussian masks instead of binary masks. In other words, the W_i 's have Gaussian entries on the diagonal. It turns out that in this case, three illuminations—instead of four—were sufficient to obtain similar performance. This

seems to be empirical support for a long-standing conjecture in quantum mechanics due to Wright (see, e.g., the concluding section of [72]). The conjecture is that there exist three unitary operators U_1, U_2, U_3 such that the phase of the (finite-dimensional) signal x is uniquely determined by the measurements $|U_1x|, |U_2x|, |U_3x|$.⁴ Our simulations suggest that one can choose $U_1 = F$, $U_2 = FW_1$, and $U_3 = FW_2$, where W_1, W_2 are diagonal matrices with i.i.d. complex normal random variables as diagonal entries. A deterministic choice, which was equally successful in our experiments and is closer to the quantum mechanical setting, is the following: $U_1 = F$, $U_2 = I$, $U_3 = FW$, where the $n \times n$ diagonal matrix W has entries $W_{j,j} = \frac{1}{\sqrt{n}} e^{-2\pi i j(j+n)/(2n)}$, $j = 0, \dots, n-1$. Furthermore, we point out that no reweighting was needed when we used six or more Gaussian masks. Expressed differently, plain trace-norm minimization succeeds with $6n$ or more intensity measurements of this kind.

We also applied the alternating projection algorithm of section 4.3, with random initial guess, to the examples above. We terminate the iterations of the alternating projection algorithm if the relative error between two successive iterates is less than 10^{-6} . When using three Gaussian masks (or the three operators I, F, FW related to the quantum mechanical setting), alternating projections always failed to recover the correct solution. The algorithm never even found an approximation with a relative MSE less than 1. With four illuminations the alternating projections algorithm computed the correct solution in about 40% of the experiments and returned a relative MSE greater than 1 in the other 60%. As we increased the number of masks, the behavior of alternating projections improved; it succeeded in about 99% of the experiments with eight Gaussian illuminations.

4.4.2. Noisy Measurements. In the next set of experiments, we consider the case when the measurements are contaminated with Poisson noise. The test signal is again a complex random signal as above. Four, six, and eight illuminations with random binary masks are used. We add random Poisson noise to the measurements for five different SNR levels, ranging from about 16dB to about 52dB. Since the solution is known, we have calculated reconstructions for various values of the parameter λ balancing the negative log-likelihood and the trace norm, and we report results for that λ giving the lowest MSE. We implemented this strategy via the standard golden section search [41]. In practice one would have to find the best λ via a strategy like cross validation (CV) or generalized cross validation (GCV). For each SNR level we repeated the experiment 10 times with different random noise and different binary masks.

Figure 4 shows the average relative MSE in dB (the values of $10 \log_{10}(\text{rel. MSE})$ are plotted) versus the SNR. The error curves show clearly the desirable linear behavior between SNR and MSE with respect to the log-log scale. The performance degrades very gracefully with decreasing SNR. Furthermore, the difference of about 5dB between the error curve associated with four measurements and the error curve associated with eight measurements corresponds to an almost twofold error reduction, which is about as much improvement as one can hope to gain by doubling the number of measurements.

We repeat this experiment with deterministic masks as described in section 3 (see (3.1)) instead of random masks. To achieve robustness vis-à-vis noise, three masks (as in Theorem 3.1) do not seem to suffice. We thus collect $7n$ measurements of

⁴Shortly after submission of this manuscript, the paper “Quantum Tomography under Prior Information” was posted on arXiv.org; cf. [36]. Based on the results in that paper, it is now assumed that Wright’s conjecture is wrong.

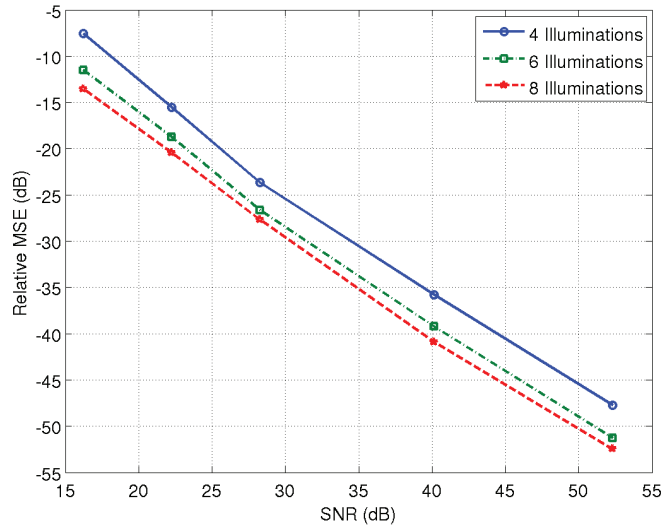


Fig. 4 Relative MSE versus SNR on a dB-scale for different numbers of illuminations with binary masks. The linear relationship between SNR and MSE is apparent.




Fig. 5 Relative MSE versus SNR on a dB-scale: seven illuminations with deterministic masks and with random masks.

the form $|F_n x|^2$, and then $\{|F_n(x + D^s x)|^2, |F_n(x - iD^s x)|^2\}$ with $s = 3, 5, 7$ as in (3.1). The recovery is very stable, and the performance curve is shown in Figure 5. For comparison we also show the performance curve corresponding to seven Gaussian random masks. Gaussian random masks yield better MSE in this example.

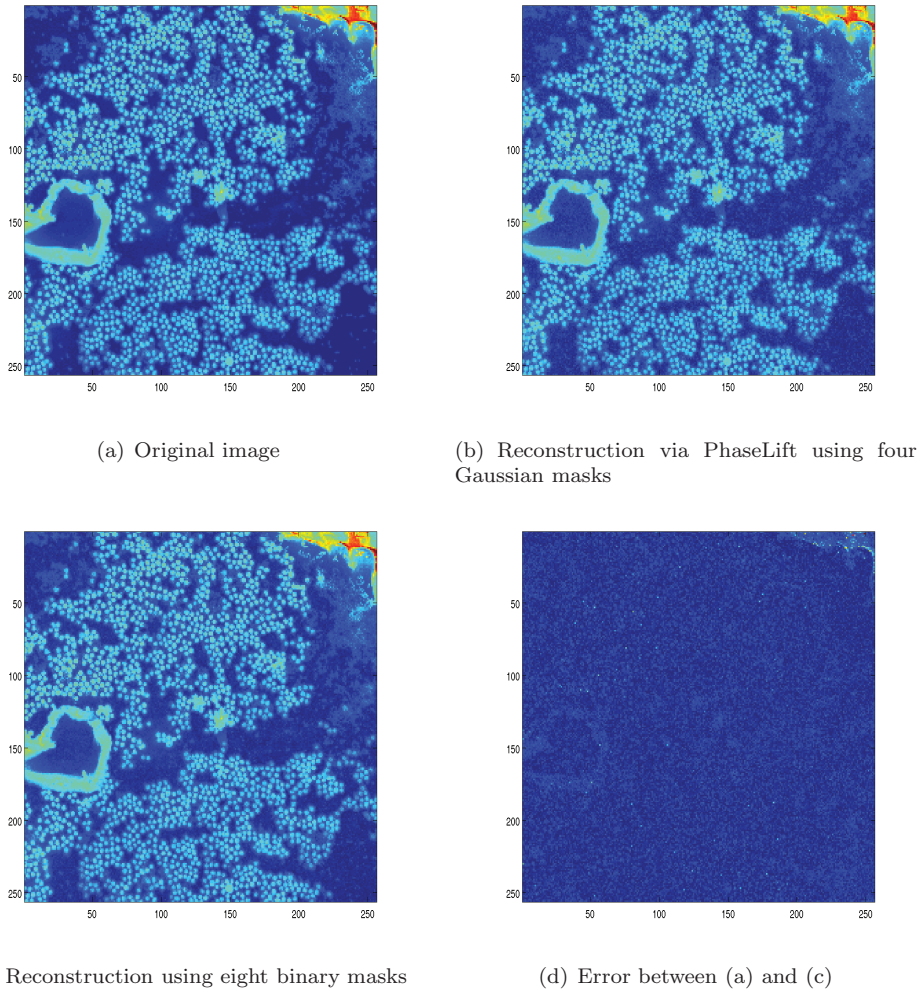


Fig. 6 Original gold balls image and reconstructions via PhaseLift.

4.5. 2D Simulations. We consider a stylized version of a setup one encounters in X-ray crystallography or diffraction imaging. The test image, shown in Figure 6(a) (magnitude), is a complex-valued image of size 256×256 , whose pixel values correspond to the complex transmission coefficients of a collection of gold balls embedded in a medium.

4.5.1. Noise-Free Measurements. In the first experiment, we demonstrate the recovery of the image shown in Figure 6(a) from noiseless measurements. We consider two different types of illuminations. The first type uses Gaussian random masks in which the coefficients on the diagonal of W_k are independent real-valued standard normal variables. We use four illuminations, one being constant, i.e., $W_1 = I$, and the other three Gaussian. Again, we choose a small value of λ set to 0.05 in $\frac{1}{2}\|b - \mathcal{A}(X)\|_2^2 + \lambda \text{Tr}(X)$ since we have no noise, and stop the reweighting iterations as soon as the residual error obeys $\|\mathcal{A}(\hat{x}_0 \hat{x}_0^*) - b\|_2 \leq 10^{-4} \|b\|_2$. The reconstruction, shown in

Figure 6(b), is visually indistinguishable from the original. Since the original image and the reconstruction are complex-valued, we display only the absolute value of each image throughout this and the next subsection.

Gaussian random masks may not be realizable in practice. Our second example uses simple random binary masks, where the entries are either 0 or 1 with equal probability. In this case, a larger number of illuminations as well as a larger number of reweighting steps are required to achieve a reconstruction of comparable quality. The result for eight binary illuminations is shown in Figure 6(c). The computation time for the recovery of such a 2D signal was on the order of several hours. Depending on the application, the current implementation of PhaseLift may be too slow for large-scale images. The development of a scalable algorithm is an important task, but is beyond the scope of this paper and will be addressed in our future research.

We repeated the experiment using as test signal an image with independent standard normal complex entries; that is, an entry is of the form $z_1 + iz_2$, where z_1 and z_2 are independent $\mathcal{N}(0, 1)$ variables. Here, four Gaussian masks were not sufficient, but we did achieve successful recovery with five Gaussian masks. We also applied the alternating projection algorithm to both test images. In the gold balls example, alternating projections succeeded both with four Gaussian masks and with eight binary masks. For the random image, however, alternating projections failed to recover the correct solution when we used five Gaussian masks. The best approximation it found has a relative MSE greater than 1. As in the 1D example, the performance of alternating projections improved as we increased the number of masks, eventually yielding consistent recovery of the correct image when we employed eight or more Gaussian masks.

4.5.2. Noisy Measurements. In the second set of experiments we consider the same test image as before but now with noisy measurements. In the first experiment the SNR is 20dB; in the second experiment the SNR is 60dB. We use 32 Gaussian random masks in each case. The resulting reconstructions are depicted in Figure 7(a) (20dB case) and Figure 7(b) (60dB case). The MSE in the 20dB case is 11.83dB.

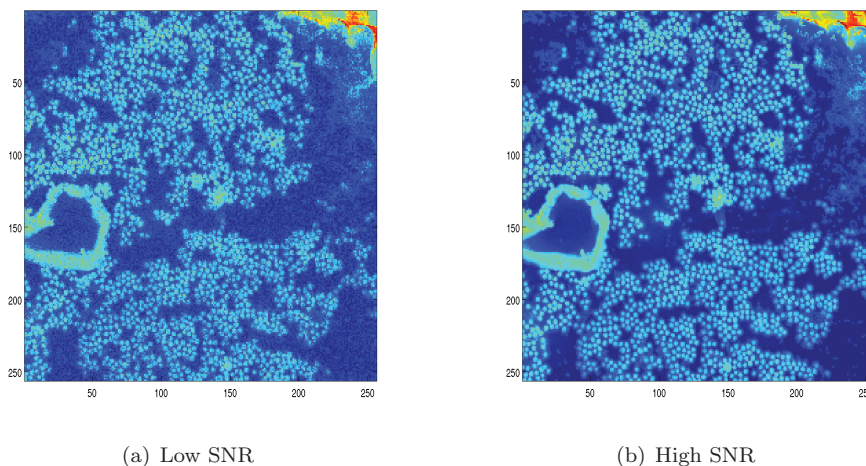


Fig. 7 Reconstructions from noisy data via PhaseLift using 32 Gaussian random masks.

While the reconstructed image appears slightly “fuzzier” than the original image, all features of the image are clearly visible. In the 60dB case the MSE is 47.96dB, and the reconstruction is virtually indistinguishable from the original image.

4.5.3. Multiple Measurements via Oversampling. Oversampling of 2D signals is widely used to overcome the nonuniqueness of the phase retrieval problem. We now explore whether this approach is viable.

Here, we consider signals with real, nonnegative values as test images, a case frequently considered in the literature; see, e.g., [58, 57, 56]. These images are of size 128×128 . We take noiseless measurements and apply PhaseLift as the alternating projection algorithm (also known as Fienup’s error reduction algorithm) [3, section 4.A]. For each method, we terminate the iterations if the relative residual error is less than 10^{-3} or if the relative error between two successive iterates is less than 10^{-6} . Since we assume that the support of the signal is known, there is no ambiguity of the solution with respect to translations. Moreover, the support is chosen to be nonsymmetric around the origin, and thus there is also no ambiguity with respect to reflections around the origin. Finally, since the signal is real-valued and positive, there is no ambiguity with respect to global phase in this case.

- The simulations show that PhaseLift yields reconstructions that fit the measured data well, yielding a small relative residual error $\|\mathcal{A}(X) - y\|_2/\|y\|_2$, yet the reconstructions are far away from the true signal. This behavior is indicative of an ill-conditioned problem.
- The iterates of the alternating projection algorithm stagnated most of the time without converging to a solution. At other times the algorithms did yield reconstructions that fit the measured data well, but in either case the reconstruction was always very different from the true signal. Moreover, the reconstructions varied widely depending on the initial (random) guess.

Table 1 displays the results of PhaseLift as well as the alternating projection algorithm as described in [3, section 4. A] (the other versions discussed in section 4 of [3] yield comparable results). The setup is this: We oversample the signal in each dimension by a factor of r , where $r = 2, 3, 4, 5$. For each oversampling rate, we run 10 experiments using a different test signal each time. The table shows the average residual errors over 10 runs as well as the average relative MSE. The ill-posedness of the problem is evident from the disconnect between small residual error and large reconstruction

Table 1 *MSE obtained by alternating projections and by PhaseLift with reweighting from oversampled DFT measurements taken on 2D real-valued and positive test images. The alternating projection algorithm does not always find a signal consistent with the data as well as the support constraint. (After the projection step in the spatial domain, the current guess does not always match the measurement in Fourier space. After “projection” in Fourier space, the signal is not the Fourier transform of a signal obeying the spatial constraints.) Our approach always finds signals matching measured data very well, and yet the reconstructions incur a large reconstruction error. This indicates severe ill-posedness since we have several distinct solutions providing an excellent fit to the measured data.*

Algorithm Oversampling	2	3	4	5
$\ \mathcal{A}(X) - y\ _2/\ y\ _2$ (Alt.Proj.)	0.0650	0.0607	0.0541	0.0713
Relative MSE (Alt.Proj.)	0.6931	0.6882	0.6736	0.6878
$\ \mathcal{A}(X) - y\ _2/\ y\ _2$ (PhaseLift)	0.0051	0.0055	0.0056	0.0052
Relative MSE (PhaseLift)	0.4932	0.4893	0.4960	0.4981

error; that is to say, we fit the data very well and yet observe a large reconstruction error. *Thus, in stark contrast to what is widely believed, our simulations indicate that oversampling by itself is not a viable strategy for phase retrieval even for nonnegative, real-valued images.*

5. Discussion. This paper introduces a novel framework for phase retrieval, combining multiple illuminations with tools from convex optimization, which has been shown to work very well in practice and bears great potential. Having said this, we note that our work also calls for improved theory, improved algorithms, and a physical implementation of these ideas. Regarding this last point, it would be interesting to design physical experiments to test our methodology on real data, and we hope to report on this in a future publication. For now, we would like to bring up important open problems.

At the theoretical level, we need to understand the following: For which families of physically implementable structured illuminations does the trace-norm heuristic succeed? How many diffraction patterns are provably sufficient for our convex programming approach to work? Also, we have shown that our approach is robust to noise in the sense that the performance degrades very gracefully as the SNR decreases. Can this be made rigorous? In a recent follow-up paper [14], it has been proved that for measurement vectors sampled independently and uniformly at random on the unit sphere, PhaseLift is indeed robust to noise, provided that the number of measurements is on the order of $n \log n$. It is natural to ask whether this line of work extends to the structured illuminations we consider in this paper. Here, it is very likely that the tools and ideas developed in the theories of compressed sensing and matrix completion will play a key role in addressing these fundamental issues.

At the algorithmic level, we need to address the fact that the lifting creates optimization problems of potentially enormous size. A tantalizing prospect is whether or not it is possible to use knowledge that the solution has low rank, e.g., rank one, to design algorithms which do not need to assemble or store very large matrices. If so, how can this be done? Here, randomized algorithms holding up a sketch of the full matrix may prove very helpful.

Finally, we would like to close by returning to another finding of this paper. Namely, oversampling the Fourier transform—this is the same as assuming finite support of the specimen—appears extremely problematic in practice, even for real-valued nonnegative signals. To be sure, we have demonstrated that there typically exist very distinct 2D signals whose moduli of the Fourier transform nearly coincide, whatever the degree of oversampling. In light of this extreme ill-posedness, we have trouble understanding why this technique is used so heavily when it does not produce useful results in the absence of very specific a priori information about the image. Moreover, our concern is compounded by the additional fact that algorithms in common use tend to return solutions that depend on an initial guess so that different runs return widely different solutions. Of course, one might also be willing to rely on image priors far stronger than finite spatial extent, real-valuedness, and positivity; they would, however, need to be specified.

Acknowledgments. We are indebted to Stefano Marchesini for inspiring and helpful discussions on the phase problem in X-ray crystallography as well as for providing us with the gold balls data set used in section 4. We want to thank Philippe Jaming for bringing to our attention Wright’s conjecture in [72].

REFERENCES

- [1] R. BALAN, B. BODMANN, P. G. CASAZZA, AND D. EDIDIN, *Painless reconstruction from magnitudes of frame coefficients*, J. Fourier Anal. Appl., 15 (2009), pp. 488–501.
- [2] R. BALAN, P. G. CASAZZA, AND D. EDIDIN, *On signal reconstruction without noisy phase*, Appl. Comput. Harmon. Anal., 20 (2006), pp. 345–356.
- [3] H. H. BAUSCHKE, P. L. COMBETTES, AND D. R. LUKE, *Phase retrieval, error reduction algorithm, and Fienup variants: A view from convex optimization*, J. Opt. Soc. Amer. A, 19 (2002), pp. 1334–1345.
- [4] A. BECK AND M. TEBoulLE, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM J. Imaging Sci., 2 (2009), pp. 183–202.
- [5] C. BECK AND R. D’ANDREA, *Computational study and comparisons of LFT reducibility methods*, in Proceedings of the American Control Conference, 1998, pp. 1013–1017.
- [6] S. BECKER, E. J. CANDÈS, AND M. GRANT, *Templates for convex cone problems with applications to sparse signal recovery*, Math. Program. Comput., 3 (2011), pp. 165–218.
- [7] A. BEN-TAL AND A. NEMIROVSKI, *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*, MOS-SIAM Ser. Optim. 2, SIAM, Philadelphia, 2001.
- [8] G. BIANCHI, F. SEGALA, AND A. VOLCIC, *The solution of the covariogram problem for plane C_+^2 convex bodies*, J. Differential Geom., 60 (2002), pp. 177–198.
- [9] M. J. BOGAN, W. H. BENNER, S. BOUTET, U. ROHNER, M. FRANK, A. BARTY, M. M. SEIBERT, F. MAIA, S. MARCHESINI, S. BAJT, B. WOODS, V. RIOT, S. P. HAU-RIEGE, M. SVENDA, E. MARKLUND, E. SPILLER, J. HADJU, AND H. N. CHAPMAN, *Single particle X-ray diffractive imaging*, Nano Lett., 8 (2008), pp. 310–316.
- [10] L. M. BRÈGMAN, *The method of successive projection for finding a common point of convex sets*, Soviet Math. Dokl., 6 (1965), pp. 688–692.
- [11] Y. M. BRUCK AND L. G. SODIN, *On the ambiguity of the image reconstruction problem*, Opt. Comm., 30 (1979), pp. 304–308.
- [12] O. BUNK, A. DIAZ, F. PFEIFFER, C. DAVID, B. SCHMITT, D. K. SATAPATHY, AND J. F. VAN DER VEEN, *Diffractive imaging for periodic samples: Retrieving one-dimensional concentration profiles across microfluidic channels*, Acta Crystallogr. A, 63 (2007), pp. 306–314.
- [13] E. J. CANDÈS AND B. RECHT, *Exact matrix completion via convex optimization*, Found. Comput. Math., 9 (2009), pp. 717–772.
- [14] E. CANDÈS, T. STROHMER, AND V. VORONINSKI, *PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming*, Comm. Pure Appl. Math., 66 (2013), pp. 1241–1274.
- [15] E. J. CANDÈS AND T. TAO, *The power of convex relaxation: Near-optimal matrix completion*, IEEE Trans. Inform. Theory, 56 (2010), pp. 2053–2080.
- [16] E. J. CANDÈS, M. B. WAKIN, AND S. P. BOYD, *Enhancing sparsity by reweighted l_1 minimization*, J. Fourier Anal. Appl., 14 (2008), pp. 877–905.
- [17] A. CARBALLAR AND M. A. MURIEL, *Phase reconstruction from reflectivity in fiber Bragg gratings*, J. Lightwave Technol., 15 (1997), pp. 1314–1322.
- [18] A. CHAI, M. MOSCOSO, AND G. PAPANICOLAOU, *Array Imaging Using Intensity-only Measurements*, Technical report, Stanford University, Stanford, CA, 2010.
- [19] C.-C. CHEN, J. MIAO, C. W. WANG, AND T. K. LEE, *Application of the optimization technique to noncrystalline X-ray diffraction microscopy: Guided hybrid input-output method (GHIO)*, Phys. Rev. B, 76 (2007), 064113.
- [20] J. V. CORBETT, *The Pauli problem, state reconstruction and quantum-real numbers*, Rep. Math. Phys., 57 (2006), pp. 53–68.
- [21] J. C. DAINTY AND J. R. FIENUP, *Phase Retrieval and Image Reconstruction for Astronomy*, in Image Recovery: Theory and Application, H. Stark, ed., Academic Press, New York, 1987, pp. 231–275.
- [22] M. DIEROLF, A. MENZEL, P. THIBAUT, P. SCHNEIDER, C. M. KEWISH, R. WEPF, O. BUNK, AND F. PFEIFFER, *Ptychographic X-ray computed tomography at the nanoscale*, Nature, 467 (2010), pp. 436–439.
- [23] H. DUADI, O. MARGALIT, V. MICO, J. A. RODRIGO, T. ALIEVA, J. GARCIA, AND Z. ZALEVSKY, *Digital holography and phase retrieval*, in Holography, Research and Technologies, J. Rosen, ed., InTech, 2011; available online at <http://www.intechopen.com/books/holography-research-and-technologies/digital-holography-and-phase-retrieval>.
- [24] A. FARIDIAN, D. HOPP, G. PEDRINI, U. EIGENTHALER, M. HIRSCHER, AND W. OSTEN, *Nanoscale imaging using deep ultraviolet digital holographic microscopy*, Opt. Express, 18 (2010), pp. 14159–14164.

- [25] M. FAZEL, *Matrix Rank Minimization with Applications*, Ph.D. thesis, Stanford University, Stanford, CA, 2002.
- [26] M. FAZEL, H. HINDI, AND S. BOYD, *Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices*, in Proceedings of the American Control Conference, 2003, pp. 2156–2162.
- [27] J. R. FIENUP, *Reconstruction of an object from the modulus of its Fourier transform*, Opt. Lett., 3 (1978), pp. 27–29.
- [28] J. R. FIENUP, *Phase retrieval algorithms: A comparison*, Appl. Opt., 21 (1982), pp. 2758–2768.
- [29] J. FINKELSTEIN, *Pure-state informationally complete and “really” complete measurements*, Phys. Rev. A (3), 70 (2004), 052107.
- [30] R. W. GERCHBERG AND W. O. SAXTON, *A practical algorithm for the determination of phase from image and diffraction plane pictures*, Optik, 35 (1972), pp. 237–246.
- [31] M. X. GOEMANS AND D. P. WILLIAMSON, *Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming*, J. ACM, 42 (1995), pp. 1115–1145.
- [32] L. GUBIN, B. POLYAK, AND E. RAIK, *The method of projections for finding the common point of convex sets*, USSR Comput. Math. Math. Phys., 7 (1967), pp. 1–24.
- [33] R. W. HARRISON, *Phase problem in crystallography*, J. Opt. Soc. Amer. A, 10 (1993), pp. 1045–1055.
- [34] H. HAUPTMAN, *The direct methods of X-ray crystallography*, Science, 233 (1986), pp. 178–183.
- [35] M. HAYES, *The reconstruction of a multidimensional sequence from the phase or magnitude of its Fourier transform*, IEEE Trans. Acoust. Speech Signal Process., 30 (1982), pp. 140–154.
- [36] T. HEINOSAARI, L. MAZZARELLA, AND M. M. WOLF, *Quantum Tomography under Prior Information*, preprint, arXiv:1109.5478v1[quant-ph], 2011.
- [37] N. HURT, *Phase Retrieval and Zero Crossings*, Kluwer Academic Publishers, Norwell, MA, 1989.
- [38] E. IP, A. P. T. LAU, D. J. F. BARROS, AND J. M. KAHN, *Coherent detection in optical fiber systems*, Opt. Express, 16 (2008), pp. 753–791.
- [39] Y. IVANKOVSKI AND D. MENDLOVIC, *High-rate long-distance fiber-optic communication based on advanced modulation techniques*, Appl. Opt., 38 (1999), pp. 5533–5540.
- [40] I. JOHNSON, K. JEFIMOV, O. BUNK, C. DAVID, M. DIEROLF, J. GRAY, D. RENKER, AND F. PFEIFFER, *Coherent diffractive imaging using phase front modifications*, Phys. Rev. Lett., 100 (2008), 155503.
- [41] J. KIEFER, *Sequential minimax search for a maximum*, Proc. Amer. Math. Soc., 4 (1953), pp. 502–506.
- [42] T. KLEINE-OSTMANN AND T. NAGATSUMA, *A review on terahertz communications research*, J. Infrared Milli. Terah. Waves, 32 (2011), pp. 143–171.
- [43] M. V. KLIBANOV, P. E. SACKS, AND A. V. TIKHONRAVOV, *The phase retrieval problem*, Inverse Problems, 11 (1995), p. 1–28.
- [44] A. LEVI AND H. STARK, *Restoration from phase and magnitude by generalized projections*, in Image Recovery: Theory and Application, H. Stark, ed., Academic Press, San Diego, CA, 1987, pp. 277–320.
- [45] Y. J. LIU, B. CHEN, E. R. LI, J. Y. WANG, A. MARCELLI, S. W. WILKENS, H. MING, Y. C. TIAN, K. A. NUGENT, P. P. ZHU, AND Z. Y. WU, *Phase retrieval in x-ray imaging based on using structured illumination*, Phys. Rev. A, 78 (2008), 023817.
- [46] E. G. LOEWEN AND E. POPOV, *Diffraction Gratings and Applications*, Marcel Dekker, New York, 1997.
- [47] Y. LU AND M. VETTERLI, *Sparse spectral factorization: Unicity and reconstruction algorithms*, in Proceedings of the 36th International Conference on Acoustics, Speech and Signal Processing (ICASSP), Prague, Czech Republic, 2011.
- [48] D. R. LUKE, *Finding best approximation pairs relative to a convex and prox-regular set in a Hilbert space*, SIAM J. Optim., 19 (2008), pp. 714–739.
- [49] D. R. LUKE, J. V. BURKE, AND R. G. LYON, *Optical wavefront reconstruction: Theory and numerical methods*, SIAM Rev., 44 (2002), pp. 169–224.
- [50] S. MARCHESINI, *Phase retrieval and saddle-point optimization*, J. Opt. Soc. Amer. A, 24 (2007), pp. 3289–3296.
- [51] S. MARCHESINI, *A unified evaluation of iterative projection algorithms for phase retrieval*, Rev. Sci. Instrum., 78 (2007), 011301.
- [52] S. MARCHESINI, *Ab Initio Compressive Phase Retrieval*, preprint, arXiv:0809.2006v1[physics.optics], 2008.
- [53] M. MESBAHI AND G. P. PAPAVALOPOULOS, *On the rank minimization problem over a positive semidefinite linear matrix inequality*, IEEE Trans. Automat. Control, 42 (1997), pp. 239–243.

- [54] J. MIAO, H. N. CHAPMAN, AND D. SAYRE, *Image reconstruction from the oversampled diffraction pattern*, Microscopy Microanalysis, 3 (Suppl. 2) (1997), pp. 1155–1156.
- [55] 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*, Annu. Rev. Phys. Chem., 59 (2008), pp. 387–410.
- [56] J. MIAO, J. KIRZ, AND D. SAYRE, *The oversampling phasing method*, Acta Cryst., D56 (2000), pp. 1312–1315.
- [57] J. MIAO, D. SAYRE, AND H. N. CHAPMAN, *Phase retrieval from the magnitude of the Fourier transforms of nonperiodic objects*, J. Opt. Soc. Amer. A, 15 (1998), pp. 1662–1669.
- [58] R. P. MILLANE, *Phase retrieval in crystallography and optics*, J. Opt. Soc. Amer. A., 7 (1990), pp. 394–411.
- [59] R. P. MILLANE, *Recent advances in phase retrieval*, in Image Reconstruction from Incomplete Data IV, P. J. Bones, M. A. Fiddy, and R. P. Millane, eds., Proc. SPIE 6316, SPIE, Bellingham, WA, 2006, 63160E.
- [60] D. L. MISELL, *A method for the solution of the phase problem in electron microscopy*, J. Phys. D: Appl. Phys., 6 (1973), pp. L6–L9.
- [61] Y. NESTEROV, *Introductory Lectures on Convex Optimization: A Basic Course*, Appl. Optim. 87, Kluwer Academic Publishers, Boston, 2004.
- [62] K. A. NUGENT, A. G. PEELE, H. N. CHAPMAN, AND A. P. MANUSCO, *Unique phase recovery for nonperiodic objects*, Phys. Rev. Lett., 91 (2003), 203902.
- [63] L. RABINER AND B. H. JUANG, *Fundamentals of Speech Recognition*, Signal Processing Series, Prentice-Hall, Upper Saddle River, NJ, 1993.
- [64] B. RECHT, M. FAZEL, AND P. A. PARRILO, *Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization*, SIAM Rev., 52 (2010), pp. 471–501.
- [65] H. REICHENBACH, *Philosophic Foundations of Quantum Mechanics*, University of California Press, Berkeley, CA, 1944.
- [66] J. M. RODENBURG, *Ptychography and related diffractive imaging methods*, Adv. Imaging Electron Phys., 150 (2008), pp. 87–184.
- [67] J. L. C. SANZ, *Mathematical considerations for the problem of Fourier transform phase retrieval from magnitude*, SIAM J. Appl. Math., 45 (1985), pp. 651–664.
- [68] G. SCAPIN, *Structural biology and drug discovery*, Curr. Pharm. Des., 12 (2006), pp. 2087–2097.
- [69] Y. SHECHTMAN, Y. C. ELДАР, A. SZAMEIT, AND M. SEGEV, *Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing*, Opt. Express, 19 (2011), pp. 14807–14822.
- [70] P. THIBAUT, M. DIEROLF, O. BUNK, A. MENZEL, AND F. PFEIFFER, *Probe retrieval in ptychographic coherent diffractive imaging*, Ultramicroscopy, 109 (2009), pp. 338–343.
- [71] D. P. VARN, G. S. CANRIGHT, AND J. P. CRUTCHFIELD, *Discovering planar disorder in close-packed structures from x-ray diffraction: Beyond the fault model*, Phys. Rev. B, 66 (2002), 174110.
- [72] A. VOGT, *Position and momentum distributions do not determine the quantum mechanical state*, in Mathematical Foundations of Quantum Theory, A. R. Marlow, ed., Academic Press, New York, 1978, pp. 365–372.
- [73] J. VON NEUMANN, *Functional Operators*, Vol. II, Ann. Math. Stud. 22, Princeton University Press, Princeton, NJ, 1950.
- [74] A. WALTHER, *The question of phase retrieval in optics*, Optical Acta, 10 (1963), pp. 41–49.
- [75] D. C. YOULA, *Mathematical theory of image restoration by the method of convex projections*, in Image Recovery: Theory and Application, H. Stark, ed., Academic Press, San Diego, CA, 1987, pp. 29–77.