

Compressive Phase Retrieval From Squared Output Measurements Via Semidefinite Programming[★]

Henrik Ohlsson^{*,**} Allen Y. Yang^{*} Roy Dong^{*}
S. Shankar Sastry^{*}

^{*} *Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, CA, USA, (email:
{ohlsson,yang,roydong,sastry}@eecs.berkeley.edu).*

^{**} *Division of Automatic Control, Department of Electrical
Engineering, Linköping University, Sweden, (e-mail:
ohlsson@isy.liu.se).*

Abstract: Given a linear system in a real or complex domain, linear regression aims to recover the model parameters from a set of observations. Recent studies in compressive sensing have successfully shown that under certain conditions, a linear program, namely, ℓ_1 -minimization, guarantees recovery of sparse parameter signals even when the system is underdetermined. In this paper, we consider a more challenging problem: when the phase of the output measurements from a linear system is omitted. Using a lifting technique, we show that even though the phase information is missing, the sparse signal can be recovered exactly by solving a semidefinite program when the sampling rate is sufficiently high. This is an interesting finding since the exact solutions to both sparse signal recovery and phase retrieval are combinatorial. The results extend the type of applications that compressive sensing can be applied to those where only output magnitudes can be observed. We demonstrate the accuracy of the algorithms through extensive simulation and a practical experiment.

Keywords: Phase Retrieval; Compressive Sensing; Semidefinite Programming.

1. INTRODUCTION

Linear models, *e.g.*, $\mathbf{y} = A\mathbf{x}$, are by far the most used and useful type of model. The main reasons for this are their simplicity of use and identification. For the identification, the least-squares (LS) estimate in a complex domain is computed by¹

$$\mathbf{x}_{\text{ls}} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{y} - A\mathbf{x}\|_2^2 \in \mathbb{C}^n, \quad (1)$$

assuming the output $\mathbf{y} \in \mathbb{C}^N$ and $A \in \mathbb{C}^{N \times n}$ are given. Further, the LS problem has a unique solution if the system is full rank and not underdetermined, *i.e.*, $N \geq n$.

Consider the alternative scenario when the system is underdetermined, *i.e.*, $n > N$. The least squares solution is no longer unique in this case, and additional knowledge has

[★] Ohlsson is partially supported by the Swedish foundation for strategic research in the center MOVIII, the Swedish Research Council in the Linnaeus center CADICS, the European Research Council under the advanced grant LEARN, contract 267381, and a postdoctoral grant from the Sweden-America Foundation, donated by ASEA's Fellowship Fund. Sastry and Yang are partially supported by an ARO MURI grant W911NF-06-1-0076. Dong is supported by the NSF Graduate Research Fellowship under grant DGE 1106400, and by the Team for Research in Ubiquitous Secure Technology (TRUST), which receives support from NSF (award number CCF-0424422).

¹ Our derivation in this paper is primarily focused on complex signals, but the results should be easily extended to real domain signals.

to be used to determine a unique model parameter. Ridge regression or Tikhonov regression [Hoerl and Kennard, 1970] is one of the traditional methods to apply in this case, which takes the form

$$\mathbf{x}_r = \underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{y} - A\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_2^2, \quad (2)$$

where $\lambda > 0$ is a scalar parameter that decides the trade off between fit (the first term) and the ℓ_2 -norm of \mathbf{x} (the second term).

Thanks to the ℓ_2 -norm regularization, ridge regression is known to pick up solutions with small energy that satisfy the linear model. In a more recent approach stemming from the LASSO [Tibsharani, 1996] and compressive sensing (CS) [Candès et al., 2006, Donoho, 2006], another convex regularization criterion has been widely used to seek the *sparsest* parameter vector, which takes the form

$$\mathbf{x}_{\ell_1} = \underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{y} - A\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1. \quad (3)$$

Depending on the choice of the weight parameter λ , the program (3) has been known as the LASSO by Tibsharani [1996], *basis pursuit denoising* (BPDN) by Chen et al. [1998], or ℓ_1 -minimization (ℓ_1 -min) by Candès et al. [2006]. In recent years, several pioneering works have contributed to efficiently solving sparsity minimization problems such as Tropp [2004], Beck and Teboulle [2009], Bruckstein et al. [2009], especially when the system parameters and observations are in high-dimensional spaces.

In this paper, we consider a more challenging problem. We still seek a linear model $\mathbf{y} = A\mathbf{x}$. Rather than assuming that \mathbf{y} is given we will assume that only the squared magnitude of the output is given

$$b_i = |y_i|^2 = |\langle \mathbf{x}, \mathbf{a}_i \rangle|^2, \quad i = 1, \dots, N, \quad (4)$$

where $A^T = [\mathbf{a}_1, \dots, \mathbf{a}_N] \in \mathbb{C}^{n \times N}$ and $\mathbf{y}^T = [y_1, \dots, y_N] \in \mathbb{C}^{1 \times N}$. This is clearly a more challenging problem since the phase of \mathbf{y} is lost when only the (squared) magnitude is available. A classical example is that \mathbf{y} represents the Fourier transform of \mathbf{x} , and that only the Fourier transform modulus is observable. This scenario arises naturally in several practical applications such as optics (Walther [1963], Millane [1990]), coherent diffraction imaging (Fienup [1987]), astronomical imaging (Dainty and Fienup [1987]), and is known as the *phase retrieval* problem.

We note that in general phase cannot be uniquely recovered regardless whether the linear model is overdetermined or not. A simple example to see this, is if $\mathbf{x}_0 \in \mathbb{C}^n$ is a solution to $\mathbf{y} = A\mathbf{x}$, then for any scalar $c \in \mathbb{C}$ on the unit circle $c\mathbf{x}_0$ leads to the same squared output \mathbf{b} . As mentioned in Candès et al. [2011a], when the dictionary A represents the unitary discrete Fourier transform (DFT), the ambiguities may represent time-reversed solutions or time-shifted solutions of the ground truth signal \mathbf{x}_0 . These global ambiguities caused by losing the phase information are considered acceptable in phase retrieval applications. From now on, when we talk about the solution to the phase retrieval problem, it is the solution up to a global phase. Accordingly, a unique solution is a solution unique up to a global phase.

Further note that since (4) is nonlinear in the unknown \mathbf{x} , $N \gg n$ measurements are in general needed for a unique solution. When the number of measurements N are fewer than necessary for a unique solution, additional assumptions are needed to select one of the solutions (just like in Tikhonov, LASSO and CS).

Finally, we note that the exact solution to either CS and phase retrieval is combinatorially expensive (Chen et al. [1998], Candès et al. [2011b]). Therefore, the goal of this work is to answer the following question: *Can we effectively recover a sparse parameter vector \mathbf{x} of a linear system up to a global ambiguity using its squared magnitude output measurements via convex programming?* The problem is referred to as *compressive phase retrieval* (CPR).

The main contribution of the paper is a *convex formulation* of the sparse phase retrieval problem. Using a lifting technique, the NP-hard problem is relaxed as a semidefinite program. Through extensive experiments, we compare the performance of our CPR algorithm with traditional CS and PhaseLift algorithms. The results extend the type of applications that compressive sensing can be applied to, namely, applications where only magnitudes can be observed.

1.1 Background

Our work is motivated by the ℓ_1 -min problem in CS and a recent PhaseLift technique in phase retrieval by Candès et al. [2011b]. On one hand, the theory of CS and ℓ_1 -min has been one of the most visible research topics in recent

years. There are several comprehensive review papers that cover the literature of CS and related optimization techniques in linear programming. The reader is referred to the works of Candès and Wakin [2008], Bruckstein et al. [2009], Loris [2009], Yang et al. [2010]. On the other hand, the fusion of phase retrieval and matrix completion is a novel topic that has recently been studied in a selected few papers, such as Chai et al. [2010], Candès et al. [2011b,a]. The fusion of phase retrieval and CS was discussed in Moravec et al. [2007]. In the rest of the section, we briefly review the phase retrieval literature and its recent connections with CS and matrix completion.

Phase retrieval has been a longstanding problem in optics and x-ray crystallography since the 1970s [Kohler and Mandel, 1973, Gonsalves, 1976]. Early methods to recover the phase signal using Fourier transform mostly relied on additional information about the signal, such as band limitation, nonzero support, real-valuedness, and nonnegativity. The Gerchberg-Saxton algorithm was one of the popular algorithms that alternates between the Fourier and inverse Fourier transforms to obtain the phase estimate iteratively [Gerchberg and Saxton, 1972, Fienup, 1982]. One can also utilize steepest-descent methods to minimize the squared estimation error in the Fourier domain [Fienup, 1982, Marchesini, 2007]. Common drawbacks of these iterative methods are that they may not converge to the global solution, and the rate of convergence is often slow. Alternatively, Balan et al. [2006] have studied a frame-theoretical approach to phase retrieval, which necessarily relied on some special types of measurements.

More recently, phase retrieval has been framed as a low-rank matrix completion problem in Chai et al. [2010], Candès et al. [2011a,b]. Given a system, a lifting technique was used to approximate the linear model constraint as a semidefinite program (SDP), which is similar to the CPR objective function (10) only without the sparsity constraint. The authors also derived the upper-bound for the sampling rate that guarantees exact recovery in the noise-free case and stable recovery in the noisy case.

We are aware of the work by Moravec et al. [2007], which has considered compressive phase retrieval on a random Fourier transform model. Leveraging the sparsity constraint, the authors proved that an upper-bound of $O(k^2 \log(4n/k^2))$ random Fourier modulus measurements to uniquely specify k -sparse signals. Moravec et al. [2007] also proposed a compressive phase retrieval algorithm. Their solution largely follows the development of ℓ_1 -min in CS, and it alternates between the domain of solutions that give rise to the same squared output and the domain of an ℓ_1 -ball with a fixed ℓ_1 -norm. However, the main limitation of the algorithm is that it tries to solve a nonconvex optimization problem which assumes the ℓ_1 -norm of the true signal is known.

2. CPR VIA SDP

In the noise free case, the phase retrieval problem takes the form of the feasibility problem:

$$\text{find } \mathbf{x} \quad \text{subj. to } \mathbf{b} = |A\mathbf{x}|^2 = \{\mathbf{a}_i^H \mathbf{x} \mathbf{x}^H \mathbf{a}_i\}_{1 \leq i \leq N}, \quad (5)$$

where $\mathbf{b}^T = [b_1, \dots, b_N] \in \mathbb{R}^{1 \times N}$. This is a combinatorial problem to solve: Even in the real domain with the sign

of the measurements $\{\alpha_i\}_{i=1}^N \subset \{-1, 1\}$, one would have to try out combinations of sign sequences until one that satisfies

$$\alpha_i \sqrt{b_i} = \mathbf{a}_i^T \mathbf{x}, \quad i = 1, \dots, N, \quad (6)$$

for some $\mathbf{x} \in \mathbb{R}^n$ has been found. For any practical size of data sets, this combinatorial problem is intractable.

Since (5) is nonlinear in the unknown \mathbf{x} , $N \gg n$ measurements are in general needed for a unique solution. When the number of measurements N are fewer than necessary for a unique solution, additional assumptions are needed to select one of the solutions. Motivated by compressive sensing, we here choose to seek the sparsest solution of CPR satisfying (5) or, equivalent, the solution to

$$\min_{\mathbf{x}} \|\mathbf{x}\|_0, \quad \text{subj. to } \mathbf{b} = |\mathbf{A}\mathbf{x}|^2 = \{\mathbf{a}_i^H \mathbf{x} \mathbf{x}^H \mathbf{a}_i\}_{1 \leq i \leq N}. \quad (7)$$

As the counting norm $\|\cdot\|_0$ is not a convex function, following the ℓ_1 -norm relaxation in CS, (7) can be relaxed as

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1, \quad \text{subj. to } \mathbf{b} = |\mathbf{A}\mathbf{x}|^2 = \{\mathbf{a}_i^H \mathbf{x} \mathbf{x}^H \mathbf{a}_i\}_{1 \leq i \leq N}. \quad (8)$$

Note that (8) is still not a linear program, as its equality constraint is not a linear equation. In the literature, a lifting technique has been extensively used to reframe problems such as (8) to a standard form in semidefinite programming, such as in Sparse PCA [d'Aspremont et al., 2007].

More specifically, given the ground truth signal $\mathbf{x}_0 \in \mathbb{C}^n$, let $X_0 \doteq \mathbf{x}_0 \mathbf{x}_0^H \in \mathbb{C}^{n \times n}$ be a rank-1 semidefinite matrix. Then the CPR problem can be cast as²

$$\begin{aligned} \min_X \quad & \|X\|_1 \\ \text{subj. to} \quad & b_i = \text{Tr}(\mathbf{a}_i^H X \mathbf{a}_i), \quad i = 1, \dots, N, \\ & \text{rank}(X) = 1, X \succeq 0. \end{aligned} \quad (9)$$

This is of course still a non-convex problem due to the rank constraint. The lifting approach addresses this issue by replacing $\text{rank}(X)$ with $\text{Tr}(X)$. For a semidefinite matrix, $\text{Tr}(X)$ is equal to the sum of the eigenvalues of X . This leads to an SDP

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

where we further denote $\Phi_i \doteq \mathbf{a}_i \mathbf{a}_i^H \in \mathbb{C}^{n \times n}$ and where $\lambda > 0$ is a design parameter. Finally, the estimate of \mathbf{x} can be found by computing the rank-1 decomposition of X via singular value decomposition. We will refer to the formulation (10) as compressive phase retrieval via lifting (CPRL).

We compare (10) to a recent solution of PhaseLift by Candès et al. [2011b]. In Candès et al. [2011b], a similar objective function was employed for phase retrieval:

$$\begin{aligned} \min_X \quad & \text{Tr}(X) \\ \text{subj. to} \quad & b_i = \text{Tr}(\Phi_i X), \quad i = 1, \dots, N, \\ & X \succeq 0, \end{aligned} \quad (11)$$

² $\|X\|_1$ for a matrix X denotes the entry-wise ℓ_1 -norm in this paper.

albeit the source signal was not assumed sparse. Using the lifting technique to construct the SDP relaxation of the NP-hard phase retrieval problem, with high probability, the program (11) recovers the exact solution (sparse or dense) if the number of measurements N is at least of the order of $O(n \log n)$. The region of success is visualized in Figure 1 as region I with a thick solid line.

If \mathbf{x} is sufficiently sparse and random Fourier dictionaries are used for sampling, Moravec et al. [2007] showed that in general the signal is uniquely defined if the number of squared magnitude output measurements \mathbf{b} exceeds the order of $O(k^2 \log(4n/k^2))$. This lower bound for the region of success of CPR is illustrated by the dash line in Figure 1.

Finally, the motivation for introducing the ℓ_1 -norm regularization in (10) is to be able to solve the sparse phase retrieval problem for N smaller than what PhaseLift requires. However, one will not be able to solve the compressive phase retrieval problem in region III below the dashed curve. Therefore, our target problems lie in region II.

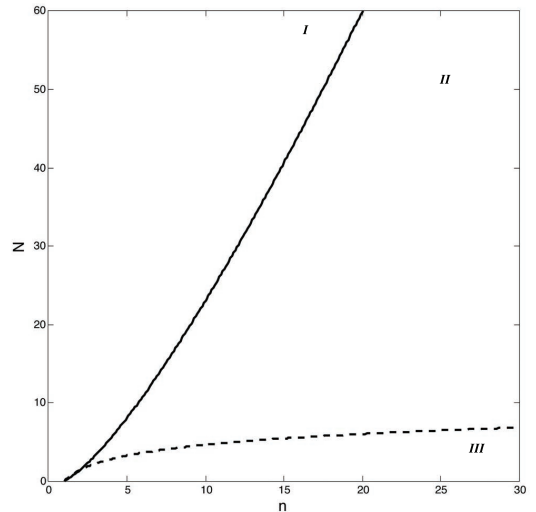


Fig. 1. An illustration of the regions in which PhaseLift and CPR are capable of recovering the ground truth solution up to a global phase ambiguity. While PhaseLift primarily targets problems in region I, CPRL operates primarily in region II.

3. NUMERICAL SOLUTIONS FOR NOISY DATA

In this section, we consider the case that the measurements are contaminated by data noise. In a linear model, typically random noise affects the output of the system as $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$, where $\mathbf{e} \in \mathbb{C}^N$ is a noise term with bounded ℓ_2 -norm: $\|\mathbf{e}\|_2 \leq \epsilon$. However, in phase retrieval, we follow closely a more special noise model used in Candès et al. [2011b]:

$$b_i = |\langle \mathbf{x}, \mathbf{a}_i \rangle|^2 + e_i. \quad (12)$$

This nonstandard model avoids the need to calculate the squared magnitude output $|\mathbf{y}|^2$ with the added noise term. More importantly, in practical phase retrieval applications, measurement noise is introduced when the squared magnitudes or intensities of the linear system are measured, not on \mathbf{y} itself (Candès et al. [2011b]).

Accordingly, we denote a linear function B of X

$$B : X \in \mathbb{C}^{n \times n} \mapsto \{\text{Tr}(\Phi_i X)\}_{1 \leq i \leq N} \in \mathbb{R}^N \quad (13)$$

that measures the noise-free squared output. Then the approximate CPR problem with bounded ℓ_2 error model (12) can be solved by the following SDP program:

$$\begin{aligned} \min \quad & \text{Tr}(X) + \lambda \|X\|_1 \\ \text{subj. to} \quad & \|B(X) - \mathbf{b}\|_2 \leq \epsilon, \\ & X \succeq 0. \end{aligned} \quad (14)$$

The estimate of \mathbf{x} , just as in noise free case, can finally be found by computing the rank-1 decomposition of X via singular value decomposition. We refer to the method as approximate CPRL. Due to the machine rounding error, in general a nonzero ϵ should be always assumed in the objective (14) and its termination condition during the optimization.

We should further discuss several numerical issues in the implementation of the SDP program. The constrained CPR problem (14) can be rewritten as an unconstrained objective function:

$$\min_{X \succeq 0} \text{Tr}(X) + \lambda \|X\|_1 + \frac{\mu}{2} \|B(X) - \mathbf{b}\|_2^2, \quad (15)$$

where $\lambda > 0$ and $\mu > 0$ are two penalty parameters.

In (15), due to the lifting process, the rank-1 condition of X is approximated by its trace function $\text{Tr}(X)$. In Candès et al. [2011b], the authors considered phase retrieval of generic (dense) signal \mathbf{x} . They proved that if the number of measurements obeys $N \geq cn \log n$ for a sufficiently large constant c , with high probability, minimizing (15) without the sparsity constraint (*i.e.*, $\lambda = 0$) recovers a unique rank-1 solution obeying $X^* = \mathbf{x}\mathbf{x}^H$.

In Section 4, we will show that using either random Fourier dictionaries or more general random projections, in practice, one needs much fewer measurements to exactly recover sparse signals if the measurements are noise-free. Nevertheless, in the presence of noise, the recovered lifted matrix X may not be exactly rank-1. In this case, one can simply use its rank-1 approximation corresponding to the largest singular value of X .

We also note that in (15), there are two main parameters λ and μ that can be defined by the user. Typically μ is chosen depending on the level of noise that affects the measurements \mathbf{b} . For λ associated with the sparsity penalty $\|X\|_1$, one can adopt a *warm start* strategy to determine its value iteratively. The strategy has been widely used in other sparse optimization, such as in ℓ_1 -min [Yang et al., 2010]. More specifically, the objective is solved iteratively with respect to a sequence of monotonically decreasing $\lambda \rightarrow 0$, and each iteration is initialized using the optimization results from the previous iteration. When λ is large, the sparsity constraint outweighs the trace constraint and the estimation error constraint, and vice versa.

Example 3.1. (Compressive Phase Retrieval). In this example, we illustrate a simple CPR example, where a 2-sparse complex signal $\mathbf{x}_0 \in \mathbb{C}^{64}$ is first transformed by the Fourier transform $F \in \mathbb{C}^{64 \times 64}$ followed by random projections $R \in \mathbb{C}^{32 \times 64}$:

$$\mathbf{b} = |RF\mathbf{x}_0|^2. \quad (16)$$

Given \mathbf{b} , F , and R , we first apply the PhaseLift algorithm [Candès et al., 2011b] with $A = RF$ to the 32 squared observations \mathbf{b} . The recovered dense signal is shown in Figure 2. PhaseLift fails to identify the 2-sparse signal.

Next, we apply CPRL (14), and the recovered sparse signal is also shown in Figure 2. CPRL correctly identifies the two nonzero elements in \mathbf{x} .

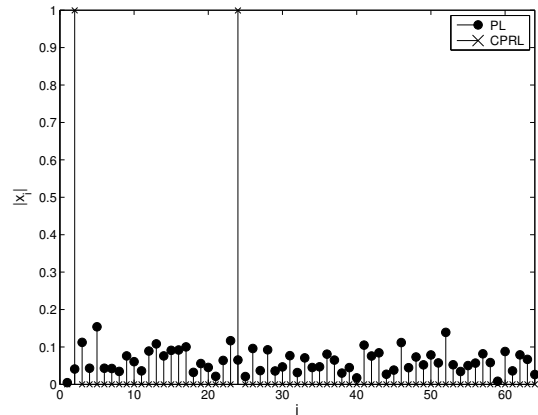


Fig. 2. The magnitude of the estimated signal provided by CPRL and PhaseLift (PL). CPRL correctly identifies elements 2 and 24 to be nonzero while PhaseLift provides a dense estimate. It is also verified that the estimate from CPRL, after a global phase shift, is approximately equal the true \mathbf{x}_0 .

4. EXPERIMENT

This section gives a number of examples. Code for the numerical illustrations can be downloaded from <http://www.rti.isy.liu.se/~ohlsson/code.html>.

4.1 Simulation

First, we repeat the simulation given in Example 3.1 for $k = 1, \dots, 5$. For each k , $n = 64$ is fixed, and we increase the measurement dimension N until CPRL recovered the true sparse support in at least 95 out of 100 trials, *i.e.*, 95% success rate. New data (\mathbf{x} , \mathbf{b} , and R) are generated in each trial. The curve of 95% success rate is shown in Figure 3.

With the same simulation setup, we compare the accuracy of CPRL with the PhaseLift approach and the CS approach in Figure 3. First, note that CS is not applicable to phase retrieval problems in practice, since it assumes the phase of the observation is also given. Nevertheless, the simulation shows CPRL via the SDP solution only requires a slightly higher sampling rate to achieve the same success rate as CS, even when the phase of the output is missing. Second, similar to the discussion in Example 3.1, without enforcing the sparsity constraint in (11), PhaseLift would fail to recover correct sparse signals in the low sampling rate regime.

It is also interesting to see the performance as n and N vary and k held fixed. We therefore use the same setup as in Figure 3 but now fixed $k = 2$ and for $n = 10, \dots, 60$, gradually increased N until CPRL recovered

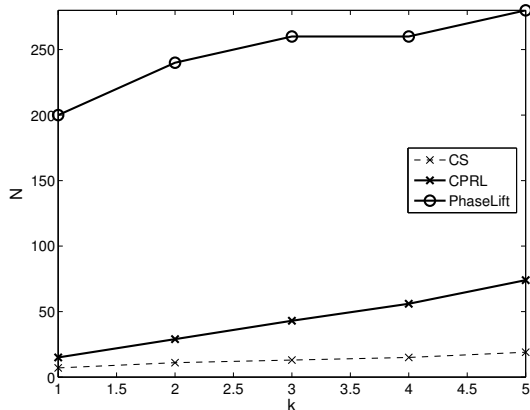


Fig. 3. The curves of 95% success rate for CPRL, PhaseLift, and CS. Note that the CS simulation is given the complete output \mathbf{y} instead of its squared magnitudes.

the true sparsity pattern with 95% success rate. The same procedure is repeated to evaluate PhaseLift and CS. The results are shown in Figure 4.

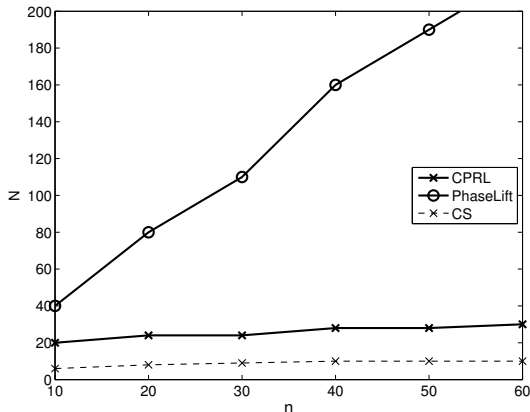


Fig. 4. The curves of 95% success rate for CPRL, PhaseLift, and CS. Note that the CS simulation is given the complete output \mathbf{y} instead of its squared magnitudes.

Compared to Figure 3, we can see that the degradation from CS to CPRL when the phase information is omitted is largely affected by the sparsity of the signal. More specifically, when the sparsity k is fixed, even when the dimension n of the signal increases dramatically, the number of squared observations to achieve accurate recovery does not increase significantly for both CS and CPRL.

4.2 CPRL Applied to Audio Signals

In this section, we further demonstrate the performance of CPRL using signals from a real-world audio recording. The timbre of a particular note on an instrument is determined by the fundamental frequency, and several overtones. In a Fourier basis, such a signal is sparse, being the summation of a few sine waves. Using the recording of a single note on an instrument will give us a naturally sparse signal, as opposed to synthesized sparse signals in the previous sections. Also, this experiment will let us analyze how robust our algorithm is in practical situations,

where effects like room ambience might color our otherwise exactly sparse signal with noise.

Our recording $\mathbf{z} \in \mathbb{R}^s$ is a real signal, which is assumed to be sparse in a Fourier basis. That is, for some sparse $\mathbf{x} \in \mathbb{C}^n$, we have $\mathbf{z} = F_{inv}\mathbf{x}$, where $F_{inv} \in \mathbb{C}^{s \times n}$ is a matrix representing a transform from Fourier coefficients into the time domain. Then, we have a randomly generated mixing matrix with normalized rows, $R \in \mathbb{R}^{N \times s}$, with which our measurements are sampled in the time domain:

$$\mathbf{y} = R\mathbf{z} = RF_{inv}\mathbf{x}. \quad (17)$$

Finally, we are only given the magnitudes of our measurements, such that $\mathbf{b} = |\mathbf{y}|^2 = |R\mathbf{z}|^2$.

For our experiment, we choose a signal with $s = 32$ samples, $N = 30$ measurements, and it is represented with $n = 2s$ (overcomplete) Fourier coefficients. Also, to generate F_{inv} , the $\mathbb{C}^{n \times n}$ matrix representing the Fourier transform is generated, and s rows from this matrix are randomly chosen.

The experiment uses part of an audio file recording the sound of a tenor saxophone. The signal is cropped so that the signal only consists of a single sustained note, without silence. Using CPRL to recover the original audio signal given \mathbf{b} , R , and F_{inv} , the algorithm gives us a sparse estimate \mathbf{x} , which allows us to calculate $\mathbf{z}_{est} = F_{inv}\mathbf{x}$. We observe that all the elements of \mathbf{z}_{est} have phases that are π apart, allowing for one global rotation to make \mathbf{z}_{est} purely real. This matches our previous statements that CPRL will allow us to retrieve the signal up to a global phase.

We also find that the algorithm is able to achieve results that capture the trend of the signal using less than s measurements. In order to fully exploit the benefits of CPRL that allow us to achieve more precise estimates with smaller errors using fewer measurements relative to s , the problem should be formulated in a much higher ambient dimension. However, using the CVX Matlab toolbox by Grant and Boyd [2010], we already ran into computational and memory limitations with the current implementation of the CPRL algorithm. These results highlight the need for a more efficient numerical implementation of CPRL.

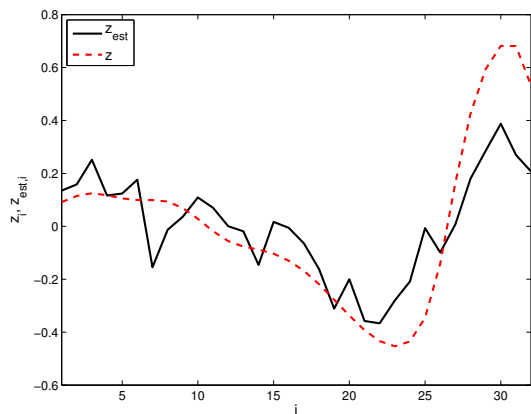


Fig. 5. The retrieved signal \mathbf{z}_{est} using CPRL versus the original audio signal \mathbf{z} .

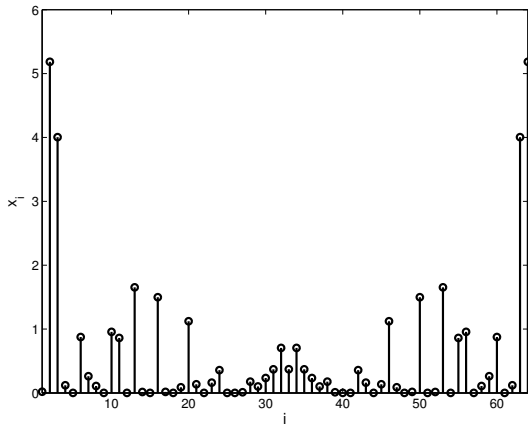


Fig. 6. The magnitude of $\hat{\mathbf{x}}$ retrieved using CPRL. The audio signal \mathbf{z}_{est} is obtained by $\mathbf{z}_{est} = F_{inv}\hat{\mathbf{x}}$.

5. CONCLUSION AND DISCUSSION

A novel method for the compressive phase retrieval problem has been presented. The method takes the form of an SDP problem and provides the means to use compressive sensing in applications where only squared magnitude measurements are available. The convex formulation gives it an edge over previous presented approaches and numerical illustrations show state of the art performance.

One of the future directions is improving the speed of the standard SDP solver, i.e., interior-point methods, currently used for the CPRL algorithm. Some preliminary results along with a more extensive study of the performance bounds of CPRL are available in Ohlsson et al. [2011].

REFERENCES

- R. Balan, P. Casazza, and D. Edidin. On signal reconstruction without phase. *Applied and Computational Harmonic Analysis*, 20:345–356, 2006.
- A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- A. Bruckstein, D. Donoho, and M. Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM Review*, 51(1):34–81, 2009.
- E. J. Candès and M. Wakin. An introduction to compressive sampling. *Signal Processing Magazine, IEEE*, 25(2):21–30, March 2008.
- E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52:489–509, February 2006.
- E. J. Candès, Y. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. Technical Report arXiv:1109.0573, Stanford University, September 2011a.
- E. J. Candès, T. Strohmer, and V. Voroninski. PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming. Technical Report arXiv:1109.4499, Stanford University, September 2011b.
- A. Chai, M. Moscoso, and G. Papanicolaou. Array imaging using intensity-only measurements. Technical report, Stanford University, 2010.
- S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.
- J. Dainty and J. Fienup. Phase retrieval and image reconstruction for astronomy. In *Image Recovery: Theory and Applications*. Academic Press, New York, 1987.
- A. d’Aspremont, L. El Ghaoui, M. Jordan, and G. Lanckriet. A direct formulation for Sparse PCA using semidefinite programming. *SIAM Review*, 49(3):434–448, 2007.
- D. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, April 2006.
- J. Fienup. Phase retrieval algorithms: a comparison. *Applied Optics*, 21(15):2758–2769, 1982.
- J. Fienup. Reconstruction of a complex-valued object from the modulus of its Fourier transform using a support constraint. *Journal of Optical Society of America A*, 4(1):118–123, 1987.
- R. Gerchberg and W. Saxton. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–246, 1972.
- R. Gonsalves. Phase retrieval from modulus data. *Journal of Optical Society of America*, 66(9):961–964, 1976.
- M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx>, August 2010.
- A. Hoerl and R. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- 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*, 63(2):126–134, 1973.
- I. Loris. On the performance of algorithms for the minimization of ℓ_1 -penalized functionals. *Inverse Problems*, 25:1–16, 2009.
- S. Marchesini. Phase retrieval and saddle-point optimization. *Journal of the Optical Society of America A*, 24(10):3289–3296, 2007.
- R. Millane. Phase retrieval in crystallography and optics. *Journal of the Optical Society of America A*, 7:394–411, 1990.
- M. Moravec, J. Romberg, and R. Baraniuk. Compressive phase retrieval. In *SPIE International Symposium on Optical Science and Technology*, 2007.
- H. Ohlsson, A. Y. Yang, R. Dong, and S. Sastry. Compressive Phase Retrieval From Squared Output Measurements Via Semidefinite Programming. Technical Report arXiv:1111.6323, University of California, Berkeley, November 2011.
- R. Tibsharani. Regression shrinkage and selection via the lasso. *Journal of Royal Statistical Society B (Methodological)*, 58(1):267–288, 1996.
- J. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242, October 2004.
- A. Walther. The question of phase retrieval in optics. *Optica Acta*, 10:41–49, 1963.
- A. Yang, A. Ganesh, Y. Ma, and S. Sastry. Fast ℓ_1 -minimization algorithms and an application in robust face recognition: A review. In *ICIP*, 2010.