GESPAR:Efficient Sparse Phase Retrieval with Application to Optics

Yoav Shechtman
Physics Department
Technion, Israel Institute of Technology
joe@tx.technion.ac.il

Amir Beck
Department of Industrial Engineering
Technion, Israel Institute of Technology
becka@ie.technion.ac.il

Yonina C. Eldar
Department of Electrical Engineering
Technion, Israel Institute of Technology
yonina@ee.technion.ac.il

***Abstract*—The problem of phase retrieval, namely, recovery of a signal from the magnitude of its Fourier transform is ill-posed since the Fourier phase information is lost. Therefore, prior information on the signal is needed in order to recover it. In this work we consider the case in which the prior information on the signal is that it is sparse, i.e., it consists of a small number of nonzero elements. We propose GESPAR: A fast local search method for recovering a sparse signal from measurements of its Fourier transform magnitude. Our algorithm does not require matrix lifting, unlike previous approaches, and therefore is potentially suitable for large scale problems such as images. Simulation results indicate that the proposed algorithm is fast and more accurate than existing techniques. We demonstrate applications in optics where GESPAR is generalized and used for finding sparse solutions to sets of quadratic measurements.**

I. Introduction

Recovery of a signal from the magnitude of its Fourier transform, also known as phase retrieval, is of great interest in applications such as optical imaging [1], crystallography [2], and more [3]. Due to the loss of Fourier phase information, the problem (in 1D) is generally ill-posed. A common approach to overcome this ill-posedeness is to exploit prior information on the signal. A variety of methods have been developed that use such prior information, which may be the signal's support, non-negativity, or the signal's magnitude [4], [5]. A popular class of algorithms is based on the use of alternate projections between the different constraints. In order to increase the probability of correct recovery, these methods require the prior information to be very precise, for example, exact/or "almost" exact knowledge of the support set. Since the projections are generally not onto convex sets, convergence to a correct recovery is not guaranteed [6]. A more recent approach is to use matrix-lifting of the problem which allows to recast phase retrieval as a semi-definite programming (SDP) problem [7]. The algorithm developed in [7] does not require prior information about the signal but instead uses multiple signal measurements (e.g., using different illumination settings, in an optical setup).

In order to obtain more robust recovery without requiring multiple measurements, we develop a method that exploits signal sparsity. Existing approaches aimed at recovering sparse signals from their Fourier magnitude belong to two main categories: SDP-based techniques [8],[9],[10] and algorithms that use alternate projections (Fienup-type methods) [11]. Phase retrieval of sparse signals can be viewed as a special case of the more general quadratic compressed sensing (QCS) problem considered in [8]. Specifically, QCS treats recovery of sparse vectors from quadratic measurements of the form $y_i = \mathbf{x}^T\mathbf{A}_i\mathbf{x},\ \ i = 1, \ldots, N$, where $\mathbf{x}$ is the unknown sparse vector to be recovered, y_i are the measurements, and $\mathbf{A}_i$ are known matrices. In (discrete) phase retrieval, $\mathbf{A} = \mathbf{F}_i^T\mathbf{F}_i$ where $\mathbf{F}_i$ is the ith row of the discrete Fourier transform (DFT) matrix.

A general approach to QCS was developed in [8] , in the context of partially incoherent imaging, based on matrix lifting. More specifically, the quadratic constraints where lifted to a higher dimension by defining a matrix variable $\mathbf{X} = \mathbf{x}\mathbf{x}^T$. The problem was then recast as an SDP involving minimization of the rank of the lifted matrix subject to the recovery constraints as well as row sparsity constraints on $\mathbf{X}$. An iterative thresholding algorithm based on a sequence of SDPs was then proposed to recover a sparse solution. Similar SDP-based ideas were recently used in the context of phase retrieval [9],[10]. However, due to the increase in dimension created by the matrix lifting procedure, the SDP approach is not suitable for large-scale problems.

Another approach for phase retrieval of sparse signals is adding a sparsity constraint to the well-known iterative error reduction algorithm of Fienup [11]. In general, Fienup-type approaches are known to suffer from convergence issues and often do not lead to correct recovery especially in 1D problems; simulation results show that even with the additional information that the input is sparse, convergence is still problematic and the algorithm often recovers erroneous solutions.

In this paper we propose an efficient method for phase retrieval which also leads to good recovery performance. Our algorithm is based on a fast 2-opt local search method (see [12] for an excellent introduction to such techniques) applied to a sparsity constrained non-linear optimization formulation of the problem. We refer to our algorithm as GESPAR: GrEedy Sparse PhAse Retrieval. Sparsity constrained nonlinear optimization problems have been considered recently in [13]; the method derived in this paper is motivated – although different in many aspects – by the local search-type techniques of [13]. We demonstrate through numerical simulations that the proposed algorithm is both efficient and more accurate than current techniques, and we present an example application in optical imaging where a modified version of GESPAR is used.

II. Problem Formulation

We are given a vector of measurements $\mathbf{y} \in \mathbb{R}^N$, that corresponds to the magnitude of an N point discrete Fourier transform of a vector $\mathbf{x} \in \mathbb{R}^N$, i.e.:

$$y_l = \left| \sum_{m=1}^{n} x_m e^{-\frac{2\pi j(m-1)(l-1)}{N}} \right|, \;\; l = 1, \ldots, N, \qquad (1)$$

where $\mathbf{x}$ was constructed by zeros padding of a vector $\bar{\mathbf{x}} \in \mathbb{R}^n$ $(n < N)$ with elements $x_i, \;\; i = 1, 2, \ldots, n$. In the simulations section we considered the setting $N = 2n$ which corresponds to oversampling the DFT of $\bar{\mathbf{x}}$ by a factor of 2. In any case, we will assume that $N \geq 2n - 1$. This allows to determine the correlation sequence of $\mathbf{x}$ from the given measurements, as we elaborate on more below. Denoting by $\mathbf{F} \in \mathbb{C}^{N \times N}$ the DFT matrix with elements $e^{-\frac{2\pi j(m-1)(l-1)}{N}}$, we can express $\mathbf{y}$ as $\mathbf{y} = |\mathbf{F}\mathbf{x}|$, where $|\cdot|$ denotes the element-wise absolute value. The vector $\mathbf{x}$ is known to be s-sparse on its support, i.e., it contains at most s nonzero elements in the first n elements. Our goal is to recover $\mathbf{x}$ given the measurements $\mathbf{y}$ and the sparsity level s.

The mathematical formulation of the problem that we consider consists of minimizing the sum of squared errors subject to the sparsity constraint:

$$\begin{array}{ll} \min_{\mathbf{x}} & \sum_{i=1}^{N} (|\mathbf{F}_i \mathbf{x}|^2 - y_i^2)^2 \\ \text{s.t.} & \|\mathbf{x}\|_0 \leq s, \\ & supp(\mathbf{x}) \subseteq \{1, 2, \ldots, n\}, \\ & \mathbf{x} \in \mathbb{R}^N, \end{array} \qquad (2)$$

where $\mathbf{F}_i$ is the ith row of the matrix $\mathbf{F}$, $\|\cdot\|_0$ stands for the zero-"norm", that is, the number of nonzero elements. Note that the unknown vector $\mathbf{x}$ can only be found up to trivial degeneracies that are the result of the loss of Fourier phase information: circular shift, global phase, and signal "mirroring".

To aid in solving the phase retrieval problem we will rely on the fact that the correlation sequence of $\bar{\mathbf{x}}$ (the first n components of $\mathbf{x}$) can be determined from $\mathbf{y}$. Specifically, let $g_m = \sum_{i=1}^{n} x_i x_{i+m}, m = -(n-1), \ldots, n-1$ denote the correlation sequence. Note that $\{g_m\}$ is a sequence of length $2n - 1$. Since the DFT length N satisfies $N \geq 2n - 1$, we can obtain $\{g_m\}$ by the inverse DFT of the squared Fourier magnitude $\mathbf{y}$. Throughout the paper, we assume that no support cancellations occur in $\{g_m\}$, namely, if $x_i \neq 0$ and $x_j \neq 0$ for some i, j, then $g_{|i-j|} \neq 0$. When the values of $\mathbf{x}$ are random, this is true with probability 1. This fact is used in the proposed algorithm in order to obtain information on the support of $\mathbf{x}$.

The information on the support is used to derive two sets, J_1 and J_2 from the correlation sequence $\{g_m\}$ in the following manner. Let J_1 be the set of indices known in advance to be in the support, from the autocorrelation sequence. In the noiseless setting which we consider, J_1 comprises two indices:

$$J_1 = \{1, i_{\max}\}.$$

Due to the existing degree of freedom relating to shift-invariance of $\mathbf{x}$, the index 1 can be assumed to be in the support, thereby removing this degree of freedom; as a consequence, the index corresponding to the last nonzero element in the autocorrelation sequence is also in the support, i.e.

$$i_{max} = 1 + \underset{i}{\operatorname{argmax}}\{i : g_i \neq 0\}.$$

We denote by J_2 the set of indices that are candidates for being in the support, meaning the indices that are *not* known in advance to be in the off-support (the complement of the support). In other words, J_2 contains the set of all indices $k \in \{1, 2, \ldots, n\}$ such that $g_{k-1} \neq 0$. Obviously, since we assume that $x_k = 0$ for $k > n$, we have $J_2 \subseteq \{1, 2, \ldots, n\}$. Defining $\mathbf{A}_i = \Re(\mathbf{F}_i)^T \Re(\mathbf{F}_i) + \Im(\mathbf{F}_i)^T \Im(\mathbf{F}_i) \in \mathbb{R}^{N \times N}$ and $c_i = y_i^2$ for $i = 1, 2, \ldots, N$, problem (2) along with the support information can be written as

$$\begin{array}{ll} \min_{\mathbf{x}} & f(\mathbf{x}) \equiv \sum_{i=1}^{N} (\mathbf{x}^T \mathbf{A}_i \mathbf{x} - c_i)^2 \\ \text{s.t.} & \|\mathbf{x}\|_0 \leq s, \\ & J_1 \subseteq supp(\mathbf{x}) \subseteq J_2, \\ & \mathbf{x} \in \mathbb{R}^N, \end{array} \qquad (3)$$

which will be the formulation to be studied.

In the next section, we propose an iterative local-search based algorithm for solving (3). We note that although in the context of phase retrieval the parameters $\mathbf{A}_i, J_1, J_2$ have special properties (e.g., $\mathbf{A}_i$ is positive semidefinite of at most rank 2, $|J_1| = 2$), we will not use these properties in the proposed method. Therefore, our approach is capable of handling general instances of (3) with the sole assumption that $\mathbf{A}_i$ is symmetric for any $i = 1, 2, \ldots, N$.

III. GrEedy Sparse PhAse Retrieval (GESPAR) Algorithm

In this section GESPAR is summarized. A more detailed description can be found in [14].

A. The Damped Gauss-Newton Method

Before describing the algorithm, we begin by presenting the damped Gauss-Newton (DGN) method [15],[16] that is in fact the core step of our approach. The DGN method is invoked in order to solve the problem of minimizing the objective function f over a *given* support $S \subseteq \{1, 2, \ldots, n\}$ $(|S| = s)$:

$$\min\{f(\mathbf{U}_S \mathbf{z}) : \mathbf{z} \in \mathbb{R}^s\}, \qquad (4)$$

where $\mathbf{U}_S \in \mathbb{R}^{n \times s}$ is the matrix consisting of the columns of the identity matrix $\mathbf{I}_N$ corresponding to the index set S. With this notation, (4) can be explicitly written as

$$\min \left\{ g(\mathbf{z}) \equiv \sum_{i=1}^{N} (\mathbf{z}^T \mathbf{U}_S^T \mathbf{A}_i \mathbf{U}_S \mathbf{z} - c_i)^2 : \mathbf{z} \in \mathbb{R}^s \right\}. \qquad (5)$$

Problem (5) is a nonlinear least-squares problem. A natural approach for tackling it is via the DGN iterations. This algorithm begins with an arbitrary vector $\mathbf{z}_0$. We choose it to be an uncorrelated random Gaussian vector with zero mean and unit variance. At each iteration, all the terms inside the squares in $g(\mathbf{z})$ are linearized around the previous guess. The linearized term is then minimized to determine the next approximation of the solution. Specifically, at each step we pick $\mathbf{y}_k$ to be the solution of

$$\underset{y}{\operatorname{argmin}} \left\{ \sum_{i=1}^{N} (\mathbf{z}_{k-1}^T \mathbf{B}_i \mathbf{z}_{k-1} - c_i + 2(\mathbf{B}_i \mathbf{z}_{k-1})^T (\mathbf{y} - \mathbf{z}_{k-1}))^2 \right\},$$

where $\mathbf{B}_i = \mathbf{U}_S^T \mathbf{A}_i \mathbf{U}_S$. This can be written as the linear least squares problem

$$\mathbf{y}_k = \arg\min \|\mathbf{M}\mathbf{y} - \mathbf{b}\|_2^2 \tag{6}$$

with the ith row of $\mathbf{M}$ being $\mathbf{M}_i = 2(\mathbf{B}_i \mathbf{z}_{k-1})^T$, and with $b_i = c_i + \mathbf{z}_{k-1}^T \mathbf{B}_i \mathbf{z}_{k-1}$ for $i = 1, 2, \ldots, N$. The solution $\mathbf{y}_k$ can therefore be calculated explicitly by the pseudo-inverse of $\mathbf{M}$, i.e. $\mathbf{y}_k = (\mathbf{M}^T \mathbf{M})^{-1} \mathbf{M}^T \mathbf{b}$. We then define a direction vector as $\mathbf{d}_k = \mathbf{y}_k - \mathbf{z}_{k-1}$. This direction is used to update the solution with an appropriate stepsize designed to guarantee the convergence of the method to a stationary point of $g(\mathbf{z})$. The stepsize is chosen via a simple backtracking procedure.

B. The 2-opt Local Search Method

The GESPAR method consists of repeatedly invoking a local-search method on an initial random support set. In this section we describe the local search procedure. At the beginning, the support is chosen to be a set of s random indices chosen to satisfy the support constraints $J_1 \subseteq S \subseteq J_2$. Then, at each iteration a swap between a support and an off-support index is performed such that the resulting solution via the DGN method improves the objective function. Since at each iteration only two elements are changed (one in the support and one in the off-support), this is a so-called "2-opt" method (see [12]). The swaps are always chosen to be between support indices corresponding to components in the current iterate with small absolute value and off-support indices corresponding to large absolute value of ∇f. This process continues as long as the objective function decreases and stops when no improvement can be made.

C. The GESPAR Algorithm

The 2-opt method can have the tendency to get stuck at local optima points. Therefore, our final algorithm, which we call GESPAR, is a restarted version of 2-opt. The 2-opt method is repeatedly invoked with different initial random support sets until the resulting objective function value is smaller than a certain threshold (success) or the number of maximum allowed total number of swaps was passed (failure). A detailed description of the method is given in Algorithm 1. One element of our specific implementation that is not described in Algorithm 1 is the incorporation of random weights added to the objective function, giving randomly different weights to the different measurements.

Algorithm 1 GESPAR

Input: $(\mathbf{A}_i, c_i, \tau, \text{ITER})$.
$\mathbf{A}_i \in \mathbb{R}^{N \times n}, i = 1, 2, \ldots, N$ - symmetric matrices.
$c_i \in \mathbb{R}, i = 1, 2, \ldots, N$.
τ - threshold parameter.
ITER - Maximum allowed total number of swaps.

Output: $\mathbf{x}$ - an optimal (or suboptimal) solution of (3).

Initialization. Set $C = 0, k = 0$.

- **Repeat**
  Invoke the 2-opt method with input $(\mathbf{A}_i, c_i, 4, 8)$ and obtain an output $\mathbf{x}$ and T. Set $\mathbf{x}_k = \mathbf{x}, C = C + T$ and advance k: $k \leftarrow k + 1$.
  Until $f(\mathbf{x}) < \tau$ or $C >$ ITER.
- The output is $\mathbf{x}_\ell$ where $\ell = \text{argmin}_{m=0,1,\ldots,k-1} f(\mathbf{x}_m)$.

IV. Numerical Simulation

In order to demonstrate the performance of GESPAR, we conducted a numerical simulation. The algorithm is evaluated both in terms of signal-recovery accuracy and in terms of computational efficiency.

A. Simulation details

We choose $\bar{\mathbf{x}}$ as a random vector of length n. The vector contains uniformly distributed values in s randomly chosen elements. The N point DFT of the signal is calculated, and its magnitude is taken as $\mathbf{y}$, the vector of measurements. The $2n - 1$ point correlation is also calculated. In order to recover the unknown vector $\mathbf{x}$, the GESPAR algorithm is used with $\tau = 10^{-4}$ and $T = 20000$, as well as two other algorithms for comparison purposes: An SDP based algorithm (Algorithm 2, [9].), and an iterative Fienup algorithm with a sparsity constraint [11]. In our simulation $n = 64$ and $N = 128$.

B. Simulation Results

Signal recovery results of the numerical simulation are shown in Fig. 1, where the probability for successful recovery is plotted for different sparsity levels. Successful recovery probability is defined as the ratio of correctly recovered signals $\mathbf{x}$ out of 100 signal-simulations. In each simulation both the support and the signal values are randomly selected. The three algorithms (GESPAR, SDP and Sparse-Fienup) are compared. The results clearly show that GESPAR outperforms the other algorithms in terms of probability of successful recovery - over 90% successful recovery up to $s = 15$, vs. $s = 8$ and $s = 5$ in the other two algorithms.

The average runtime performance of the three algorithms was also compared for several sparsity levels ($s = 3, 5, 8$), and the results are shown in table I. GESPAR is shown to perform much faster than the SDP based method, and comparable in time to the Sparse-Fienup method, while outperforming both in terms of signal recovery.

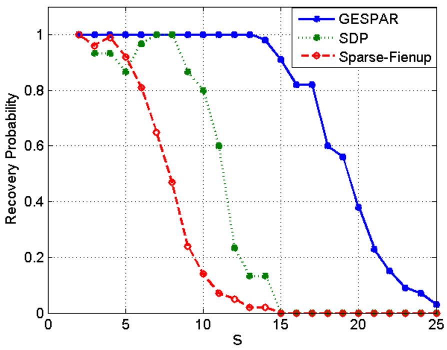

Fig. 1. Recovery probability vs. sparsity (s)

TABLE I
RUNTIME COMPARISON

	SDP	Sparse-Fienup	GESPAR
$s = 3$	1.32 sec	0.09 sec	0.12 sec
$s = 5$	1.78 sec	0.12 sec	0.12 sec
$s = 8$	3.85 sec	0.50 sec	0.23 sec

V. Applications in Optics

As an example of one of the recent applications of GESPAR in optical problems, where it is modified to handle more general quadratic problems, we present Coherent Diffractive Imaging (CDI) for sparsely varying objects. CDI [17] is an imaging method used usually in the x-ray domain, where a small object is illuminated by a coherent plane wave, and the far-field diffraction intensity pattern is measured. The measured intensity corresponds to the 2D Fourier transform of the object. Discretization of the problem followed by appropriate scaling of coordinates yields: $y_i = \mathbf{x}^T \mathbf{A}_i \mathbf{x}, \ \ i = 1, \ldots, N$, where y_i are the far-field intensity measurements, $\mathbf{x}$ is the object to be recovered, and as before - $\mathbf{A}_i = \mathbf{F}_i^T \mathbf{F}_i$. We shall now focus on an example where a dynamic scene is being imaged - e.g. a moving object - so that sequential intensity patterns are being captured at a certain frame rate. If the difference in the object between the consecutive frames $\Delta_k = \mathbf{x}_k - \mathbf{x}_{k-1}$ is sparse (even if the object itself is not) - then recovering the frame difference becomes the problem of finding a sparse solution Δ_k to $y_k^i = (\mathbf{x}_{k-1} + \Delta_k)^T \mathbf{A}^i (\mathbf{x}_{k-1} + \Delta_k)$. Given The result of the previous frame $\mathbf{x}_{k-1}$,this is a quadratic problem in Δ_k, and a modified version of GESPAR is used to solve it. An example recovery is shown in Figure 2- where a comparison to standard frame by frame Fienup HIO [4] recovery without using sparsity is made. In this example there is added noise (SNR=30) and the first frame is assumed to be known (e.g. $\mathbf{y}_0$ is measured with a sufficient number of measurements).

VI. Conclusion

We proposed and demonstrated GESPAR - a fast algorithm to recover a sparse vector from its Fourier magnitude. We showed via simulations that GESPAR outperforms alternative approaches suggested for this problem. The algorithm does not require matrix-lifting, and therefore is potentially suitable for large scale problems such as 2D images, and we demonstrate its application for a more general quadratic imaging problem.

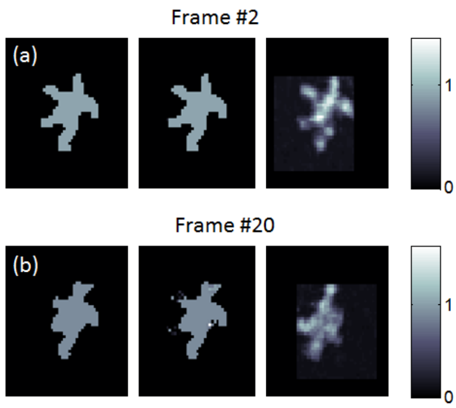

Fig. 2. Sparsely varying CDI example - True object (Left) is being recovered from noisy Fourier magnitude (SNR=30), using sparsity of frame differences (GESPAR - center) and without (Fienup HIO algorithm - right).

References

[1] A. Walther. "The question of phase retrieval in optics". Opt. Acta, 10:41–49, 1963.
[2] R.W. Harrison, "Phase problem in crystallography". J. Opt. Soc. Am. A, 10(5):1045–1055, 1993.
[3] N. Hurt. "Phase retrieval and zero crossings," Kluwer Academic Publishers, Norwell, MA, 1989.
[4] J.R. Fienup, "Phase retrieval algorithms: a comparison," Applied Optics 21, 2758-2769, 1982.
[5] R.W. Gerchberg and W.O. Saxton, "Phase retrieval by iterated projections," Optik 35, 237, 1972.
[6] 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. Am. A, 19(7):1334–1345, 2002.
[7] E. J. Candes, Y. C. Eldar, T. Strohmer and V. Voroninski, "Phase retrieval via matrix completion," arXiv:1109.0573, Sep. 2011.
[8] Y. Shechtman, Y.C. Eldar, A. Szameit, and M. Segev, "Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing," Optics Express 19, 16, 14807-24822, 2011.
[9] K. Jaganathan, S.Oymak, and B. Hassibi, "Recovery of sparse 1-D signals from the magnitudes of their Fourier transform," arXiv:1206.1405v1, June 2012.
[10] H. Ohlsson, A. Y. Yang, R. Dong, S. S. Sastry, "Compressive phase retrieval from squared output measurements via semidefinite programming," arXiv:1111.6323v3 , March 2012.
[11] S. Mukherjee and C. S. Seelamantula, "An iterative algorithm for phase retrieval with sparsity constaints: Application to frequency-domain optical-coherence tomography," ICASSP 2012
[12] C. H. Papadimitriou. *Combinatorial optimization: algorithms and complexity*. Prentice-Hall 1982, with Ken Steiglitz; second edition by Dover, 1998.
[13] A. Beck, Y. C. Eldar, "Sparsity constrained nonlinear optimization: Optimality conditions and algorithms," arXiv:1203.4580v1,
[14] Y. Shechtman et al., "GESPAR: Efficient phase retrieval of sparse signals", arXiv:1301.1018 (2013).
[15] D. P. Bertsekas. *Nonlinear Programming*. Belmont MA: Athena Scientific, second edition, 1999. March 2012.
[16] A. Björck, *Numerical Methods for Least-Squares Problems*, Philadelphia, PA: SIAM, 1996.
[17] J. Miao et al., "Extending the methodology of X-ray crystallography to allow imaging of micrometer-sized non-crystalline specimens," Nature, vol. 400, pp. 342–344, (1999).