Published in final edited form as:

IEEE Trans Comput Imaging. 2015 December 1; 1(4): 247–258. doi:10.1109/TCI.2015.2498402.

Undersampled Phase Retrieval with Outliers

Daniel S. Weller [Member, IEEE],

Charles L. Brown Department of Electrical and Computer Engineering, University of Virginia,
Charlottesville, VA 22904 USA

Ayelet Pnueli,

Electrical Engineering Department, Technion, Israel Institute of Technology, Haifa 32000, Israel

Gilad Divon,

Electrical Engineering Department, Technion, Israel Institute of Technology, Haifa 32000, Israel

Ori Radzyner,

Electrical Engineering Department, Technion, Israel Institute of Technology, Haifa 32000, Israel

Yonina C. Eldar [Fellow, IEEE], and

Electrical Engineering Department, Technion, Israel Institute of Technology, Haifa 32000, Israel

Jeffrey A. Fessler [Fellow, IEEE]

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor,
MI 48109 USA

Daniel S. Weller: dweller@virginia.edu; Ayelet Pnueli: ayelet.pnueli@gmail.com; Gilad Divon: giladd44@gmail.com; Ori Radzyner: radzy@campus.technion.ac.il; Yonina C. Eldar: yonina@ee.technion.ac.il; Jeffrey A. Fessler: fessler@umich.edu

Abstract

This paper proposes a general framework for reconstructing sparse images from undersampled (squared)-magnitude data corrupted with outliers and noise. This phase retrieval method uses a layered approach, combining repeated minimization of a convex majorizer (surrogate for a nonconvex objective function), and iterative optimization of that majorizer using a preconditioned variant of the alternating direction method of multipliers (ADMM). Since phase retrieval is nonconvex, this implementation uses multiple initial majorization vectors. The introduction of a robust 1-norm data fit term that is better adapted to outliers exploits the generality of this framework. The derivation also describes a normalization scheme for the regularization parameter and a known adaptive heuristic for the ADMM penalty parameter. Both 1D Monte Carlo tests and 2D image reconstruction simulations suggest the proposed framework, with the robust data fit term, reduces the reconstruction error for data corrupted with both outliers and additive noise, relative to competing algorithms having the same total computation.

Index Terms

phase retrieval; sparsity; majorize-minimize; alternating direction method of multipliers

I. Introduction

Phase retrieval [1]–[3] refers to the problem of recovering a signal or image from magnitude-only measurements of a transform of that signal. This problem appears in crystallography [4]–[7], optics [8], astronomy [9], and other areas [10]–[14].

Phase retrieval is inherently ill-posed, as many signals may share the same magnitude spectrum [15]. To address this issue, existing phase retrieval algorithms incorporate different sources of prior information. The Gerchberg-Saxton error reduction method [16] of alternating projections uses magnitude information about both an image and its Fourier spectrum. Fienup's hybrid input-output algorithm [17], [18] generalizes the image domain projection of error reduction to other constraints such as image boundary and support information [19]–[24]. More recently, the alternating projections framework [25] has been extended to sparse reconstruction [26]–[28]; examples include compressive phase retrieval [29], the message-passing method PR-GAMP [30], and the sparse Fienup method [31]. Other formulations approach phase retrieval differently. One method uses rough phase estimates [32] to dramatically improve reconstruction quality. Another uses a matrix lifting scheme [33], [34] to construct a semidefinite relaxation of the phase retrieval problem [35] that may be combined with sparsity-promoting regularization [33], [36]–[41]. Graph-based and convex optimization methods in [42] and greedy algorithms like GESPAR [43] also employ sparsity for phase retrieval.

Measurements can be very noisy at the resolution desired in many phase retrieval imaging applications. Many existing methods either ignore measurement noise or use quadratic data fit terms. The proposed method, based on [44], employs a robust 1-norm data fit term, corresponding to the negative log-likelihood of a Laplace distribution, to improve robustness to outliers. This data fit term can also be found in some matrix lifting phase retrieval methods [40], [41], at the expense of much larger memory and computational resources. Fast convergence of the proposed reconstruction can be achieved through a new optimization framework nesting two iterative components: alternating direction method of multipliers (ADMM) iterations inside each step of an outer majorize-minimize (MM) algorithm. This framework accommodates both the desired 1-norm data fit term and sparsity-promoting regularization. More specifically, majorization yields a tight convex surrogate for the original nonconvex objective. Introducing an auxiliary variable enables efficient minimization of this majorizer via a more easily separable preconditioned variant of ADMM (ordinary ADMM was used in [44]).

This paper is organized as follows. Section II presents a robust cost function for the phase retrieval problem. Section III introduces a convex majorizer for this optimization problem, and Section IV describes the use of ADMM to solve this convex subproblem. This section also introduces an optional regularization parameter normalization factor for Monte Carlo simulations and an existing adaptive heuristic for the ADMM penalty parameter [45] to greatly reduce manual tuning of these parameters. Experiments in Section V validate the parameter selection approach, compare convergence against a conventional algorithm applied to the robust phase retrieval problem, and evaluate the proposed method against existing sparsity-promoting phase retrieval methods, including a 1-norm variant of sparse

Fienup [31], the message-passing method PR-GAMP [30], and GESPAR [43]. Supplementary material includes a comparison with CPRL matrix lifting [37]; however, extreme memory requirements prevented CPRL from inclusion in the experiments with larger signals. Both 1D Monte Carlo and 2D simulations demonstrate that the proposed approach improves reconstruction quality versus all four competing methods when measurements contain both outliers and additive noise. Section VI discusses the proposed framework and algorithm and future extensions.

Code is online at <http://people.virginia.edu/~dsw8c/sw.html>. A supplement with additional experiments and derivations is available from IEEE Xplore.

II. Problem Statement

The following forward model describes the acquisition of M squared-magnitude measurements $\mathbf{y} = [y_1, \dots, y_M]^T$ from a general $M \times N$ linear transform $\mathbf{A}$ of a length- N (complex-valued) signal $\mathbf{x}$:

$$y_m = |[\mathbf{A}\mathbf{x}]_m|^2 + \nu_m, \quad m=1, \dots, M, \quad (1)$$

where $[\mathbf{A}\mathbf{x}]_m = \sum_{n=1}^N A_{mn}x_n$, and $[\nu_1, \dots, \nu_M]^T$ is a vector of white Gaussian noise added to the squared-magnitude data. In contrast to adding noise to the complex $\mathbf{A}\mathbf{x}$ before taking the magnitude, as in [17], [30], this paper uses the post-magnitude noise model found in [25], [33], [35], [37], [43]. The vector $\mathbf{x}$ may represent either a 1D signal or a higher dimensional image, columnized.

Expanding beyond the conventional model in (1), the proposed framework aims to minimize the sum of negative log-likelihood functions $\sum_{m=1}^M -\ell(y_m; |[\mathbf{A}\mathbf{x}]_m|^q)$, for $q \geq 1$. The system may measure the magnitude $|[\mathbf{A}\mathbf{x}]_m|$ ($q = 1$), its square ($q = 2$), or a more general power ($q \geq 1$). More importantly, the data fit term extends more broadly to negative log-likelihood functions of the form $f(h([\mathbf{A}\mathbf{x}]_m; y_m))$, where $f(\cdot)$ is convex and nondecreasing (on $\mathbb{R}_+$), and the function

$$h(t; y) \triangleq |y - |t|^q| \quad (2)$$

of $t \in \mathbb{C}$ is the data fit error for fixed $y \in \mathbb{R}$. For this class of log-likelihood functions, the majorizer derived in Section III is convex in $\mathbf{x}$. To account for outliers in squared-magnitude measurements, this paper explores using the negative log-likelihood of a Laplace distribution:

$$-\ell(y_m; |[\mathbf{A}\mathbf{x}]_m|^2) \propto |y_m - |[\mathbf{A}\mathbf{x}]_m|^2|. \quad (3)$$

This data fit term takes the form of a 1-norm and has a long history of providing robustness to outliers, even if the measurement noise does not follow a Laplace distribution [46].

In this work, the 1-norm $\|\mathbf{x}\|_1$ regularizes the ill-posed phase retrieval problem, promoting image sparsity. Including a synthesis transform in the sensing matrix $\mathbf{A}$ directly extends this prior to synthesis-form sparsity. The proposed phase retrieval approach seeks a minimizer $\hat{\mathbf{x}} \in \mathbb{C}^N$ of

$$\arg \min_{\mathbf{x} \in \mathbb{C}^N} \Psi(\mathbf{x}) \triangleq \sum_{m=1}^M f(h([\mathbf{A}\mathbf{x}]_m; y_m)) + \beta \|\mathbf{x}\|_1, \quad (4)$$

where $\beta > 0$ is the regularization penalty parameter, and $h(\cdot; y_m)$ is given by (2). The reconstructed signal $\hat{\mathbf{x}}$ should be approximately sparse and roughly consistent with the data.

The proposed formulation in (4) shares a 1-norm data fit term with recent matrix lifting phase retrieval methods [40], [41], but with greatly reduced memory requirements. Many other existing approaches implicitly (via projections) or explicitly minimize the quadratic negative log-likelihood representing a Gaussian distribution and are not designed to accommodate this data fit term, limiting their robustness to outliers. The competing GESPAR method [43] also is restricted to 0-“norm” sparsity (counts the number of nonzeros).

III. Majorization of the Measurement Objective

The inverse problem formulation of phase retrieval is particularly difficult to solve because having only magnitude information makes the data fit term in the objective function $\Psi(\mathbf{x})$ in (4) nonconvex. Although conventional methods like nonlinear conjugate gradients (NLCG) [47] can approximately minimize $\Psi(\mathbf{x})$, the more sophisticated approach proposed in this section facilitates much more rapid convergence. This approach begins by constructing a convex majorizer for $\Psi(\mathbf{x})$. Section IV describes an iterative method for minimizing this majorizer effectively.

A. Derivation of the Majorizer

A majorizer $\varphi(t; s, y)$ of the function $h(t; y)$ of t in (2) satisfies two properties: $\varphi(s; s, y) = h(s; y)$, and $\varphi(t; s, y) \geq h(t; y)$, for all t . Decreasing the majorizer value also reduces the value of the original function [48], so $h(t; y) < h(s; y)$ if t satisfies $\varphi(t; s, y) < \varphi(s; s, y)$. Assuming $f(\cdot)$ is convex and nondecreasing, and the majorizer $\varphi(t; s, y)$ is convex in its argument t , $f(\varphi(t; s, y))$ is also convex in t and majorizes $f(h(t; y))$ [49]. The approach below for finding $\varphi(t; s, y)$ is related to the concave-convex procedure [50], [51].

Let $h_+(t; y) = |t|^q - y$, and $h_-(t; y) = y - |t|^q$ be functions of t . Then, $h(t; y) = \max\{h_+(t; y), h_-(t; y)\}$. As $q \geq 1$, $h_+(t; y)$ is already convex in t , but $h_-(t; y)$ is concave in t . When $y \leq 0$, $h(t; y) = h_+(t; y)$. Otherwise, a majorizer $\varphi_-(t; s, y)$, convex in t , replaces $h_-(t; y)$. In this case, $\varphi(t; s, y) \triangleq \max\{h_+(t; y), \varphi_-(t; s, y)\}$ is convex in t and majorizes $h(t; y)$.

Since $h_-(t; y)$ is concave in t , its tangent plane about some point $s \in \mathbb{C}$ is a suitable convex majorizer:

$$\begin{aligned}\phi_{-}(t; s, y) &= (y - |s|^q) + (-q|s|^{q-1}) \Re\{e^{i\angle s}(t-s)\} \\ &= y + (q-1)|s|^q - q|s|^{q-1} \Re\{te^{-i\angle s}\}.\end{aligned}\quad (5)$$

When $|s|^q < y$, $\phi_{-}(t; s, y)$ is tight among convex majorizers. However, when $|s|^q > y$, s is in the convex region of $h(\cdot; y)$, and the tangent plane for $s \triangleq y^{1/q}e^{i\angle s}$ majorizes $h_{-}(t; y)$ more tightly in the range of $|t|^q - y$.

In summary, the majorizer for the function $h(t; y)$ of t is

$$\phi(t; s, y) = \begin{cases} h_{+}(t; y), & y \leq 0, \\ \max\{h_{+}(t; y), \phi_{-}(t; s, y)\}, & |s|^q < y, \\ \max\{h_{+}(t; y), \phi_{-}(t; \bar{s}, y)\}, & 0 < y \leq |s|^q. \end{cases} \quad (6)$$

In the first case, $h(t; y)$ is already convex in t . The second and third cases correspond to s being in the concave and convex regions of $h(\cdot; y)$, respectively. Figure 1 portrays examples of the function $h(t; y)$ and its surrogate $\phi(t; s, y)$ in both the second (s in concave region) and third (s in convex region) cases. Substituting $\phi(t; s, y)$ for $h(t; y)$ in the objective $\Psi(\mathbf{x})$ in (4) yields its majorizer $\Phi(\mathbf{x}; s)$, convex in $\mathbf{x}$:

$$\Phi(\mathbf{x}; s) = \sum_{m=1}^M f(\phi([A\mathbf{x}]_m; s_m, y_m)) + \beta \|\mathbf{x}\|_1. \quad (7)$$

Having constructed $\Phi(\mathbf{x}; s)$, the sequel describes how to minimize $\Phi(\mathbf{x})$ using this function.

Algorithm 1

Majorize-minimize scheme for solving (4).

Require: $I_{\text{mm}}, \varepsilon_{\text{mm}}$, random $s^0 \in \mathbb{C}^M$.

for $i = 1: I_{\text{mm}}$ **do**

$$\mathbf{x}^i \leftarrow \arg \min_{\mathbf{x}} \Phi(\mathbf{x}; s^{i-1}). \quad (8)$$

$$s^i \leftarrow A\mathbf{x}^i. \quad (9)$$

if $\|s^i - s^{i-1}\| < \varepsilon_{\text{mm}}$ **then break**

end if

end for

B. Majorize-Minimize (MM) Algorithm

The proposed approach to solving (4) uses the majorize-minimize (MM) scheme [48], [52] outlined in Algorithm 1. Each iteration of this MM method decreases $\Psi(\mathbf{x})$ by minimizing $\Phi(\mathbf{x}; s)$ over $\mathbf{x}$, converging to a critical point of $\Psi(\mathbf{x})$ when $\Psi(\cdot)$ and $\Phi(\cdot; s)$ are differentiable

at every non-critical majorization point $\mathbf{x} = \mathbf{s}$. Running the algorithm for multiple different initial choices of $\mathbf{s}^0$ increases the chance of finding a global optimum of the original nonconvex problem. Many phase retrieval methods also employ multiple initializations, as do nonconvex solvers more generally.

IV. Solving the Majorized Objective with ADMM

Jointly minimizing M pairwise maximum functions to minimize (7) directly would be combinatorially hard. Instead, introducing an auxiliary vector $\mathbf{u} = \mathbf{Ax}$, each function in the summation in (7) depends only on a single $u_m = [\mathbf{u}]_m$. The constrained problem using this auxiliary variable is

$$\{\mathbf{x}^{i+1}, \mathbf{u}\} \leftarrow \arg \min_{\mathbf{x}, \mathbf{u}} \sum_{m=1}^M f(\phi(u_m; s_m, y_m)) + \beta \|\mathbf{x}\|_1, \quad \text{s. t. } u_m = [\mathbf{Ax}]_m, \quad m=1, \dots, M. \quad (10)$$

The alternating direction method of multipliers (ADMM) framework [45], [53]–[55] uses the augmented Lagrangian of this constrained problem:

$$\mathcal{L}_A(\mathbf{x}, \mathbf{u}; \mathbf{b}) \triangleq \sum_{m=1}^M f(\phi(u_m; s_m, y_m)) + \beta \|\mathbf{x}\|_1 + \frac{\mu}{2} \|\mathbf{Ax} - \mathbf{u} + \mathbf{b}\|_2^2, \quad (11)$$

where $\mathbf{b} \in \mathbb{C}^M$ and $\mu > 0$ are the scaled dual vector (Lagrange multipliers) and augmented Lagrangian penalty parameter, respectively. The implementation of ADMM in Algorithm 2 minimizes (11), subject to $\mathbf{u} = \mathbf{Ax}$. To simplify notation here and in subsequent sections, define $d_m = [\mathbf{Ax} + \mathbf{b}]_m$. Initially, $\mathbf{x}^0$, $\mathbf{u}^0$, and $\mathbf{b}^0$ are set to $\mathbf{0}$. In later iterations, the last $\mathbf{x}$, $\mathbf{u}$, and $\mathbf{b}$ from the previous run of ADMM “warm-start” the next run. Methods for updating $\mathbf{x}$ and $\mathbf{u}$ depend on the specific $\mathbf{A}$ and $f(\cdot)$ used. This paper provides details for general $\mathbf{A}$ with the 1-norm data fit term.

Algorithm 2

ADMM method for solving (11).

Require: I_{ADMM} , $\varepsilon_{\text{ADMM}}$, $\mathbf{x}^0$, $\mathbf{u}^0$, $\mathbf{b}^0$, $\mathbf{y}$, β , μ .

for $i = 1: I_{\text{ADMM}}$ **do**

$$\mathbf{x}^i \leftarrow \arg \min_{\mathbf{x}} \beta \|\mathbf{x}\|_1 + \frac{\mu}{2} \|\mathbf{Ax} - (\mathbf{u}^{i-1} - \mathbf{b}^{i-1})\|_2^2. \quad (12)$$

for $m = 1: M$ **do**

$$d_m \leftarrow [\mathbf{Ax}^i + \mathbf{b}^{i-1}]_m.$$

$$\mathbf{u}_m^i \leftarrow \arg \min_{\mathbf{u}} f(\phi(u; s_m, y_m)) + \frac{\mu}{2} |u - d_m|^2. \quad (13)$$

end for

$$b^i \leftarrow b^{i-1} + Ax^i - u^i. \quad (14)$$

```

    if  $\|x^i - x^{i-1}\| < \varepsilon_{\text{ADMM}}$  then break
    end if
end for

```

A. Updating x

The update for x in the preceding ADMM framework has the extensively studied synthesis form of compressed sensing (CS) [56]–[59]. Various CS algorithms may be appropriate, depending on the structure of A .

If A is left-unitary, so that $A'A = I$, then the least-squares term in (12) simplifies to

$\|x - A'(u^i - b^i)\|_2^2$ plus a constant term. In this case, updating x becomes soft thresholding: $x_n^{i+1} \leftarrow \text{soft}([A'(u^i - b^i)]_n; \frac{\beta}{\mu})$, where

$$\text{soft}(x; \tau) = \frac{x}{|x|} \max\{|x| - \tau, 0\}. \quad (15)$$

Otherwise, an iterative algorithm like FISTA [60] could be embedded within the ADMM method [44]. Instead, we use “preconditioned” ADMM (PADMM) [61], [62] accelerated¹ using Nesterov momentum [64], essentially using a single FISTA step as the x -update in (12):

$$x^i \leftarrow \text{soft}(z^{i-1} - \frac{1}{c}A'(Az^{i-1} - u^{i-1} + b^{i-1}); \frac{\beta}{\mu c}). \quad (16)$$

$$t^i \leftarrow (1 + \sqrt{1 + 4(t^{i-1})^2})/2. \quad (17)$$

$$z^i \leftarrow x^i + \frac{t^{i-1} - 1}{t^i}(x^i - x^{i-1}). \quad (18)$$

The scalar c must satisfy $cI \succeq A'A$; it can be precomputed using power iterations, or found directly in many cases. For example, $c = 1$ for the undersampled unitary discrete Fourier transform (DFT) used in the experiments in this paper. “Gradient”-based adaptive restarting [65] can help avoid divergence: when the momentum term $x^i - x^{i-1}$ points away from $x^i - z^{i-1}$, the momentum is reset ($t^{i-1} = 0$, $t^i = 1$). While PADMM does not possess the same convergence guarantees as regular ADMM, faster convergence may be possible by adjusting the dual update in (14); see [66].

¹This method differs from accelerated ADMM [63] that applies momentum without introducing the separable majorizer simplifying the quadratic augmented Lagrangian penalty in (11) we depend on here.

B. Updating u

Because of the proposed variable-splitting, updating the auxiliary vector u can be performed element-by-element. Since $f(\cdot)$ is monotone nondecreasing, and $\varphi(u_m; s_m, y_m)$ is the pointwise maximum of two functions (for $y_m > 0$), $f(\varphi(u_m; s_m, y_m)) = \max\{f_+(u_m), f_-(u_m)\}$, where

$$\begin{aligned} f_+(u_m) &\triangleq \frac{\mu}{2}|u_m - d_m|^2 + f(h_+(u_m; y_m)), \\ f_-(u_m) &\triangleq \frac{\mu}{2}|u_m - d_m|^2 \end{aligned} \quad (19)$$

$$+ \begin{cases} 0, & y_m \leq 0, \\ f(\phi_-(u_m; s_m, y_m)), & |s_m|^q < y_m, \\ f(\phi_-(u_m; \bar{s}_m, y_m)), & 0 < y_m \leq |s_m|^q, \end{cases} \quad (20)$$

and $d_m = [Ax + b]_m$. Updating u_m is equivalent to solving

$$\arg \min_{u, T} T, \text{ s. t. } f_+(u) \leq T, f_-(u) \leq T. \quad (21)$$

The minimizing T corresponds to the value of $f(\varphi(u; s_m, y_m))$ at its minimum (with respect to u). The Lagrangian of (21) is $T + \gamma_+(f_+(u) - T) + \gamma_-(f_-(u) - T)$, with Lagrange multipliers $\gamma_+, \gamma_- \geq 0$. Differentiating yields $\gamma_+ + \gamma_- = 1$. Three possibilities exist:

1. $\gamma_+ = 1, \gamma_- = 0$: The optimal $u = u_+$ minimizes $f_+(u)$ and satisfies $f_+(u_+) > f_-(u_+)$.
2. $\gamma_+ = 0, \gamma_- = 1$: The optimal $u = u_-$ minimizes $f_-(u)$ and satisfies $f_-(u_-) > f_+(u_-)$.
3. $\gamma_+, \gamma_- > 0$: Both $f_+(u)$ and $f_-(u)$ equal T . The optimal $u = u_{\pm}$ minimizes both of these functions along the curve $f_+(u) = f_-(u)$.

For $f(\cdot)$ corresponding to the 1-norm data fit term in (3) on squared-magnitude measurements ($q = 2$), the optimal values of u for each case for the m th measurement are

$$u_+ = \frac{\mu}{2+\mu} d_m, \quad (22)$$

$$u_- = \frac{2s_m}{\mu} + d_m, \text{ and } \quad (23)$$

$$u_{\pm} = \sqrt{2(y_m + |s_m|^2)} e^{i\angle((2+\mu)s_m + \mu d_m)} - s_m. \quad (24)$$

When $|s_m|^q \rightarrow y$, we replace s_m above with s_m^- . The functions $f_+(u)$ and $f_-(u)$ are evaluated for each case to determine which of the three cases applies. These expressions, and corresponding expressions for quadratic $f(\cdot)$, are derived in the supplement.

C. Computational Complexity

The proposed algorithm consists of nested layers of iterative methods, adding complexity compared to simpler methods like nonlinear conjugate gradients (NLCG). Multiple initial values of s^0 are tested to increase the likelihood of finding a global minimum. For each initial value, several iterations of the MM algorithm in Algorithm 1 are run. Finally, for each outer iteration of the MM method, several inner iterations of ADMM (or PADMM) are performed.

Each iteration of ADMM/PADMM involves updating $\mathbf{x}$, $\mathbf{u}$, and $\mathbf{b}$. Updating $\mathbf{x}$ involves two matrix-vector products with $\mathbf{A}$ or $\mathbf{A}'$. Reusing the calculated value of $\mathbf{A}\mathbf{x}$ avoids recomputing it through the remainder of the iteration. When $\mathbf{A}$ is a DFT matrix, the cost is roughly $O(N \log N)$ for each. At least for the 1-norm data fit term with squared-magnitude measurements, each candidate u_m is a simple function of d_m , s_m , and y_m , so that the cost of updating $\mathbf{u}$ is roughly $O(M)$. Updating $\mathbf{b}$ is a simple addition, again scaling as $O(M)$. The overall cost of an ADMM iteration is $O(N \log N + M)$.

Without acceleration, the error in $\mathbf{x}$ converges roughly as $O(1/I_{\text{ADMM}})$ for preconditioned ADMM (I_{ADMM} is the number of iterations) [61]. Empirical convergence behavior of our ADMM implementation is established in the automatic ADMM parameter tuning experiment in Section V-B. Computational costs are reported along with the simulations in Section V. When transitioning from relatively small 1D experiments to a much larger 2D experiment, the number of MM iterations (I_{mm}) only increases modestly, and the number of PADMM iterations and initializations remains constant.

D. Parameter Selection

The regularization parameter β controls the level of sparsity in the reconstructed signal. Additionally, the ADMM penalty parameter μ impacts the convergence rate of the inner ADMM/PADMM algorithm. Introducing an adaptive heuristic for μ and a normalization factor for β avoids manual tuning of these parameters for every experiment.

For ADMM penalty parameter μ , the automatic heuristic in [45] and quadratic-optimal strategy in [67] provide alternatives to adjusting μ manually. The chosen adaptive method, described in [45], starts at some initial value and adapts μ every 10 ADMM iterations by comparing the residual $\mathbf{u}^i - \mathbf{A}\mathbf{x}^i$ and dual residual $\mu\mathbf{A}'(\mathbf{u}^i - \mathbf{u}^{i-1})$. This method is compared against using fixed (manually tuned) μ in Section V-B.

The choice of regularization parameter β , which reflects prior knowledge about the sparsity of the desired signal, also greatly influences the reconstruction. All the competing methods investigated in this paper use this type of parameter, or the related sparsity factor K . While K may be more-or-less known, learning β from K is not straightforward [59]. In the Monte Carlo simulations that follow, the optimum value of β varies based on the true 1-norm of $\mathbf{x}$ and the actual data discrepancy. Not knowing these a priori, this algorithm uses a simple normalization framework for β that requires only the measurements $\mathbf{y}$ and the approximate noise level/number of outliers. Differentiating $\Sigma_m f(h([\mathbf{A}\mathbf{x}]_m; y_m))$ with respect to $\mathbf{x}$, obtains (for a 1-norm data fit term with $q = 2$)

$$2\mathbf{A}'\mathbf{D}_{\text{noise}}\mathbf{A}\mathbf{x}, \quad (25)$$

where $\mathbf{D}_{\text{noise}}$ is a diagonal matrix with entries $[\mathbf{D}_{\text{noise}}]_{m,m} = \text{sign}([\mathbf{A}\mathbf{x}]_m/2 - y_m)$. To make this expression as consistent as possible as the noise level or number of measurements changes, the data fit term is normalized according to the 2-norm of (25). When $\mathbf{A}$ is an undersampled (unitary) DFT, the 2-norm becomes

$$\left(\sum_{m: |[\mathbf{A}\mathbf{x}]_m|^2 \neq y_m} |[\mathbf{A}\mathbf{x}]_m|^2 \right)^{1/2}.$$

Assuming zero-mean noise, the expected value of $[\mathbf{A}\mathbf{x}]_m/2$ is y_m . When y_m is an outlier, this is not the case, and $[\mathbf{A}\mathbf{x}]_m/2$ is approximated by the average value of the measurements not likely to be outliers. Assuming the M_{out} largest measurements are the most likely outliers, and $\bar{y}$ represents the arithmetic mean of the remaining measurements, the normalizer becomes

$$\left(M_{\text{out}}\bar{y} + \sum_{m: y_m \text{ not outlier}} y_m \right)^{1/2} = (M_{\text{out}}\bar{y} + (M - M_{\text{out}})\bar{y})^{1/2} = (M\bar{y})^{1/2}. \quad (26)$$

With this normalization, the proposed algorithm can be applied to a whole set of signals without manually tuning β for each one. Although outliers are unknown a priori, the estimation error of $\bar{y}$ should be small when $M_{\text{out}} \ll M$.

V. Experimental Setup and Results

Simulations throughout this paper consist of generating a length- N sparse signal with K nonzero coefficients, acquiring M samples of the squared-magnitude DFT of that signal, reconstructing the signal using the proposed and/or competing algorithms listed in Table I, and comparing the reconstructed signals against the true signal.

A. Experimental Setup

This section describes the general setup common to all experiments. These experiments are simulations, generating the sparse support of each true signal at random, and randomly sampling the amplitude and phase of each nonzero coefficient uniformly between 0 and 1 (amplitude) and 0 and 2π (phase).

For each simulated signal, M noise-free measurements are randomly selected from the squared-magnitude of the signal's DFT coefficients. Randomly selected outliers are set to have an amplitude between one and two times the maximum measurement. Additionally, Gaussian or Laplace noise (40 dB SNR unless stated otherwise) are added to all the measurements.

The reconstructions are performed using multiple initializations, and the “best” reconstructed signal for each method is retained. For the proposed method, 50 initializations are performed per trial, 100 for the fully-sampled ($M = N$) case, and the lowest value of $\Psi(\mathbf{x})$ determines the best reconstruction. The regularization parameter β is held fixed for the Monte Carlo experiments; the ADMM penalty parameter μ is automatically adapted [45], not manually tuned. Other reconstruction parameters are provided in Table II. Competing methods include the GESPAR greedy method [43], the L_1 -Fienup method (sparse Fienup [31] with the image-domain projection modified to project the signal onto the ℓ_1 -ball with radius β_{sf} , like [29]), and the message passing algorithm PR-GAMP [30]. These other methods are run for at least 50 initializations, but often more to allow for the same total amount of computation (measured via tracking the number of multiplies by $\mathbf{A}$ or $\mathbf{A}'$). The best reconstructions are chosen for L_1 -Fienup, GESPAR, and PR-GAMP according to the smallest 2-norm data discrepancy. In the supplement, the proposed method is compared with compressive matrix lifting (CPRL) [37]. As CPRL requires significantly more memory to run, with a length-128 complex signal requiring upwards of 17 GB of memory, the experiment featuring CPRL uses a much smaller signal ($N = 64$).

Sparsity and Fourier coefficient magnitudes are invariant to spatial shifts, reversal, and global phase. Thus, the error computation is relative to the best alignment/reversal and global phase for each reconstructed signal. The best alignment is identified for both the reconstructed signal and its reversed version by cross-correlation with the true signal. A global phase term is then estimated from the version with the best alignment. Reconstruction errors are reported relative to the true signal using the median of the squared errors (normalized by N) over the set of trials. This peak-signal-to-error ratio (PSER) is converted to dB scale:

$$\text{PSER} = -10 \log_{10}(\text{median squared error}), \quad (27)$$

where the maximum true signal amplitude is one.

B. Validating Parameter Selection Methods

The first experiment compares convergence of the majorizer objective value in (7) for automatically adapted μ and fixed, manually tuned μ . Simulations of a 1D signal are repeated for both $K = 6$ and $K = 8$ sparse coefficients, and both $M = 64$ (undersampled case, using PADMM) and $M = N = 128$ (fully-sampled case, using ADMM) squared-magnitude measurements, corrupted by both additive Gaussian noise (40 dB SNR) and 5 outliers. For each experiment, we run one set of ADMM/PADMM iterations with the 1-norm data fit distribution, some with a fixed penalty parameter μ (only the best are shown), and others with the adaptive method, starting from different initial values. For these experiments, the regularization parameter β is chosen to portray a range of convergence behaviors, not to optimize the reconstruction.

For sparsity $K = 6$ and both $M = 64$ and $M = N = 128$ noisy measurements, Figure 2 portrays the objective function convergence rates over $I_{\text{ADMM}} = 100$ ADMM/PADMM iterations for the three best fixed choices of μ , relative to the best overall objective function value observed after 400 iterations. These are compared against the adaptive method starting at the

best ($\mu = 1$) and a suboptimal ($\mu = 0.1$) initial value. This experiment verifies that the adaptive method achieves nearly as good convergence as the best fixed method in both the undersampled (using PADMM) and fully-sampled (using ADMM) cases, even when not initialized to the best choice of μ . The same experiment for different sparsity $K = 8$ yields similar results to the example shown. Since the adaptive method appears to ensure rapid convergence across varying degrees of measurements and sparsities, this adaptive heuristic scheme with initial $\mu = 1$ is employed throughout the experiments that follow, without any additional tuning.

To observe how sensitive the regularization parameter β with the proposed normalization factor is as the sparsity level K or number of measurements M varies, the proposed algorithm is evaluated on sets of 50 simulated signals, each of whose squared-magnitude measurements are corrupted with additive Gaussian noise (40 dB SNR) and 5 outliers. In [44], β scales roughly linearly with the number of measurements for the proposed method without normalization. With normalization, the optimal β appears to remain fairly constant between 0.1 and $10^{-0.9}$. Figure 3 plots the median squared error (lines) and error quartiles (boxes) versus the regularization parameter β for (a) different sparsity levels $K = 3, 5, 6, 8$, holding $M = N = 128$ fixed, and for (b) different measurements $M = 32, 64, 96, 128$, holding $K = 3$ fixed. The β values found in this experiment are fixed and reused in all the Monte Carlo experiments, regardless of noise level or type, with no further tuning.

To ensure competing methods are not at a disadvantage, both GESPAR and PR-GAMP are provided the true sparsity (K) for each signal. For the L_1 -Fienup method, the radius β_{sf} of the ℓ_1 -ball constraint is set to the 1-norm of the true signal.

C. Rapid Convergence with Preconditioned ADMM

The robust phase retrieval problem described in (4) can be solved via conventional methods including nonlinear conjugate gradients (NLCG), if the 1-norm is approximated by a differentiable function. However, close approximations to the 1-norm have a high curvature that slow convergence of NLCG. We compared empirically the convergence rates of NLCG and the proposed algorithm. Representative length-128 signals, one with sparsity $K = 6$ and $M = 64$ noisy measurements, and the other with sparsity $K = 8$ and $M = 128$ noisy data (both 40 dB SNR Gaussian noise and 5 outliers), are reconstructed using both methods. First, the MM method with adaptive preconditioned ADMM is run for 50 initializations, and the best result (minimum objective value) is kept. Then, the NLCG method is run for that same best initializer, for a number of iterations equivalent to the total number of inner iterations of the preconditioned ADMM method. The objective function in (4) is plotted for each NLCG iteration (solid line) and every MM iteration (circles) in Fig. 4. The plotted objective functions converge at very different rates, with a distinct advantage to the proposed MM algorithm with adaptive preconditioned ADMM.

D. Monte Carlo Comparisons (1D)

This section compares the proposed phase retrieval method against the competing methods listed in Table I via 50-trial Monte Carlo simulations with different values of sparsity K , number of measurements M , and noise/outlier levels and types. All the comparisons in this

section involve length-128 1D signals and measurements corrupted with both outliers and either Gaussian or Laplace noise. The same β values identified in Section V-B are reused here for all types of noise.

The first test evaluates the proposed algorithm on measurements corrupted by Gaussian noise (40 dB SNR) and 5 outliers. Figure 5 depicts PSER values corresponding to median squared errors for the proposed and competing methods. Equivalent comparisons for measurements corrupted by Laplace noise (also 40 dB SNR) and 5 outliers are shown in Figure 6. The median squared errors for the proposed method show significant improvement over competing methods in both cases. The supplement depicts the PSER values for the mean squared errors in both experiments. Regarding runtimes, as measured by the multiplications by $\mathbf{A}$ or $\mathbf{A}'$ (the dominant computations), the proposed method runs for the same number of iterations for all sparsities K and measurements M , except the number of initializations is doubled for $M = N$ for robustness (this is responsible for the upper limit on the range of multiplies for all the methods). Otherwise, the amount of computation remains nearly constant.

To see how the noise level or number of outliers affects reconstruction quality, Monte Carlo simulations are conducted for the Gaussian noise + outliers case, varying the number and variance of outliers and SNR of the additive noise. Figure 7 shows median PSER values for $K = 3$ sparse signals ($N = 128$), whose measurements are corrupted by 2, 4, 8, 16 outliers, with a range $[1, 2]$ times the maximum measurement value, holding the Gaussian noise SNR fixed at 40 dB. Supplementary material contains a similar figure for a smaller outlier range $[1, \sqrt{2}]$ and for $K = 5$ sparse signals. Figure 8 depicts improvements for $K = 5$ sparse signals ($N = 128$) with measurements with 20, 30, 40, 50, 60 dB SNR Gaussian noise, holding the number and range of outliers fixed at 5 and $[1, 2]$, respectively. The supplement contains the corresponding plot for $K = 3$. The improvement in squared error appears significant over a wide range of noise levels and numbers of outliers.

E. Image Comparisons (2D)

This experiment examines image reconstruction with undersampled measurements corrupted by outliers and additive Gaussian noise. The $N = 512 \times 512$ -pixel star of David phantom used in [44] is inspired by the real example image shown in [68]. The pattern in the image is constructed using 30 discs, each 21-pixels wide. A dictionary of these discs (at all 512×512 positions) is used as the synthesis transform for all the reconstructions. Since the dictionary is shift-invariant, implementing the dictionary via multiplication in frequency saves computation and storage for all the methods. The squared-magnitudes of the 2D DFT of this image are randomly undersampled by a factor of two ($M = N/2 = 131,072$). One percent of the measurements are changed to outliers, and 60 dB SNR additive Gaussian noise is added to all measurements. The phantom is reconstructed using both the proposed and competing algorithms, resulting in the images in Figure 9.

To conserve space, the blank image produced by the PR-GAMP method is not shown here. Reconstructions of the same image with fewer outliers are provided in the supplementary material. In terms of scalability, the proposed method works well without much adjustment;

only the number of MM iterations I_{mm} changes (from 10 to 20), doubling the number of multiplies versus the proposed method in the 1D case. The ADMM penalty parameter remains the same (adaptive) as in the 1D experiments, but we reexamine the regularization parameters for all the methods. The true K and β_{sf} are fixed, while several values for β between 0.1 and 0.4 are tested to account for the smaller K/N of the 2D image. The advantage of the proposed method is clear, as none of the competing methods recovered the true image, even when running GESPAR, PR-GAMP, and L1-Fienup for at least as many initializations and at least as much (often, much more) computation as the proposed method. In the supplement, this improvement in quality is apparent even with extremely few outliers.

VI. Discussion

Undersampled phase retrieval relies heavily on side information to reproduce a quality image. Employing sparsity in the image domain, or dictionary-based sparsity, helps identify the best image among all those that share the same magnitude Fourier spectrum. Resolving this ambiguity becomes even more challenging in the face of measurement noise, especially outliers. The proposed method using a 1-norm data fit term excels at reconstructing images despite these conditions, greatly improving upon other techniques for such data, even after giving faster methods equivalent computation (via more initializations).

The proposed method differs from existing work in two ways: a robust data fit model and a nested MM+ADMM algorithm for reconstruction with this model. Although this algorithm can be generalized, including to the conventional quadratic data fit term, preliminary experiments (not shown) do not portray the same level of robustness with the ℓ_2 approach as the proposed method with the ℓ_1 data fit term. Thus, the benefit likely derives from the data fit term. This hypothesis is consistent with the fact that competing methods perform well in settings without outliers. However, experimental results comparing the proposed algorithm to a conventional gradient method also suggest that the algorithm is important, as the gradient method converges very slowly and would not yield a quality result with the same amount of computation. Although existing methods may possess theoretic convergence guarantees, the faster empirical convergence of the objective function in (4) using the proposed method means that the model and algorithmic contributions are intertwined, and both are needed to achieve robustness with outliers. In SNR-limited applications like point spread function estimation in super-resolution optical microscopy, the additional robustness provided by the 1-norm model versus the Gaussian model will greatly simplify the acquisition and reduce noise-related errors in the phase retrieval reconstruction.

Some existing methods automatically tune parameters, like PR-GAMP [30]. With the normalization and adaptive methods for parameter selection we describe, the Monte Carlo simulations reveal significantly reduced errors versus other methods, even without extensive manual parameter tuning. A complete solution to parameter selection would rely on more sophisticated automatic methods [69]. Further experiments on larger, real datasets are necessary to fully describe parameter selection and assess real performance of the proposed method.

Paired with parameter selection, multiple initializations are also important to overcome the nonconvexity of the inverse problem and find a reasonable global solution. Although recently proposed techniques like Wirtinger flow [70] show promise for the oversampled case, randomly selecting multiple initial majorization vectors s^0 appears to be more robust for the proposed method. As using multiple initial choices for s^0 proportionally increases computation time, the overall reconstruction time may be an issue in higher dimensions. However, a suitable 2D image reconstruction is obtained with the same number of initializations (50) as in the 1D case. Still, the 2D reconstructions all took over two hours on a modern workstation using MATLAB. Compared to much faster, simpler methods like alternating projections, that can recover an image in seconds or minutes (in the absence of outliers), the proposed method is suitable when obtaining a quality reconstruction is paramount, or when those faster methods fail to recover the true image (like in Figure 9).

In both the 1D and 2D cases, the proposed method clearly outperforms the L_1 -modified sparse Fienup method, GESPAR, and even PR-GAMP, when outliers are present in the data, even when controlling for computation time. As the reduced squared error is prominent for the extremely sparse signals evaluated here, the 1-norm sparsity term should allow for similar improvements for signals that are less sparse or compressible. This quality gain is not without cost, as the mean squared error (in supplementary material) shows greater variability than the median numbers, suggesting noticeable errors are generally larger versus other methods. In the future, this framework will be extended to image domain constraints like nonnegativity and other forms of regularization, including analysis-form sparsity. These additions should facilitate reconstruction of real images.

VII. Conclusion

The key contributions of this paper are two-fold and intertwined. A general framework is proposed that extends phase retrieval reconstruction to measurements corrupted by outliers. A new implementation of this general framework is described featuring multiple initializations, majorization-minimization, and (preconditioned) ADMM. In addition, using normalization and existing adaptive heuristics, the proposed method is made robust without manual tuning as noise levels/types or numbers of outliers change. A direct comparison against competing methods establishes quantitative and visible advantages over existing methods, over a wide range of simulations.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

DSW was funded by National Institutes of Health (NIH) grant F32 EB015914. JAF is funded in part by NIH grant U01 EB018753 and an equipment donation from Intel. YCE is funded in part by Israel Science Foundation Grant 170/10, SRC, and Intel Collaborative Research Institute for Computational Intelligence.

The authors would like to acknowledge Yoav Shechtman for insights relating to phase retrieval and coherent diffraction imaging, and for sharing image data, and James Fienup for general discussions on phase retrieval. A special thanks also goes out to the reviewers, for providing substantive and constructive feedback on multiple iterations of this paper.

Biographies

Daniel S. Weller (S'05-M'12) joined the University of Virginia in Charlottesville, VA, in August 2014, where he is an Assistant Professor of Electrical and Computer Engineering, with courtesy appointments in Biomedical Engineering and in Radiology and Medical Imaging. From 2012–2014, he was a post-doctoral research fellow in the Electrical Engineering - Systems program at the University of Michigan, in Ann Arbor, MI, supported by the National Institutes of Health via a Ruth L. Kirschstein National Research Service Award. He received his Ph.D. and S.M. in electrical engineering from the Massachusetts Institute of Technology, in Cambridge, MA, in June 2012 and June 2008, respectively, where he worked in the Signal Transformation and Information Representation Group in the Research Laboratory of Electronics. Previously, he received his B.S. in electrical and computer engineering, with honors, from Carnegie Mellon University, in Pittsburgh, PA, in May 2006.

At Carnegie Mellon, he was a research assistant with the General Motors Collaborative Research Laboratory. He also worked at Texas Instruments, Apple, and ATK. He is a member of Tau Beta Pi, Eta Kappa Nu, IEEE, and the International Society for Magnetic Resonance in Medicine (ISMRM). He currently serves as an associate editor for the IEEE Transactions on Medical Imaging, and has reviewed for numerous IEEE and other professional journals and conferences. He was a finalist in the Student Paper Competition at the 2011 IEEE International Symposium on Biomedical Imaging. He received a National Defense Science and Engineering Graduate fellowship and a graduate research fellowship from the National Science Foundation. His research interests include medical imaging, especially magnetic resonance imaging, signal processing and estimation, and nonideal sampling and reconstruction.

Ayelet Pnueli has a PhD in Physics from the Technion (Israel Institute of Technology). She worked many years in various R&D positions in the industry including IAI (Israeli aircraft industry), Applied Materials Inc., and HP labs. Currently she is a designer of digital educational games.

Gilad Divon received his B.Sc degree in Electrical Engineering and his B.A degree in physics from the Technion - Israel institute of technology, in Haifa, Israel, in 2015. During his studies, he took part in research, conducted in the Signal Acquisition Modeling and Processing Lab, supervised by Prof Yonina Eldar. He also worked as a student and as a system engineer at Elbit Systems Ltd, in the civil aviation the electro-optical R&D groups. In 2015 he started his master degree in electrical engineering in the field of computer vision and graphic, dealing with saliency of 3D models.

Ori Radzyner received his B.Sc degree in Electrical Engineering and his B.A degree in physics from the Technion - Israel institute of technology, in Haifa, Israel, in 2015. During his studies, he took part in research, conducted in the Signal Acquisition Modeling and Processing Lab. He also worked as a student at Elbit Systems Ltd. in one of the electro-optical R&D groups. Nowadays he works as an engineer in the semiconductor industry.

Yonina C. Eldar (S'98-M'02-SM'07-F'12) received the B.Sc. degree in Physics in 1995 and the B.Sc. degree in Electrical Engineering in 1996 both from Tel-Aviv University (TAU), Tel-Aviv, Israel, and the Ph.D. degree in Electrical Engineering and Computer Science in 2002 from the Massachusetts Institute of Technology (MIT), Cambridge.

From January 2002 to July 2002 she was a Postdoctoral Fellow at the Digital Signal Processing Group at MIT. She is currently a Professor in the Department of Electrical Engineering at the Technion - Israel Institute of Technology, Haifa, Israel, where she holds the Edwards Chair in Engineering. She is also a Research Affiliate with the Research

Laboratory of Electronics at MIT and was a Visiting Professor at Stanford University, Stanford, CA. Her research interests are in the broad areas of statistical signal processing, sampling theory and compressed sensing, optimization methods, and their applications to biology and optics.

Dr. Eldar has received numerous awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award (2013), the IEEE/AESS Fred Nathanson Memorial Radar Award (2014), and the IEEE Kiyo Tomiyasu Award (2016). She was a Horev Fellow of the Leaders in Science and Technology program at the Technion and an Alon Fellow. She received the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, the Wolf Foundation Krill Prize for Excellence in Scientific Research, the Henry Taub Prize for Excellence in Research (twice), the Hershel Rich Innovation Award (three times), the Award for Women with Distinguished Contributions, the Andre and Bella Meyer Lectureship, the Career Development Chair at the Technion, the Muriel & David Jacknow Award for Excellence in Teaching, and the Technions Award for Excellence in Teaching (2 times). She received several best paper awards and best demo awards together with her research students and colleagues including the SIAM outstanding Paper Prize and the IET Circuits, Devices and Systems Premium Award, and was selected as one of the 50 most influential women in Israel.

She is a member of the Young Israel Academy of Science and Humanities and the Israel Committee for Higher Education, and an IEEE Fellow. She is the Editor in Chief of Foundations and Trends in Signal Processing, a member of the IEEE Sensor Array and Multichannel Technical Committee and serves on several other IEEE committees. In the past, she was a Signal Processing Society Distinguished Lecturer, member of the IEEE Signal Processing Theory and Methods and Bio Imaging Signal Processing technical committees, and served as an associate editor for the IEEE Transactions On Signal Processing, the EURASIP Journal of Signal Processing, the SIAM Journal on Matrix Analysis and Applications, and the SIAM Journal on Imaging Sciences. She was Co-Chair and Technical Co-Chair of several international conferences and workshops.

She is author of the book "Sampling Theory: Beyond Bandlimited Systems" and co-author of the books "Compressed Sensing" and "Convex Optimization Methods in Signal Processing and Communications", all published by Cambridge University Press.

Jeffrey A. Fessler (F'06) received the BSEE degree from Purdue University in 1985, the MSEE degree from Stanford University in 1986, and the M.S. degree in Statistics from Stanford University in 1989. From 1985 to 1988 he was a National Science Foundation

Graduate Fellow at Stanford, where he earned a Ph.D. in electrical engineering in 1990. He has worked at the University of Michigan since then. From 1991 to 1992 he was a Department of Energy Alexander Hollaender Post-Doctoral Fellow in the Division of Nuclear Medicine. From 1993 to 1995 he was an Assistant Professor in Nuclear Medicine and the Bioengineering Program. He is now a Professor in the Departments of Electrical Engineering and Computer Science, Radiology, and Biomedical Engineering. He became a Fellow of the IEEE in 2006, for contributions to the theory and practice of image reconstruction. He received the Francois Erbsmann award for his IPMI93 presentation, and the Edward Hoffman Medical Imaging Scientist Award in 2013. He has served as an associate editor for IEEE Transactions on Medical Imaging, the IEEE Signal Processing Letters, and the IEEE Transactions on Image Processing, and is currently serving as an associate editor for the IEEE Transactions on Computational Imaging. He has chaired the IEEE T-MI Steering Committee and the ISBI Steering Committee. He was co-chair of the 1997 SPIE conference on Image Reconstruction and Restoration, technical program co-chair of the 2002 IEEE International Symposium on Biomedical Imaging (ISBI), and general chair of ISBI 2007. His research interests are in statistical aspects of imaging problems, and he has supervised doctoral research in PET, SPECT, X-ray CT, MRI, and optical imaging problems.

References

1. Fienup JR. Phase retrieval algorithms: a personal tour [Invited]. *Appl Optics*. Jan; 2013 52(1):45–56.
2. Shechtman Y, Eldar YC, Cohen O, Chapman HN, Miao J, Segev M. Phase retrieval with application to optical imaging: A contemporary overview. *IEEE Signal Processing Magazine*. May; 2015 32(3): 87–109.
3. Klibanov MV, Sacks PE, Tikhonravov AV. The phase retrieval problem. *Inverse Prob*. Feb; 1995 11(1):1–28.
4. Sayre D. Some implications of a theorem due to Shannon. *Acta Cryst*. 1952; 5:843.
5. Millane RP. Phase retrieval in crystallography and optics. *J Opt Soc Am A*. Mar; 1990 7(3):394–411.
6. Hauptman HA. The phase problem of X-ray crystallography. *Reports on Progress in Physics*. Nov; 1991 54(11):1427–54.
7. Harrison RW. Phase problem in crystallography. *J Opt Soc Am A*. May; 1993 10(5):1046–55.
8. Walther A. The question of phase retrieval in optics. *Optica Acta: Intl J of Optics*. 1963; 10(1):41–9.
9. Fienup, JR.; Dainty, JC. Phase retrieval and image reconstruction for astronomy. In: Stark, H., editor. *Image Recovery: Theory and Application*. San Diego: Academic; 1987. p. 231–275.
10. Miao J, Charalambous P, Kirz J, Sayre D. Extending the methodology of X-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature*. Jul; 1999 400(6742):342–4.
11. Balan R, Casazza P, Edidin D. On signal reconstruction without phase. *Applied and Computational Harmonic Analysis*. May; 2006 20(3):345–356.
12. Setsompop K, Wald LL, Alagappan V, Gagoski BA, Adalsteinsson E. Magnitude least squares optimization for parallel radio frequency excitation design demonstrated at 7 Tesla with eight channels. *Mag Res Med*. Apr; 2008 59(4):908–15.
13. Chai A, Moscoso M, Papanicolaou G. Array imaging using intensity-only measurements. *Inverse Prob*. Jan.2011 27(1):015005.
14. Latychevskaia T, Longchamp JN, Fink HW. Novel Fourier-domain constraint for fast phase retrieval in coherent diffraction imaging. *Optics Express*. Sep; 2011 19(20):19 330–9.
15. Oppenheim AV, Lim JS. The importance of phase in signals. *Proc IEEE*. May; 1981 69(5):529–41.

16. Gerchberg RW, Saxton WO. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*. Apr.1972 35:237–46.
17. Fienup JR. Reconstruction of an object from the modulus of its Fourier transform. *Optics Letters*. Jul; 1978 3(1):27–9. [PubMed: 19684685]
18. Fienup JR. Phase retrieval algorithms: a comparison. *Appl Optics*. Aug; 1982 21(15):2758–69.
19. Hayes MH, Quatieri TF. Recursive phase retrieval using boundary conditions. *J Opt Soc Am*. Nov; 1983 73(11):1427–1433.
20. Fienup JR. Phase retrieval using boundary conditions. *J Opt Soc Am A*. Feb; 1986 3(2):284–288.
21. Fienup JR. Reconstruction of a complex valued object from the modulus of its Fourier transform using a support constraint. *J Opt Soc Am A*. Jan; 1987 4(1):118–23.
22. Elser V. Phase retrieval by iterated projections. *J Opt Soc Am A*. Jan; 2003 20(1):40–55.
23. Marchesini S. A unified evaluation of iterative projection algorithms for phase retrieval [Invited]. *Review of Scientific Instruments*. Mar.2007 78(1):011301. [PubMed: 17503899]
24. Marchesini S. Phase retrieval and saddle-point optimization. *J Opt Soc Am A*. Oct; 2007 24(10):3289–3296.
25. Bauschke HH, Combettes PL, Luke DR. Hybrid projection-reflection method for phase retrieval. *J Opt Soc Am A*. Jun; 2003 20(6):1025–34.
26. Ohlsson H, Eldar YC. On conditions for uniqueness in sparse phase retrieval. *Proc IEEE Conf Acoust Speech Sig Proc*. 2014:1841–5.
27. Ranieri, J.; Chebira, A.; Lu, YM.; Vetterli, M. Phase retrieval for sparse signals: Uniqueness conditions. 2013. arxiv 1308.3058. [Online]. Available: <http://arxiv.org/abs/1308.3058>
28. Eldar YC, Mendelson S. Phase retrieval: Stability and recovery guarantees. *Applied and Computational Harmonic Analysis*. May; 2014 36(3):473–94.
29. Moravec ML, Romberg JK, Baraniuk RG. Compressive phase retrieval. *Proc SPIE 6701 Wavelets XII*. 2007:670120.
30. Schniter P, Rangan S. Compressive phase retrieval via generalized approximate message passing. *Proc 50th Allerton Conf on Comm, Control, and Computing*. 2012:815–822.
31. Mukherjee S, Seelamantula CS. An iterative algorithm for phase retrieval with sparsity constraints: application to frequency domain optical coherence tomography. *Proc IEEE Conf Acoust Speech Sig Proc*. 2012:553–6.
32. Osherovich E, Zibulevsky M, Yavneh I. Approximate fourier phase information in the phase retrieval problem: what it gives and how to use it. *J Opt Soc Am A*. Oct; 2011 28(10):2124–2131.
33. Candès EJ, Strohmer T, Voroninski V. PhaseLift: exact and stable signal recovery from magnitude measurements via convex programming. *Comm Pure Appl Math*. Aug; 2013 66(8):1241–74.
34. Candès EJ, Eldar YC, Strohmer T, Voroninski V. Phase retrieval via matrix completion. *SIAM J Imaging Sci*. 2013; 6(1):199–225.
35. Demanet, L.; Jugnon, V. Convex recovery from interferometric measurements. 2013. arxiv 1307.6864. [Online]. Available: <http://arxiv.org/abs/1307.6864>
36. Shechtman Y, Eldar YC, Szameit A, Segev M. Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing. *Optics Express*. Aug; 2011 19(16):14807–22.
37. Ohlsson, H.; Yang, AY.; Dong, R.; Sastry, SS. Compressive phase retrieval from squared output measurements via semidefinite programming. 2012. arxiv 1111.6323. [Online]. Available: <http://arxiv.org/abs/1111.6323>
38. Li X, Voroninski V. Sparse signal recovery from quadratic measurements via convex programming. *SIAM J Math Anal*. 2013; 45(5):3019–33.
39. Waldspurger I, d’Aspremont A, Mallat S. Phase recovery, maxcut and complex semidefinite programming. *Mathematical Programming*. Dec.2013 :1–35.
40. Candès EJ, Li X. Solving quadratic equations via PhaseLift when there are about as many equations as unknowns. *Foundations of Computational Mathematics*. 2014; 14(5):1017–26.
41. Hand, P. PhaseLift is robust to a constant fraction of arbitrary errors. 2015. arxiv 1502.04241v1. [Online]. Available: <http://arxiv.org/abs/1502.04241v1>

42. Jaganathan K, Oymak S, Hassibi B. Recovery of sparse 1-d signals from the magnitudes of their Fourier transform. *Intl Symp on Information Theory*. 2012:1473–7.
43. Shechtman Y, Beck A, Eldar YC. GESPAR: efficient phase retrieval of sparse signals. *IEEE Trans Sig Proc*. Feb; 2014 62(4):928–38.
44. Weller DS, Pnueli A, Radzyner O, Divon G, Eldar YC, Fessler JA. Phase retrieval of sparse signals using optimization transfer and ADMM. *Proc IEEE Intl Conf on Image Processing*. 2014:1342–6.
45. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found & Trends in Machine Learning*. 2010; 3(1):1–122.
46. Tibshirani R. Regression shrinkage and selection via the lasso. *J Royal Stat Soc Ser B*. 1996; 58(1):267–88.
47. Fletcher R, Reeves CM. Function minimization by conjugate gradients. *Comput J*. 1964; 7(2):149–54.
48. Jacobson MW, Fessler JA. An expanded theoretical treatment of iteration-dependent majorize-minimize algorithms. *IEEE Trans Im Proc*. Oct; 2007 16(10):2411–22.
49. Boyd, S.; Vandenberghe, L. *Convex optimization*. Cambridge; 2004.
50. Yuille AL, Rangarajan A. The concave-convex procedure. *Neural Computation*. Apr; 2003 15(4): 915–36. [PubMed: 12689392]
51. Kim K, Son YD, Bresler Y, Cho ZH, Ra JB, Ye JC. Dynamic PET reconstruction using temporal patch-based low rank penalty for ROI-based brain kinetic analysis. *Phys Med Biol*. Mar.2015 60(5):2019. [PubMed: 25675392]
52. Lange K, Hunter DR, Yang I. Optimization transfer using surrogate objective functions. *J Computational and Graphical Stat*. Mar; 2000 9(1):1–20.
53. Glowinski R, Marrocco A. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de dirichlet non linéaires. *Modélisation Mathématique et Analyse Numérique*. 1975; 9(R2):41–76.
54. Gabay D, Mercier B. A dual algorithm for the solution of nonlinear variational problems via finite-element approximations. *Comput Math Appl*. 1976; 2(1):17–40.
55. Eckstein J, Bertsekas DP. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*. Apr; 1992 55(1–3):293–318.
56. Chen SS, Donoho DL, Saunders MA. Atomic decomposition by basis pursuit. *SIAM J Sci Comp*. 1998; 20(1):33–61.
57. Donoho DL, Elad M. Optimally sparse representation in general (nonorthogonal) dictionaries via 1 minimization. *Proc Natl Acad Sci*. Mar; 2003 100(5):2197–2202. [PubMed: 16576749]
58. Tropp JA. Greed is good: algorithmic results for sparse approximation. *IEEE Trans Info Theory*. Oct; 2004 50(10):2231–42.
59. Eldar, Y.; Kutyniok, G. *Compressed sensing: Theory and applications*. Cambridge; 2012.
60. Beck A, Teboulle M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J Imaging Sci*. 2009; 2(1):183–202.
61. Chambolle A, Pock T. A first-order primal-dual algorithm for convex problems with applications to imaging. *J Math Im Vision*. 2011; 40(1):120–145.
62. Ouyang Y, Chen Y, Lan G, Pasiliao E Jr. An accelerated linearized alternating direction method of multipliers. *SIAM J Imaging Sci*. 2015; 8(1):644–81.
63. Goldstein T, O’Donoghue B, Setzer S. Fast alternating direction optimization methods. *SIAM J Imaging Sci*. 2014; 7(3):1588–1623.
64. Nesterov Y. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Math Dokl*. 1983; 27(2):372–76.
65. O’Donoghue B, Candès E. Adaptive restart for accelerated gradient schemes. *Found Computational Math*. 2014
66. Fazel M, Pong TK, Sun D, Tseng P. Hankel matrix rank minimization with applications to system identification and realization. *SIAM J Matrix Anal Appl*. 2013; 34(3):946–77.

67. Ghadimi E, Teixeira A, Shames I, Johansson M. Optimal parameter selection for the alternating direction method of multipliers (ADMM): quadratic problems. *IEEE Trans Auto Control*. Mar; 2015 60(3):644–58.
68. Szameit A, Shechtman Y, Osherovich E, Bulkich E, Sidorenko P, Dana H, Steiner S, Kley EB, Gazit S, Cohen-Hyams T, Shoham S, Zibulevsky M, Yavneh I, Eldar YC, Cohen O, Segev M. Sparsity-based single-shot subwavelength coherent diffractive imaging. *Nature Materials*. Apr. 2012 11:4559.
69. Weller DS, Ramani S, Nielsen JF, Fessler JA. Monte Carlo SURE-based parameter selection for parallel magnetic resonance imaging reconstruction. *Mag Res Med*. May; 2014 71(5):1760–70.
70. Candès EJ, Li X, Soltanolkotabi M. Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Trans Info Theory*. Apr; 2015 61(4):1985–2007.

Fig. 1.

The data fit error $h(t; y)$ (blue solid line) and the convex majorizer $\phi(t; s, y)$ (red dashed line) are plotted for real t , $y = 1$, and $q = 2$. Circles highlight the majorization points s for both examples. In the left figure, s is in the concave region of $h(\cdot; y)$, so the tangent plane at s is used in this region. In the right figure, s is located in the convex region of $h(\cdot; y)$, and the tangent plane at $y^{1/q}e^{i\angle s}$ is used instead.

Fig. 2.

The objective function $\Phi(x^i; s)$, relative to converged value Φ^* , is plotted versus ADMM/PADMM iteration i for ADMM/PADMM using fixed (thin, color lines) and adaptive (thick, black lines) penalty parameters (μ).

Fig. 3.

The median (lines) and quartiles (boxes) of 50 trials reconstructed using the proposed method with a 1-norm data fit term are plotted versus regularization parameter β for varying (a) signal sparsity levels K , and (b) measurements M . The signal length $N = 128$.

Fig. 4.

The objective function $\Psi(\mathbf{x}^i)$ in (4) is plotted versus NLCG iteration i and the equivalent MM iteration for both NLCG (solid line) and MM with adaptive preconditioned ADMM (circles), for K -sparse length- N signals from length- M noisy data (40 dB SNR AWGN noise, 5 outliers).

Fig. 5.

The PSER of 50 trials reconstructed using GESPAR, the proposed method, PR-GAMP, and L_1 -Fienup, for a range of measurement (M/N) and sparsity fractions (K/N), for measurements with 40 dB SNR Gaussian noise and 5 outliers. Computations (1000's of multiplies by A, A'): 110–1354 (GESPAR), 94–218 (proposed), 111–250 (PR-GAMP), 103–218 (L_1 -Fienup).

Fig. 6.

The PSER of 50 trials reconstructed using GESPAR, the proposed method, PR-GAMP, and L₁-Fienup, for a range of measurement (M/N) and sparsity fractions (K/N), for measurements with 40 dB SNR Laplace noise and 5 outliers. Computations (1000's of multiplies by A, A'): 110–1361 (GESPAR), 94–218 (proposed), 110–294 (PR-GAMP), 103–218 (L₁-Fienup).

Fig. 7.

The PSER of 50 trials reconstructed using GESPAR, the proposed method, PR-GAMP, and L_1 -Fienup, for a range of measurement (M/N) and outliers, for $K = 3$ and measurements with 40 dB SNR Gaussian noise. Computations (1000's of multiplies by $\mathbf{A}, \mathbf{A}'$): 109–537 (GESPAR), 103–217 (proposed), 110–248 (PR-GAMP), 108–218 (L_1 -Fienup).

Fig. 8.

The PSER of 50 trials reconstructed using GESPAR, the proposed method, PR-GAMP, and L_1 -Fienup, for a range of measurement (M/N) and Gaussian noise SNRs, for $K = 5$ and measurements with 5 outliers. Computations (1000's of multiplies by A, A'): 106–497 (GESPAR), 94–215 (proposed), 107–266 (PR-GAMP), 106–218 (L_1 -Fienup).

Fig. 9.

The best reconstruction (with regularization parameter β) for the proposed method is compared against competing methods with the optimal (true) values of β_{sf} or K . These images are shown for the 512×512 -pixel star of David phantom, from $M = N/2$ measurements, with 60 dB AWGN noise and 1% (1311) outliers. The PR-GAMP method (not shown) converged to a blank (all zero) image. Computations (1000's of multiplies): 584 (GESPAR), 218 (proposed), 218 (PR-GAMP), 218 (L_1 -Fienup).

TABLE I

Comparison of Reconstruction Methods

Method	Implementation	Sparsity	Data Fit Term
L_1 -Fienup [31]	alternating projections	1-norm	quadratic (projection)
GESPAR [43]	greedy	0-“norm”	quadratic
PR-GAMP [30]	message-passing	0-“norm”	quadratic ²
Proposed	MM, (P)ADMM	1-norm	1-norm (ℓ_1)

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

²The Po-GAMP method is implemented for noise applied before taking the (squared)-magnitude, unlike the others here that assume noise is added to the (squared)-magnitude measurements.

TABLE II**Reconstruction Method Parameters**

All methods	$q = 2$, 50 Monte Carlo trials, 50 inits each
L ₁ -Fienup	50 iters/init, 5 conjugate gradient iterations for data projection (2D recon only), 10^{-4} stop tol
GESPAR	100 Gauss-Newton iters per GESPAR step, 10^{-5} stop tol, random measurement weights on
PR-GAMP	20 expectation-maximization iters, 200 inner iters each, 10^{-4} stop tol
Proposed	$\mu = 1$ start, $I_{\text{ADMM}} = 100$ (PADMM for $M < N$), adapt μ every 10 iters, $I_{\text{mm}} = 10$, 10^{-10} stop tol