

Quadratic Basis Pursuit

Henrik Ohlsson, *Member, IEEE*, Allen Y. Yang, *Member, IEEE*, Roy Dong, Michel Verhaegen, S. Shankar Sastry, *Fellow, IEEE*

Abstract—In many compressive sensing problems today, the relationship between the measurements and the unknowns could be nonlinear. Traditional treatment of such nonlinear relationships have been to approximate the nonlinearity via a linear model and the subsequent un-modeled dynamics as noise. The ability to more accurately characterize nonlinear models has the potential to improve the results in both existing compressive sensing applications and those where a linear approximation does not suffice, e.g., phase retrieval. In this paper, we extend the classical compressive sensing framework to a second-order Taylor expansion of the nonlinearity. Using a lifting technique and a method we call quadratic basis pursuit, we show that the sparse signal can be recovered exactly when the sampling rate is sufficiently high. We further present efficient numerical algorithms to recover sparse signals in second-order nonlinear systems, which are considerably more difficult to solve than their linear counterparts in sparse optimization.

I. INTRODUCTION

Consider the problem of finding the sparsest signal \mathbf{x} satisfying a system of linear equations:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^n} \quad & \|\mathbf{x}\|_0 \\ \text{subj. to} \quad & y_i = \mathbf{b}_i^\top \mathbf{x}, \quad y_i \in \mathbb{R}, \mathbf{b}_i \in \mathbb{R}^n, i = 1, \dots, N. \end{aligned} \quad (1)$$

This problem is known to be combinatorial and NP-hard [2] and a number of approaches to approximate its solution have been proposed. One of the most well known approaches is to relax the zero norm and replace it with the ℓ_1 -norm:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{x}\|_1 \quad \text{subj. to} \quad y_i = \mathbf{b}_i^\top \mathbf{x}, \quad i = 1, \dots, N. \quad (2)$$

This approach is often referred to as *basis pursuit* (BP) [3].

The ability to recover the optimal solution to (1) is essential in the theory of *compressive sensing* (CS) [4], [5] and a tremendous amount of work has been dedicated to solving and analyzing the solution of (1) and (2) in the last decade. Today CS is regarded as a powerful tool in signal processing and widely used in many applications. For a detailed review of the literature, the reader is referred to several recent publications such as [6], [7].

Ohlsson, Yang, Dong, and Sastry are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA. Ohlsson is also with the Division of Automatic Control, Department of Electrical Engineering, Linköping University, Sweden. Verhaegen is with the Delft Center for Systems and Control, Delft University, Delft 2628CD, The Netherlands. Corresponding author: Henrik Ohlsson, Cory Hall, University of California, Berkeley, CA 94720. Email: ohlsson@eecs.berkeley.edu.

The authors gratefully acknowledge support by the Swedish Research Council in the Linnaeus center CADICS, the European Research Council under the advanced grant LEARN, contract 267381, by a postdoctoral grant from the Sweden-America Foundation, donated by ASEA's Fellowship Fund, by a postdoctoral grant from the Swedish Research Council, and by ARO 63092-MA-II and DARPA FA8650-11-1-7153.

This paper was presented in part at NIPS 2012, Lake Tahoe, USA, Dec 3-6, 2012, [1].

It has recently been shown that CS can be extended to nonlinear models. More specifically, the relatively new topic of *nonlinear compressive sensing* (NLCS) deals with a more general problem of finding the sparsest signal \mathbf{x} to a nonlinear set of equations:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^n} \quad & \|\mathbf{x}\|_0 \\ \text{subj. to} \quad & y_i = f_i(\mathbf{x}), \quad y_i \in \mathbb{R}, i = 1, \dots, N, \end{aligned} \quad (3)$$

where each $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is a continuously differentiable function. Compared to CS, the literature on NLCS is still very limited. The interested reader is referred to [8], [9] and references therein.

In this paper, we will restrict our attention from rather general nonlinear systems, and instead focus on nonlinearities that depends quadratically on the unknown \mathbf{x} . More specifically, we consider the following problem formulated in the complex domain:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{C}^n} \quad & \|\mathbf{x}\|_0 \\ \text{subj. to} \quad & y_i = a_i + \mathbf{b}_i^H \mathbf{x} + \mathbf{x}^H \mathbf{c}_i + \mathbf{x}^H \mathbf{Q}_i \mathbf{x}, \\ & i = 1, \dots, N, \end{aligned} \quad (4)$$

where $a_i \in \mathbb{C}$, $\mathbf{b}_i, \mathbf{c}_i \in \mathbb{C}^n$, $y_i \in \mathbb{C}$, and $\mathbf{Q}_i \in \mathbb{C}^{n \times n}$, $i = 1, \dots, N$. In a sense, being able to solve (4) would make it possible to apply the principles of CS to a second-order Taylor expansion of the nonlinear relationship in (3), while traditional CS mainly considers its linear approximation or first-order Taylor expansion. In particular, in the most simple case, when a second order Taylor expansion is taken around zero (i.e., a Maclaurin expansion), let $a_i = f_i(0)$, $\mathbf{b}_i = \mathbf{c}_i = \nabla_{\mathbf{x}}^T f_i(0)/2$ and $\mathbf{Q}_i = \nabla_{\mathbf{x}}^2 f_i(0)/2$, $i = 1, \dots, N$, with $\nabla_{\mathbf{x}}$ and $\nabla_{\mathbf{x}}^2$ denoting the gradient and Hessian with respect to \mathbf{x} . In this case, \mathbf{Q} is a Hermitian matrix. Nevertheless, we note that our derivations in the paper does not depend on the matrix \mathbf{Q} to be symmetric in the real domain or Hermitian in the complex domain.

In another motivating example, we consider the well-known phase retrieval problem in x-ray crystallography, see for instance [10], [11], [12], [13], [14], [15]. The underlying principal of x-ray crystallography is that the information about the crystal structure can be obtained from its diffraction pattern by hitting the crystal by an x-ray beam. Due to physical limitations, typically only the intensity of the diffraction pattern can be measured but not its phase. This leads to a nonlinear relation

$$y_i = |\mathbf{a}_i^H \mathbf{x}|^2 = \mathbf{x}^H \mathbf{a}_i \mathbf{a}_i^H \mathbf{x}, \quad i = 1, \dots, N, \quad (5)$$

between the measurements $y_1, \dots, y_N \in \mathbb{R}$ and the structural information contained in $\mathbf{x} \in \mathbb{C}^n$. The complex vectors

$\mathbf{a}_1, \dots, \mathbf{a}_N \in \mathbb{C}^n$ are known and \mathbf{H} denotes the conjugate transpose. The mathematical problem of recovering \mathbf{x} from y_1, \dots, y_N , and $\mathbf{a}_1, \dots, \mathbf{a}_N$ is referred to as the phase retrieval problem. The traditional phase retrieval problem is known to be combinatorial [16].

If \mathbf{x} is sparse under an appropriate basis in (5), the problem is referred to as *compressive phase retrieval* (CPR) in [17], [18] or *quadratic compressed sensing* (QCS) in [19]. These algorithms can be applied to several important imaging applications, such as diffraction imaging [20], astronomy [21], [22], optics [23], x-ray tomography [24], microscopy [25], [26], [27], and quantum mechanics [28], to mention a few. As we will later show, our solution as a convex relaxation of (4), called *quadratic basis pursuit* (QBP), can be readily applied to solve this type of problems, namely, let $a_i = b_i = c_i = 0$, $\mathbf{Q}_i = \mathbf{a}_i \mathbf{a}_i^H$, $i = 1, \dots, N$.

A. Contributions

The main contribution of this paper is a novel convex technique for solving the sparse quadratic problem (4), namely, QBP. The proposed framework is not a greedy algorithm and inherits desirable properties, *e.g.*, perfect recovery, from BP and the traditional CS results. In comparison, most of the existing solutions for sparse nonlinear problems are greedy algorithms, and therefore their ability to give global convergence guarantees is limited.

Another contribution is an efficient numerical algorithm that solves the QBP problem and compares favorably to other existing sparse solvers in convex optimization. The algorithm is based on *alternating direction method of multipliers* (ADMM). Applying the algorithm to the complex CPR problem, we show that the QBP approach achieves the state-of-the-art result compared to other phase retrieval solutions when the measurements are under-sampled.

In Section II, we will first develop the main theory of QBP. In Section III, we present the ADMM algorithm. Finally, in Section IV, we conduct comprehensive experiments to validate the performance of the new algorithm on both synthetic and more practical imaging data.

B. Literature Review

To the best of our knowledge, this paper is the first work focusing on recovery of sparse signals from systems of general quadratic equations. Overall, the literature on nonlinear sparse problems and NLCS is also very limited. One of the first papers discussing these topics is [29]. They present a greedy gradient based algorithm for estimating the sparsest solution to a general nonlinear equation system. A greedy approach was also proposed in [30] for the estimation of sparse solutions of nonlinear equation systems. The work of [8] proposed several iterative hard-thresholding and sparse simplex pursuit algorithms. As the algorithms are nonconvex greedy solutions, the analysis of the theoretical convergence only concerns about their local behavior. In [9], the author also considered a generalization of the *restricted isometry property* (RIP) to support the use of similar iterative hard-thresholding algorithms for solving general NLCS problems.

Our paper is inspired by several recent works on CS applied to the phase retrieval problem [17], [31], [32], [19], [18], [33], [27], [34], [35], [36]. In particular, the generalization of compressive sensing to CS was first proposed in [17]. In [19], the problem was also referred to as QCS. These methods typically do not consider a general quadratic constraint as in (4) but a pure quadratic form (*i.e.*, $a_i = b_i = c_i = 0$, $i = 1, \dots, N$, in (4)).

In terms of the numerical algorithms that solves the CPR problem, most of the existing methods are greedy algorithms, where a solution to the underlying non-convex problem is sought by a sequence of local decisions [17], [31], [19], [33], [27], [36]. In particular, the QCS algorithm in [19] used a *lifting* technique similar to that in [37], [38], [39], [40] and *iterative rank minimization* resulting in a series of semidefinite programs (SDPs) that would converge to a local optimum.

The first work that applied the lifting technique to the PR and CPR problems was presented in [32]. Extensions of similar techniques were also studied in [41], [34]. The methods presented in our previous publications [1], [18] were also based on the lifting technique. It is important to note that the algorithms proposed in [32], [1], [18] are non-greedy global solutions, which are different from the previous local solutions [17], [19]. Our work was inspired by the solutions to phase retrieval via low-rank approximation in [32], [16], [42]. Given an oversampled phase retrieval problem, a lifting technique was used to relax the nonconvex problem with a SDP. The authors of [16], [42] also derived an upper-bound for the sampling rate that guarantees exact recovery in the noise-free case and stable recovery in the noisy case. Nevertheless, the work in [16], [42] only addressed the oversampled phase retrieval problem but not CPR or NLCS. The only similarities between our work and theirs are the lifting technique and convex relaxation. This lifting technique has also been used in other topics to convert nonconvex quadratic problems to SDPs, see for instance [43], [34]. The work presented in [32] and our previous contributions [1], [18] only discussed the CPR problem.

Finally, in [35], a *message passing* algorithm similar to that in CS was proposed to solve the compressive phase retrieval problem. The work in [44] further considered stability and uniqueness in real phase retrieval problems. CPR has also been shown useful in practice and we refer the interested reader to [17], [27] for two very nice contributions. Especially fascinating we find the work presented in [27] where the authors show how CPR can be used to facilitate sub-wavelength imaging in microscopy.

C. Notation and Assumptions

In this paper, we will use bold face to denote vectors and matrices and normal font for scalars. We denote the transpose of a real vector by \mathbf{x}^T and the conjugate transpose of a complex vector by \mathbf{x}^H . $\mathbf{X}_{i,j}$ is used to denote the (i, j) th element, $\mathbf{X}_{i,:}$ the i th row and $\mathbf{X}_{:,j}$ the j th column of a matrix \mathbf{X} , respectively. We will use the notation $\mathbf{X}_{i_1:i_2, j_1:j_2}$ to denote a submatrix constructed from rows i_1 to i_2 and columns j_1 to j_2 of \mathbf{X} . Given two matrices \mathbf{X} and \mathbf{Y} ,

we use the following fact that their product in the trace function commutes, namely, $\text{Tr}(\mathbf{X}\mathbf{Y}) = \text{Tr}(\mathbf{Y}\mathbf{X})$, under the assumption that the dimensions match. $\|\cdot\|_0$ counts the number of nonzero elements in a vector or matrix; similarly, $\|\cdot\|_1$ denotes the element-wise ℓ_1 -norm of a vector or matrix, *i.e.*, the sum of the magnitudes of the elements; whereas $\|\cdot\|$ represents the ℓ_2 -norm for vectors and the spectral norm for matrices.

II. QUADRATIC BASIS PURSUIT

A. Convex Relaxation via Lifting

As optimizing the ℓ_0 -norm function in (4) is known to be a combinatorial problem, in this section, we first introduce a convex relaxation of (4).

It is easy to see that the general quadratic constraint of (4) can be rewritten as the quadratic form:

$$y_i = [1 \quad \mathbf{x}^H] \begin{bmatrix} a_i & \mathbf{b}_i^H \\ \mathbf{c}_i & \mathbf{Q}_i \end{bmatrix} \begin{bmatrix} 1 \\ \mathbf{x} \end{bmatrix} \in \mathbb{C}, \quad i = 1, \dots, N. \quad (6)$$

Since each y_i is a scalar, we further have

$$y_i = \text{Tr} \left([1 \quad \mathbf{x}^H] \begin{bmatrix} a_i & \mathbf{b}_i^H \\ \mathbf{c}_i & \mathbf{Q}_i \end{bmatrix} \begin{bmatrix} 1 \\ \mathbf{x} \end{bmatrix} \right) \quad (7)$$

$$= \text{Tr} \left(\begin{bmatrix} a_i & \mathbf{b}_i^H \\ \mathbf{c}_i & \mathbf{Q}_i \end{bmatrix} \begin{bmatrix} 1 \\ \mathbf{x} \end{bmatrix} [1 \quad \mathbf{x}^H] \right). \quad (8)$$

Define $\Phi_i = \begin{bmatrix} a_i & \mathbf{b}_i^H \\ \mathbf{c}_i & \mathbf{Q}_i \end{bmatrix}$ and $\mathbf{X} = \begin{bmatrix} 1 \\ \mathbf{x} \end{bmatrix} [1 \quad \mathbf{x}^H]$, both matrices of dimensions $(n+1) \times (n+1)$. The operation that constructs \mathbf{X} from the vector $\begin{bmatrix} 1 \\ \mathbf{x} \end{bmatrix}$ is known as the *lifting* operator [37], [38], [39], [40]. By definition, \mathbf{X} is a Hermitian matrix, and it satisfies the constraints that $\mathbf{X}_{1,1} = 1$ and $\text{rank}(\mathbf{X}) = 1$. Hence, (4) can be rewritten as

$$\begin{aligned} & \min_{\mathbf{X}} \|\mathbf{X}\|_0 \\ \text{subj. to} & \quad y_i = \text{Tr}(\Phi_i \mathbf{X}), \quad i = 1, \dots, N, \\ & \quad \text{rank}(\mathbf{X}) = 1, \mathbf{X}_{1,1} = 1, \mathbf{X} \succeq 0. \end{aligned} \quad (9)$$

When the optimal solution \mathbf{X}^* is found, the unknown \mathbf{x} can be obtained by the rank-1 decomposition of \mathbf{X}^* via *singular value decomposition* (SVD).

The above problem is still non-convex and combinatorial. Therefore, solving it for any moderate size of n is impractical. Inspired by recent literature on matrix completion [45], [32], [16], [42] and sparse PCA [46], we relax the problem into the following convex *semidefinite program* (SDP):

$$\begin{aligned} & \min_{\mathbf{X}} \text{Tr}(\mathbf{X}) + \lambda \|\mathbf{X}\|_1 \\ \text{subj. to} & \quad y_i = \text{Tr}(\Phi_i \mathbf{X}), \quad i = 1, \dots, N, \\ & \quad \mathbf{X}_{1,1} = 1, \mathbf{X} \succeq 0, \end{aligned} \quad (10)$$

where $\lambda \geq 0$ is a design parameter. In particular, the trace of \mathbf{X} is a convex surrogate of the low-rank condition and $\|\mathbf{X}\|_1$ is the well-known convex surrogate for $\|\mathbf{X}\|_0$ in (9). We refer to the approach as *quadratic basis pursuit* (QBP).

One can further consider a noisy counterpart of the QBP problem, where some deviation between the measurements and the estimates is allowed. More specifically, we propose the

following *quadratic basis pursuit denoising* (QBPD) problem:

$$\begin{aligned} & \min_{\mathbf{X}} \text{Tr}(\mathbf{X}) + \lambda \|\mathbf{X}\|_1 \\ \text{subj. to} & \quad \sum_i^N \|y_i - \text{Tr}(\Phi_i \mathbf{X})\|^2 \leq \epsilon, \\ & \quad \mathbf{X}_{1,1} = 1, \mathbf{X} \succeq 0, \end{aligned} \quad (11)$$

for some $\epsilon > 0$.

B. Theoretical Analysis

In this section, we highlight some theoretical results derived for QBP. The analysis follows that of CS, and is inspired by derivations given in [16], [4], [32], [5], [47], [48], [6]. For further analysis on special cases of QBP and its noisy counterpart QBPD, please refer to [18].

First, it is convenient to introduce a linear operator B :

$$B : \mathbf{X} \in \mathbb{C}^{n \times n} \mapsto \{\text{Tr}(\Phi_i \mathbf{X})\}_{1 \leq i \leq N} \in \mathbb{C}^N. \quad (12)$$

We consider a generalization of the *restricted isometry property* (RIP) of the linear operator B .

Definition 1 (RIP). A linear operator $B(\cdot)$ as defined in (12) is (ϵ, k) -RIP if

$$\left| \frac{\|B(\mathbf{X})\|^2}{\|\mathbf{X}\|^2} - 1 \right| < \epsilon \quad (13)$$

for all $\|\mathbf{X}\|_0 \leq k$ and $\mathbf{X} \neq 0$.

We can now state the following theorem:

Theorem 2 (Recoverability/Uniqueness). Let $\bar{\mathbf{x}} \in \mathbb{C}^n$ be a solution to (4). If $\mathbf{X}^* \in \mathbb{C}^{(n+1) \times (n+1)}$ satisfies $\mathbf{y} = B(\mathbf{X}^*)$, $\mathbf{X}^* \succeq 0$, $\text{rank}(\mathbf{X}^*) = 1$, $\mathbf{X}_{1,1}^* = 1$ and if $B(\cdot)$ is a $(\epsilon, 2\|\mathbf{X}^*\|_0)$ -RIP linear operator with $\epsilon < 1$ then \mathbf{X}^* and $\bar{\mathbf{x}}$ are unique and $\mathbf{X}_{2:n+1,1}^* = \bar{\mathbf{x}}$.

Proof: Assume the contrary *i.e.*, $\mathbf{X}_{2:n+1,1}^* \neq \bar{\mathbf{x}}$ and hence that $\mathbf{X}^* \neq \begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H]$. It is clear that $\left\| \begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H] \right\|_0 \leq \|\mathbf{X}^*\|_0$ and hence $\left\| \begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H] - \mathbf{X}^* \right\|_0 \leq 2\|\mathbf{X}^*\|_0$. Since $\left\| \begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H] - \mathbf{X}^* \right\|_0 \leq 2\|\mathbf{X}^*\|_0$, we can apply the RIP inequality (13) on $\begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H] - \mathbf{X}^*$. If we use that $\mathbf{y} = B(\mathbf{X}^*) = B\left(\begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H]\right)$ and hence $B\left(\begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H] - \mathbf{X}^*\right) = 0$, we are led to the contradiction $1 < \epsilon$. We therefore conclude that $\mathbf{X}^* = \begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H]$, $\mathbf{X}_{2:n+1,1}^* = \bar{\mathbf{x}}$ and that \mathbf{X}^* and $\bar{\mathbf{x}}$ are unique. ■

We can also give a bound on the sparsity of $\bar{\mathbf{x}}$:

Theorem 3 (Bound on $\|\bar{\mathbf{x}}\|_0$ from above). Let $\bar{\mathbf{x}}$ be the sparsest solution to (4) and let $\tilde{\mathbf{X}}$ be the solution of QBPD (10). If $\tilde{\mathbf{X}}$ has rank 1 then $\|\tilde{\mathbf{X}}_{2:n+1,1}\|_0 \geq \|\bar{\mathbf{x}}\|_0$.

Proof: Let $\tilde{\mathbf{X}}$ be a rank-1 solution of QBP (10). By contradiction, assume $\|\tilde{\mathbf{X}}_{2:n+1,1}\|_0 < \|\bar{\mathbf{x}}\|_0$. Since $\tilde{\mathbf{X}}_{2:n+1,1}$ satisfies the constraints of (4), it is a feasible solution of (4). As assumed, $\tilde{\mathbf{X}}_{2:n+1,1}$ also gives a lower objective value than $\bar{\mathbf{x}}$ in

(4). This is a contradiction since $\bar{\mathbf{x}}$ was assumed to be the solution of (4). Hence we must have that $\|\tilde{\mathbf{X}}_{2:n+1,1}\|_0 \geq \|\bar{\mathbf{x}}\|_0$. ■

The following result now holds trivially:

Corollary 4 (Guaranteed recovery using RIP). *Let $\bar{\mathbf{x}}$ be the sparsest solution to (4). The solution of QBP $\tilde{\mathbf{X}}$ is equal to $\begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H]$ if it has rank 1 and $B(\cdot)$ is $(\epsilon, 2\|\tilde{\mathbf{X}}\|_0)$ -RIP with $\epsilon < 1$.*

Proof: This follows trivially from Theorem 2 by realizing that $\tilde{\mathbf{X}}$ satisfy all properties of \mathbf{X}^* . ■

Given the RIP analysis, it may be that the linear operator $B(\cdot)$ does satisfy the RIP property defined in Definition 1 with a small enough ϵ , as pointed out in [16]. In these cases, RIP-1 may be considered:

Definition 5 (RIP-1). *A linear operator $B(\cdot)$ is (ϵ, k) -RIP-1 if*

$$\left| \frac{\|B(\mathbf{X})\|_1}{\|\mathbf{X}\|_1} - 1 \right| < \epsilon \quad (14)$$

for all matrices $\mathbf{X} \neq 0$ and $\|\mathbf{X}\|_0 \leq k$.

Theorems 2–3 and Corollary 4 all hold with RIP replaced by RIP-1 and will not be restated in detail here. Instead, we summarize the most important property in the following theorem:

Theorem 6 (Upper bound and recoverability using RIP-1). *Let $\bar{\mathbf{x}}$ be the sparsest solution to (4). The solution of QBP (10), $\tilde{\mathbf{X}}$, is equal to $\begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H]$ if it has rank 1 and $B(\cdot)$ is $(\epsilon, 2\|\tilde{\mathbf{X}}\|_0)$ -RIP-1 with $\epsilon < 1$.*

Proof: The proof follows trivially from the proof of Theorem 2. ■

The RIP-type argument may be difficult to check for a given matrix and are more useful for claiming results for classes of matrices/linear operators. For instance, it has been shown that random Gaussian matrices satisfy the RIP with high probability. However, given realization of a random Gaussian matrix, it is indeed difficult to check if it actually satisfies the RIP. Two alternative arguments are the *spark condition* [3] and the *mutual coherence* [49], [50]. The spark condition usually gives tighter bounds but is known to be difficult to compute as well. On the other hand, mutual coherence may give less tight bounds, but is more tractable. We will focus on mutual coherence, which is defined as:

Definition 7 (Mutual coherence). *For a matrix \mathbf{A} , define the mutual coherence as*

$$\mu(\mathbf{A}) = \max_{1 \leq i, j \leq n, i \neq j} \frac{|\mathbf{A}_{:,i}^H \mathbf{A}_{:,j}|}{\|\mathbf{A}_{:,i}\| \|\mathbf{A}_{:,j}\|}. \quad (15)$$

Let \mathbf{B} be the matrix satisfying $\mathbf{y} = \mathbf{B}\mathbf{X}^s = B(\mathbf{X})$ with \mathbf{X}^s being the vectorized version of \mathbf{X} . We are now ready to state the following theorem:

Theorem 8 (Recovery using mutual coherence). *Let $\bar{\mathbf{x}}$ be the sparsest solution to (4). The solution of QBP (10), $\tilde{\mathbf{X}}$, is equal*

to $\begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H]$ if it has rank 1 and $\|\tilde{\mathbf{X}}\|_0 < 0.5(1+1/\mu(\mathbf{B}))$.

Proof: It follows from [49] [6, Thm. 5] that if

$$\|\tilde{\mathbf{X}}\|_0 < \frac{1}{2} \left(1 + \frac{1}{\mu(\mathbf{B})} \right) \quad (16)$$

then $\tilde{\mathbf{X}}$ is the sparsest solution to $\mathbf{y} = B(\mathbf{X})$. Since $\begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H]$ is by definition the sparsest rank 1 solution to $\mathbf{y} = B(\mathbf{X})$, it follows that $\tilde{\mathbf{X}} = \begin{bmatrix} 1 \\ \bar{\mathbf{x}} \end{bmatrix} [1 \quad \bar{\mathbf{x}}^H]$. ■

III. NUMERICAL ALGORITHMS

In addition to the above analysis of guaranteed recovery properties, a critical issue for practitioners is the efficiency of numerical solvers that can handle moderate-sized SDP problems. Several numerical solvers used in CS may be applied to solve nonsmooth SDPs, which include interior-point methods, e.g., used in CVX [51], gradient projection methods [52], and augmented Lagrangian methods (ALM) [52]. However, interior-point methods are known to scale badly to moderate-sized convex problems in general. Gradient projection methods also fail to meaningfully accelerate QBP due to the complexity of the projection operator. Alternatively, nonsmooth SDPs can be solved by ALM. However, the augmented primal and dual objective functions are still SDPs, which are equally expensive to solve in each iteration. There also exist a family of iterative approaches, often referred to as *outer approximation methods*, that successively approximate the solution of an SDP by solving a sequence of linear programs (see [53]). These methods approximate the positive semidefinite cone by a set of linear constraints and refine the approximation in each iteration by adding a new set of linear constraints. However, we have experienced slow convergence using these type of methods. In summary, QBP as a nonsmooth SDP is categorically more expensive to solve compared to the linear programs underlying CS, and the task exceeds the capability of many popular sparse optimization techniques.

In this paper, we propose a novel solver to the nonsmooth SDP underlying QBP via the *alternating directions method of multipliers* (ADMM, see for instance [54] and [55, Sec. 3.4]) technique. The motivation to use ADMM is two-fold:

- 1) It scales well to large data sets.
- 2) It is known for its fast convergence.

There are also a number of strong convergence results which further motivates the choice [54].

To set the stage for ADMM, let n denote the dimension of \mathbf{x} , and let N denote the number of measurements. Then, rewrite (10) to the equivalent SDP

$$\begin{aligned} \min_{\mathbf{X}_1, \mathbf{X}_2, \mathbf{Z}} \quad & f_1(\mathbf{X}_1) + f_2(\mathbf{X}_2) + g(\mathbf{Z}), \\ \text{subj. to} \quad & \mathbf{X}_1 - \mathbf{Z} = 0, \quad \mathbf{X}_2 - \mathbf{Z} = 0, \end{aligned} \quad (17)$$

where $\mathbf{X}_1 = \mathbf{X}_1^H \in \mathbb{C}^{(n+1) \times (n+1)}$, $\mathbf{X}_2 = \mathbf{X}_2^H \in$

$\mathbb{C}^{(n+1) \times (n+1)}$, $\mathbf{Z} = \mathbf{Z}^H \in \mathbb{C}^{(n+1) \times (n+1)}$, and

$$f_1(\mathbf{X}) \triangleq \begin{cases} \text{Tr}(\mathbf{X}) & \text{if } y_i = \text{Tr}(\Phi_i \mathbf{X}), i = 1, \dots, N \\ & \text{and } \mathbf{X}_{1,1} = 1 \\ \infty & \text{otherwise} \end{cases}$$

$$f_2(\mathbf{X}) \triangleq \begin{cases} 0 & \text{if } \mathbf{X} \succeq 0 \\ \infty & \text{otherwise} \end{cases}$$

$$g(\mathbf{Z}) \triangleq \lambda \|\mathbf{Z}\|_1.$$

Define two matrices \mathbf{Y}_1 and \mathbf{Y}_2 as the Lagrange multipliers of the two equality constraints in (17), respectively. Then the update rules of ADMM lead to the following:

$$\begin{aligned} \mathbf{X}_i^{l+1} &= \arg \min_{\mathbf{X}=\mathbf{X}^H} f_i(\mathbf{X}) + \text{Tr}(\mathbf{Y}_i^l (\mathbf{X} - \mathbf{Z}^l)) \\ &\quad + \frac{\rho}{2} \|\mathbf{X} - \mathbf{Z}^l\|^2, \\ \mathbf{Z}^{l+1} &= \arg \min_{\mathbf{Z}=\mathbf{Z}^H} g(\mathbf{Z}) + \sum_{i=1}^2 \text{Tr}(\mathbf{Y}_i^l \mathbf{Z}) \\ &\quad + \frac{\rho}{2} \|\mathbf{X}_i^{l+1} - \mathbf{Z}\|^2, \\ \mathbf{Y}_i^{l+1} &= \mathbf{Y}_i^l + \rho(\mathbf{X}_i^{l+1} - \mathbf{Z}^{l+1}), \end{aligned} \quad (18)$$

for $i = 1, 2$, where $\rho \geq 0$ is a parameter that enforces consensus between \mathbf{X}_1 , \mathbf{X}_2 , and \mathbf{Z} . Each of these steps has a tractable calculation. After some simple manipulations, we have:

$$\begin{aligned} \mathbf{X}_1^{l+1} &= \arg \min_{\mathbf{X}=\mathbf{X}^H} \|\mathbf{X} - (\mathbf{Z}^l - \frac{\mathbf{I} + \mathbf{Y}_1^l}{\rho})\|, \\ &\quad \text{subj. to } y_i = \text{Tr}(\Phi_i \mathbf{X}), \quad i = 1, \dots, N, \\ &\quad \mathbf{X}_{1,1} = 1. \end{aligned} \quad (19)$$

Let $\tilde{B} : \mathbb{C}^{(n+1) \times (n+1)} \rightarrow \mathbb{C}^{N+1}$ be the augmented linear operator such that $\tilde{B}(\mathbf{X}) = \begin{bmatrix} B(\mathbf{X}) \\ \mathbf{X}_{1,1} \end{bmatrix}$, where B is the linear operator defined by (12). Assuming that a feasible solution exists, and defining $\Pi_{\tilde{B}}$ as the orthogonal projection onto the convex set given by the linear constraints, i.e., $\begin{bmatrix} \mathbf{y} \\ 1 \end{bmatrix} = \tilde{B}(\mathbf{X})$,

the solution is: $\mathbf{X}_1^{l+1} = \Pi_{\tilde{B}}(\mathbf{Z}^l - \frac{\mathbf{I} + \mathbf{Y}_1^l}{\rho})$. This matrix-valued problem can be solved by converting the linear constraint on Hermitian matrices into an equivalent constraint on real-valued vectors.

Next,

$$\mathbf{X}_2^{l+1} = \arg \min_{\mathbf{X} \succeq 0} \left\| \mathbf{X} - \left(\mathbf{Z}^l - \frac{\mathbf{Y}_2^l}{\rho} \right) \right\| = \Pi_{PSD} \left(\mathbf{Z}^l - \frac{\mathbf{Y}_2^l}{\rho} \right), \quad (20)$$

where Π_{PSD} denotes the orthogonal projection onto the positive-semidefinite cone, which can easily be obtained via eigenvalue decomposition.

Finally, let $\bar{\mathbf{X}}^{l+1} = \frac{1}{2} \sum_{i=1}^2 \mathbf{X}_i^{l+1}$ and similarly $\bar{\mathbf{Y}}^l$. Then, the \mathbf{Z} update rule can be written:

$$\begin{aligned} \mathbf{Z}^{l+1} &= \arg \min_{\mathbf{Z}=\mathbf{Z}^H} \lambda \|\mathbf{Z}\|_1 + \rho \|\mathbf{Z} - (\bar{\mathbf{X}}^{l+1} + \frac{\bar{\mathbf{Y}}^l}{\rho})\|^2 \\ &= \text{soft}(\bar{\mathbf{X}}^{l+1} + \frac{\bar{\mathbf{Y}}^l}{\rho}, \frac{\lambda}{2\rho}) \end{aligned} \quad (21)$$

where $\text{soft}(\cdot)$ in the complex domain is defined with respect to a positive real scalar q as:

$$\text{soft}(x, q) = \begin{cases} 0 & \text{if } |x| \leq q, \\ \frac{|x| - q}{|x|} x & \text{otherwise.} \end{cases} \quad (22)$$

Note that if the first argument is a complex value, the soft operator is defined in terms of the magnitude rather than the sign and if it is a matrix, the the soft operator acts element-wise.

Setting $l = 1, \mathbf{X}_1^l = \mathbf{X}_2^l = \mathbf{Z}^l = \mathbf{I}$, where \mathbf{I} denotes the identity matrix, and $\rho^l = 1$, setting $l = 0$, the Hermitian matrices $\mathbf{X}_i^{l+1}, \mathbf{Z}_i^{l+1}, \mathbf{Y}_i^{l+1}$ can now be iteratively computed using the ADMM iterations (18). The stopping criterion of the algorithm is given by:

$$\|r^l\| \leq n\epsilon^{abs} + \epsilon^{rel} \max(\|\bar{\mathbf{X}}^l\|, \|\mathbf{Z}^l\|), \quad (23)$$

$$\|s^l\| \leq n\epsilon^{abs} + \epsilon^{rel} \|\bar{\mathbf{Y}}^l\|, \quad (24)$$

where $\epsilon^{abs}, \epsilon^{rel}$ are algorithm parameters set to 10^{-3} and r^l and s^l are the primal and dual residuals, respectively, as:

$$r^l = [\mathbf{X}_1^l - \mathbf{Z}^l \quad \mathbf{X}_2^l - \mathbf{Z}^l], \quad (25)$$

$$s^l = -\rho [\mathbf{Z}^l - \mathbf{Z}^{l-1} \quad \mathbf{Z}^l - \mathbf{Z}^{l-1}]. \quad (26)$$

We also update ρ according to the rule discussed in [54]:

$$\rho^{l+1} = \begin{cases} \tau_{incr} \rho^l & \text{if } \|r^l\| > \mu \|s^l\|, \\ \rho^l / \tau_{decr} & \text{if } \|s^l\| > \mu \|r^l\|, \\ \rho^l & \text{otherwise,} \end{cases} \quad (27)$$

where τ_{incr}, τ_{decr} , and μ are algorithm parameters. Values commonly used are $\mu = 10$ and $\tau_{incr} = \tau_{decr} = 2$.

In terms of the computational complexity of the ADMM algorithm, its inner loop calculates the updates of \mathbf{X}_i, \mathbf{Z} , and $\mathbf{Y}_i, i = 1, 2$. It is easy to see that its complexity is dominated by (19) and (20), which is bounded by $\mathcal{O}(n^2 N^2 + n^3)$, while the calculation of \mathbf{Z} and \mathbf{Y}_i is linear with respect to the number of their elements.

IV. EXPERIMENTS

In this section, we provide comprehensive experiments to validate the efficacy of the QBP algorithms in solving several representative nonlinear CS which depends quadratically on the unknown. We compare their performance primarily with two existing algorithms. As we mentioned in Section I, if an underdetermined nonlinear system is approximated up to the first order, the classical sparse solver in CS is basis pursuit. In NLCS literature, several greedy algorithms have been proposed for nonlinear systems. In this section, we choose to compare with the *iterative hard thresholding* (IHT) algorithm in [8] in Section IV-A and another greedy algorithm demonstrated in [27] in Section IV-C.¹

A. Nonlinear Compressive Sensing in Real Domain

In this experiment, we illustrate the concept of nonlinear compressive sensing. Assume that there is a cost associated with sampling and that we would like to recover $\mathbf{z}_0 \in \mathbb{R}^m$, related to our samples $y_i \in \mathbb{R}, i = 1, \dots, N$, via

$$y_i = f_i(\mathbf{z}_0), \quad i = 1, \dots, N, \quad (28)$$

¹Besides the comparisons shown here, we have also compared to a number of CPR algorithms [17], [36]. Not surprisingly, they performed badly on the general quadratic problems since they do not account for the linear term.

using as few samples as possible. Also, assume that there is a sparsifying basis $\mathbf{D} \in \mathbb{R}^{m \times n}$, possibly overcomplete, such that

$$\mathbf{z}_0 = \mathbf{D}\mathbf{x}_0, \quad \text{with } \mathbf{x}_0 \text{ sparse.} \quad (29)$$

Hence, we have


$$y_i = f_i(\mathbf{D}\mathbf{x}_0), \quad i = 1, \dots, N, \quad (30)$$

with \mathbf{x}_0 a sparse vector. If we approximate the nonlinear equation system (30) using a second order Maclaurin expansion we end up with a set of quadratic equations,


$$y_i = f_i(0) + \nabla f_i(0)\mathbf{D}\mathbf{x}_0 + \mathbf{x}_0^T \mathbf{D}^T \frac{\nabla^2 f_i(0)}{2} \mathbf{D}\mathbf{x}_0, \quad i = 1, \dots, N. \quad (31)$$

Hence, we can use QBP to recover \mathbf{x}_0 given $\{f_i(\mathbf{x}), y_i\}_{i=1}^N$ and \mathbf{D} .


In particular, let $\mathbf{D} = \mathbf{I}$, $n = m = 20$, $N = 25$, $f_i(\mathbf{x}) = a_i + \mathbf{b}_i^T \mathbf{x} + \mathbf{x}^T \mathbf{Q}_i \mathbf{x}$, $i = 1, \dots, N$, and generate $\{y_i\}_{i=1}^N$ by sampling $\{a_i, \mathbf{b}_i, \mathbf{Q}_i\}_{i=1}^N$ from a unitary Gaussian distribution. Let \mathbf{x}_0 be a binary vector with three elements different than zero. Given $\{y_i, a_i, \mathbf{b}_i, \mathbf{Q}_i\}_{i=1}^N$, the task is now to recover \mathbf{x}_0 . The results of this simulation are shown in Figure 1.




(a) Ground truth.




(b) QBP with $\lambda = 50$.



(c) QBP with $\lambda = 0$.



(d) Basis pursuit.



(e) Iterative hard thresholding.

Fig. 1. Estimated 20-D sparse signals measured in a simulated quadratic system of equations. The QBP solution perfectly recovers the ground truth with $\lambda = 50$, while the remaining algorithms fail to recover the correct sparse coefficients.

First, as the noiseless measurements are generated by a quadratic system of equations, it is not surprising that QBP perfectly recovers the sparse signal \mathbf{x}_0 when $\lambda = 50$. One may wonder whether in the 25-D ambient space, the solution \mathbf{x}_0 is unique. To show that the solution is not unique, we let $\lambda = 0$ and again apply QBP. As shown in Figure 1 (c), the solution is dense and it also satisfies the quadratic constraints. Therefore, we have verified that the system is underdetermined and there exist multiple solutions.

Second, in Figure 1 (d), we approximate (31) only up to the first order and set $\mathbf{Q}_i = 0, i = 1, \dots, N$. The approximation enables us to employ the classical basis pursuit algorithm in CS to seek the best 3-sparse estimate \mathbf{x} . As expected, the approximation is not accurate enough, and the estimate is far from the ground truth.

Third, we implement the iterative hard thresholding (IHT) algorithm in [8], and the correct number of nonzero coefficients in \mathbf{x}_0 is also provided to the algorithm. Its estimate is given in Figure 1 (e). As IHT is a greedy algorithm, its performance is affected by the initialization. In Figure 1 (e), the initial value is set by $\mathbf{x} = 0$, and the estimate is incorrect.

Finally, we note that the advantage of using general CS theory is that fewer samples are needed to recover a source signal from its observations. This remains true for NLCS presented in this paper. However, as (28) and (31) are nonlinear equation systems, typically $N \gg m$ measurements are required for recovering a unique solution. In the same simulation shown in Figure 1, one could ignore the sparsity constraint (*i.e.*, by letting $\lambda = 0$ in Figure 1 (c)), and it would require $N' = 40$ observations for QBP to recover the unique solution, which is exactly the ground-truth signal.

Clearly, Figure 1 is only able to illustrate one set of simulation results. To more systematically demonstrate the accuracy of the four algorithms in probability, a Monte Carlo simulation is performed that repeats the above simulation but with different randomly generated \mathbf{x}_0 and $\{a_i, \mathbf{b}_i, \mathbf{Q}_i\}$. Table I shows the rates of successful recovery. We can see QBP achieves the highest success rate, which is followed by IHT. BP and the dense QBP solution basically fail to return enough good results. $\lambda = 50$ was used in all trials.

TABLE I
THE PERCENTAGE OF CORRECTLY RECOVERING \mathbf{x}_0 IN 100 TRIALS.

Method	QBP ($\lambda = 50$)	QBP ($\lambda = 0$)	BP	IHT
Success rate	79%	5%	3%	54%

B. The Shepp-Logan Phantom

In this experiment, we consider recovery of images from random samples. More specifically, we formulate an example of the CPR problem in the QBP framework using the Shepp-Logan phantom. Our goal is to show that using the QBP algorithm provides approximate solutions that are visually close to the ground-truth images.

Consider the ground-truth image in Figure 2. This 30×30 Shepp-Logan phantom has a 2D Fourier transform with 100 nonzero complex coefficients. We generate N linear combinations of pixels, and then measure the square of the measurements. This relationship can be written as:

$$\mathbf{y} = |\mathbf{A}\mathbf{x}|^2 = \{\mathbf{x}^H \mathbf{a}_i \mathbf{a}_i^H \mathbf{x}\}_{1 \leq i \leq N}, \quad (32)$$

where $\mathbf{A} = \mathbf{R}\mathbf{F}$ is the concatenation of a random matrix \mathbf{R} and the Fourier basis \mathbf{F} , and the image $\mathbf{F}\mathbf{x}$ is represented as a stacked vector in the 900-D complex domain. The CPR problem minimizes the following objective function:

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \quad \text{subj. to} \quad \mathbf{y} = |\mathbf{A}\mathbf{x}|^2 \in \mathbb{R}^N. \quad (33)$$

Previously, an SDP solution to the non sparse phase retrieval problem was proposed in [16], which is called *PhaseLift*. In a sense, PhaseLift can be viewed as a special case of the QBP solution in (10) where $\lambda = 0$, namely, the sparsity constraint is not enforced. In Figure 2 (b), the recovered result using PhaseLift is shown with $N = 2400$.

To compare visually the performance of the QBP solution when the sparsity constraint is properly enforced, two recovered results are shown in Figure 2 (c) and (d) with $N = 2400$ and 1500, respectively. Note that the number of measurements with respect to the sparsity in \mathbf{x} is too low for both QBP and PhaseLift to perfectly recover \mathbf{x} . Therefore, in this case, we employ the noisy version of the algorithm QBDP to recover the image. We can clearly see from the illustrations that QBDP provides a much better approximation and outperforms PhaseLift visually even though it uses considerably fewer measurements.




Fig. 2. Recovery of a Shepp-Logan Image by PhaseLift and QBDP.

C. Subwavelength Imaging

In this example, we present an example in sub-wavelength coherent diffractive imaging. The experiment and the data collection were conducted by [27].

Let y_i , $i = 1, \dots, N$, be intensity samples of a 2D diffraction pattern. The diffraction pattern is the result of a 532 nm

laser beam passing through an arrangement of holes made on a opaque piece of glass. The task is to decide the location of the holes out of a number of possible locations.

It can be shown that the relation between the intensity measurements and the arrangements of holes is of the following type:

$$y_i = |\mathbf{a}_i^H \mathbf{x}|^2, \quad i = 1, \dots, N, \quad (34)$$

where $y_i \in \mathbb{R}$, $i = 1, \dots, N$, are intensity measurements, $\mathbf{a}_i \in \mathbb{C}^n$, $i = 1, \dots, N$, are known complex vectors and $\mathbf{x} \in \mathbb{R}^n$, is the sought entity, each element giving the likelihood of a hole at a given location.

We use QBDP with $\varepsilon = 0.0012$ and $\lambda = 100$. 89 measurements were selected by taking every 200th intensity measurement from the dataset of [27]. The quantity \mathbf{x} is from the setup of the experiment known to be real and $a_i = b_i = c_i = 0$. We hence have

$$y_i = \mathbf{x}^T \mathbf{Q}_i \mathbf{x} = |\mathbf{a}_i^H \mathbf{x}|^2, \quad i = 1, \dots, N, \quad (35)$$

with $\mathbf{Q}_i = \mathbf{a}_i \mathbf{a}_i^H \in \mathbb{C}^{n \times n}$, $i = 1, \dots, N$, and $\mathbf{x} \in \mathbb{R}^n$.

The resulting estimate is given to the left in Figure 3. The result deviates from the ground truth and the result presented in [27] (shown in Figure 3 right), and it actually finds a more sparse pattern. It is interesting to note that both estimates are however within the noise level estimated in [27]:

$$\frac{1}{N} \sum_i (y_i - |\mathbf{a}_i^H \mathbf{x}|^2)^2 \leq 1.8 \times 10^{-6}. \quad (36)$$

Therefore, under the same noise assumptions, the two solutions are equally likely to lead to the same observations \mathbf{y} . However, knowing that there is a solution within the noise level that is indeed sparser than the ground-truth pattern, it should *not* be the optimal solution to have recovered the ground truth, since there exists a sparser solution.




Fig. 3. The estimated sparse vector \mathbf{x} . The crosses mark possible positions for holes, while the dots represent the recovered nonzero coefficients. **Left:** Recovered pattern by QBDP. Note that this estimate is sparser than the ground truth but within the estimated noise level. **Right:** Recovered pattern by the compressive phase retrieval method used in [27].

V. CONCLUSION

Classical compressive sensing assumes a linear relation between samples and the unknowns. The ability to more accurately characterize nonlinear models has the potential to improve the results in both existing compressive sensing applications and those where a linear approximation does not suffice, *e.g.*, phase retrieval.

This paper presents an extension of classical compressive sensing to quadratic relations or second order Taylor expansions of the nonlinearity relating measurements and the unknowns. The novel extension is based on lifting and convex relaxations and the final formulation takes the form of a SDP. The proposed method, quadratic basis pursuit, inherits properties of basis pursuit and classical compressive sensing and conditions for perfect recovery etc are derived. We also give an efficient numerical implementation.

ACKNOWLEDGEMENT

The authors would like to acknowledge useful discussions and inputs from Yonina C. Eldar, Mordechai Segev, Laura Waller, Filipe Maia, Stefano Marchesini and Michael Lustig. We also want to acknowledge the authors of [27] for kindly sharing their data with us.

Ohlsson is partially supported by the Swedish Research Council in the Linnaeus center CADICS, the European Research Council under the advanced grant LEARN, contract 267381, by a postdoctoral grant from the Sweden-America Foundation, donated by ASEA's Fellowship Fund, and by a postdoctoral grant from the Swedish Research Council. Yang is supported in part by ARO 63092-MA-II and DARPA FA8650-11-1-7153.

REFERENCES

- [1] H. Ohlsson, A. Yang, R. Dong, and S. S. Sastry, "CPRL — an extension of compressive sensing to the phase retrieval problem," in *Advances in Neural Information Processing Systems* 25, P. Bartlett, F. Pereira, C. Burges, L. Bottou, and K. Weinberger, Eds., 2012, pp. 1376–1384.
- [2] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM Journal on Computing*, vol. 24, no. 2, pp. 227–234, 1995.
- [3] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1998.
- [4] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, pp. 489–509, Feb. 2006.
- [5] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [6] A. Bruckstein, D. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Review*, vol. 51, no. 1, pp. 34–81, 2009.
- [7] Y. C. Eldar and G. Kutyniok, *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012.
- [8] A. Beck and Y. C. Eldar, "Sparsity constrained nonlinear optimization: Optimality conditions and algorithms," Tech. Rep. arXiv:1203.4580, 2012.
- [9] T. Blumensath, "Compressed sensing with nonlinear observations and related nonlinear optimization problems," Tech. Rep. arXiv:1205.1650, 2012.
- [10] D. Kohler and L. Mandel, "Source reconstruction from the modulus of the correlation function: a practical approach to the phase problem of optical coherence theory," *Journal of the Optical Society of America*, vol. 63, no. 2, pp. 126–134, 1973.
- [11] R. Gonsalves, "Phase retrieval from modulus data," *Journal of Optical Society of America*, vol. 66, no. 9, pp. 961–964, 1976.
- [12] R. Gerchberg and W. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik*, vol. 35, pp. 237–246, 1972.
- [13] J. Fienup, "Phase retrieval algorithms: a comparison," *Applied Optics*, vol. 21, no. 15, pp. 2758–2769, 1982.
- [14] S. Marchesini, "Phase retrieval and saddle-point optimization," *Journal of the Optical Society of America A*, vol. 24, no. 10, pp. 3289–3296, 2007.
- [15] R. Balan, P. Casazza, and D. Edidin, "On signal reconstruction without phase," *Applied and Computational Harmonic Analysis*, vol. 20, pp. 345–356, 2006.
- [16] E. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," Stanford University, Tech. Rep. arXiv:1109.4499, Sep. 2011.
- [17] M. Moravec, J. Romberg, and R. Baraniuk, "Compressive phase retrieval," in *SPIE International Symposium on Optical Science and Technology*, 2007.
- [18] H. Ohlsson, A. Y. Yang, R. Dong, and S. Sastry, "Compressive Phase Retrieval From Squared Output Measurements Via Semidefinite Programming," University of California, Berkeley, Tech. Rep. arXiv:1111.6323, Nov. 2011.
- [19] Y. Shechtman, Y. C. Eldar, A. Szameit, and M. Segev, "Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing," *Opt. Express*, vol. 19, no. 16, pp. 14807–14822, Aug 2011.
- [20] 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 Crystallographica Section A*, vol. 63, no. 4, pp. 306–314, Jul. 2007.
- [21] J. Dainty and J. Fienup, "Phase retrieval and image reconstruction for astronomy," in *Image Recovery: Theory and Application*, e. H. Stark, Ed. Academic Press, New York, 1987.
- [22] J. R. Fienup, J. C. Marron, T. J. Schulz, and J. H. Seldin, "Hubble space telescope characterized by using phase-retrieval algorithms," *Applied Optics*, vol. 32, no. 10, pp. 1747–1767, Apr 1993.
- [23] A. Walther, "The question of phase retrieval in optics," *Optica Acta*, vol. 10, pp. 41–49, 1963.
- [24] M. Dierolf, A. Menzel, P. Thibault, P. Schneider, C. M. Kewish, R. Wepf, O. Bunk, and F. Pfeiffer, "Ptychographic x-ray computed tomography at the nanoscale," *Nature*, vol. 467, pp. 436–439, 2010.
- [25] 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," *Annual Review of Physical Chemistry*, vol. 59, no. 1, pp. 387–410, 2008.
- [26] J. Antonello, M. Verhaegen, R. Fraanje, T. van Werkhoven, H. C. Gerritsen, and C. U. Keller, "Semidefinite programming for model-based sensorless adaptive optics," *J. Opt. Soc. Am. A*, vol. 29, no. 11, pp. 2428–2438, Nov. 2012.
- [27] A. Szameit, Y. Shechtman, E. Osherovich, E. Bullkich, P. Sidorenko, H. Dana, S. Steiner, E. B. Kley, S. Gazit, T. Cohen-Hyams, S. Shoham, M. Zibulevsky, I. Yavneh, Y. C. Eldar, O. Cohen, and M. Segev, "Sparsity-based single-shot subwavelength coherent diffractive imaging," *Nature Materials*, vol. 11, no. 5, pp. 455–459, May 2012.
- [28] J. Corbett, "The Pauli problem, state reconstruction and quantum-real numbers," *Reports on Mathematical Physics*, vol. 57, no. 1, pp. 53–68, 2006.
- [29] T. Blumensath and M. E. Davies, "Gradient pursuit for non-linear sparse signal modelling," in *European Signal Processing Conference*, Lausanne, Switzerland, Apr. 2008.
- [30] L. Li and B. Jafarpour, "An iteratively reweighted algorithm for sparse reconstruction of subsurface flow properties from nonlinear dynamic data," *CoRR*, vol. abs/0911.2270, 2009.
- [31] S. Marchesini, "Ab Initio Undersampled Phase Retrieval," *Microscopy and Microanalysis*, vol. 15, Jul. 2009.
- [32] A. Chai, M. Moscoso, and G. Papanicolaou, "Array imaging using intensity-only measurements," Stanford University, Tech. Rep., 2010.
- [33] E. Osherovich, Y. Shechtman, A. Szameit, P. Sidorenko, E. Bullkich, S. Gazit, S. Shoham, E. Kley, M. Zibulevsky, I. Yavneh, Y. Eldar, O. Cohen, and M. Segev, "Sparsity-based single-shot subwavelength coherent diffractive imaging," in *2012 Conference on Lasers and Electro-Optics (CLEO)*, San Jose, CA, USA, May 2012.
- [34] K. Jaganathan, S. Oymak, and B. Hassibi, "Recovery of Sparse 1-D Signals from the Magnitudes of their Fourier Transform," *ArXiv e-prints*, Jun. 2012.
- [35] P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate message passing," in *Proceedings of Allerton Conference on Communication, Control, and Computing*, Monticello, IL, USA, Oct. 2012.
- [36] Y. Shechtman, A. Beck, and Y. C. Eldar, "GESPAR: Efficient Phase Retrieval of Sparse Signals," *ArXiv e-prints*, Jan. 2013.
- [37] N. Shor, "Quadratic optimization problems," *Soviet Journal of Computer and Systems Sciences*, vol. 25, pp. 1–11, 1987.
- [38] L. Lovász and A. Schrijver, "Cones of matrices and set-functions and 0-1 optimization," *SIAM Journal on Optimization*, vol. 1, pp. 166–190, 1991.

- [39] Y. Nesterov, "Semidefinite relaxation and nonconvex quadratic optimization," *Optimization Methods & Software*, vol. 9, pp. 141–160, 1998.
- [40] M. X. Goemans and D. P. Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *J. ACM*, vol. 42, no. 6, pp. 1115–1145, Nov. 1995.
- [41] X. Li and V. Voroninski, "Sparse Signal Recovery from Quadratic Measurements via Convex Programming," *ArXiv e-prints*, Sep. 2012.
- [42] E. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," Stanford University, Tech. Rep. arXiv:1109.0573, Sep. 2011.
- [43] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase Recovery, MaxCut and Complex Semidefinite Programming," *ArXiv e-prints*, Jun. 2012.
- [44] Y. C. Eldar and S. Mendelson, "Phase Retrieval: Stability and Recovery Guarantees," *ArXiv e-prints*, Nov. 2012.
- [45] E. Candès and B. Recht, "Exact matrix completion via convex optimization," *CoRR*, vol. abs/0805.4471, 2008.
- [46] A. d'Aspremont, L. El Ghaoui, M. Jordan, and G. Lanckriet, "A direct formulation for Sparse PCA using semidefinite programming," *SIAM Review*, vol. 49, no. 3, pp. 434–448, 2007.
- [47] E. Candès, "The restricted isometry property and its implications for compressed sensing," *Comptes Rendus Mathématique*, vol. 346, no. 9–10, pp. 589–592, 2008.
- [48] R. Berinde, A. Gilbert, P. Indyk, H. Karloff, and M. Strauss, "Combining geometry and combinatorics: A unified approach to sparse signal recovery," in *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, Sep. 2008, pp. 798–805.
- [49] D. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 -minimization," *PNAS*, vol. 100, no. 5, pp. 2197–2202, Mar. 2003.
- [50] E. Candès, X. Li, Y. Ma, and J. Wright, "Robust Principal Component Analysis?" *Journal of the ACM*, vol. 58, no. 3, 2011.
- [51] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21," <http://cvxr.com/cvx>, Aug. 2010.
- [52] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific, 1999.
- [53] H. Konno, J. Gotoh, T. Uno, and A. Yuki, "A cutting plane algorithm for semi-definite programming problems with applications to failure discriminant analysis," *Journal of Computational and Applied Mathematics*, vol. 146, no. 1, pp. 141–154, 2002.
- [54] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, 2011.
- [55] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997.