Robust Phase Retrieval with Sparsity under Nonnegativity Constraints

Daniel S. Weller
Department of Electrical and Computer Engineering
University of Virginia
Charlottesville, Virginia 22904-4743
Email: d.s.weller@ieee.org

***Abstract*—High-resolution phase retrieval is challenging due to the low signal-to-noise ratio of measurements. This work utilizes variable splitting and alternating minimization to simultaneously enforce a 1-norm data fit penalty, an analysis-form sparse regularizer, and nonnegativity or real-valued image constraints to resolve an image from squared-magnitude measurements. The reconstruction algorithm incorporates real-valued and nonnegativity constraints via enforcing Hermitian symmetry on Fourier transform coefficients and projecting the reconstructed image onto the nonnegative orthant. The proposed method preserves image quality more robustly than unconstrained analysis-form sparse phase retrieval on the Shepp-Logan phantom, for a range of noise levels and outlier measurements.**

I. Introduction

Recent phase retrieval methods such as [1], [2], [3], [4] target signal recovery from magnitude transform measurements in environments with low signal-to-noise. This paper targets such large scale phase retrieval problems found in imaging, where in addition to being sparse or compressible in some domain, the image pixels also typically are nonnegative (e.g., intensities) or real-valued rather than being complex-valued quantities. Such constraints on the reconstructed signal values are already extensively used in classical phase retrieval algorithms [5]. Their presence in compressive model-based reconstruction algorithms also is not new [6], [7]. Such nonnegativity constraints would also be applicable to "robust" phase retrieval reconstructions, such as those incorporating sparsity and 1-norm-based data fit terms [3], [4], [8]. This paper adds nonnegativity constraints to the optimization problem and restricts the domain of the image to real-valued signals. The primal-dual optimization-based formulation based on the alternating directions method of multipliers (ADMM) [9], [10], [11], [12] used in [3], [8]. for finding a locally optimal solution. This approach is applied to large-scale 2D images to demonstrate the advantages of nonnegative-constrained optimization over robust phase retrieval without such constraints.

To set up the phase retrieval problem, consider an N-element signal $\mathbf{x}$, such as an image represented as a columnized vector of pixels. In applications of phase retrieval like telescope calibration (in astronomy), coherent diffraction imaging, and Fourier ptychography, a vector of M Fourier transform measurements $\mathbf{y}$ are acquired without their complex phase information. In the undersampled case, where $M < N$, the samples $\mathbf{y}$ correspond to the set $\{y_m : m \in \mathcal{M}\}$, where $|\mathcal{M}| = M$. In addition, physical measurements are corrupted by noise. Various types of noise are considered in the literature; this work follows the post-magnitude additive noise/outlier model used in [3]. Denoting the Fourier transform $\boldsymbol{\mathcal{F}}$, the discrete forward measurement model for the mth subsampled coefficient becomes

$$y_m = |[\boldsymbol{\mathcal{F}}\mathbf{x}]_m|^2 + \text{noise}, \quad m \in \mathcal{M}. \tag{1}$$

A nonnegative-constrained phase retrieval inverse problem is

$$\operatorname*{arg\,min}_{\mathbf{x} \geq 0} \sum_{m \in \mathcal{M}} \left| |[\boldsymbol{\mathcal{F}}\mathbf{x}]_m|^2 - y_m \right| + \beta R(\mathbf{x}), \tag{2}$$

where $\beta > 0$ and $R(\mathbf{x})$ are the parameter and penalty function for regularizing the solution. The nonnegativity constraint is written in the domain of the minimization problem. Implicit in this constraint is the real-valued domain of $\mathbf{x}$, which is also not considered in the author's earlier work [3], [8].

A number of regularizers made popular in the literature are motivated by prior information. Compressed sensing and phase retrieval are combined in numerous ways, including alternating projections [13], semidefinite programming [14], message passing [2], graph-based convex methods [15], and greedy algorithms [1]. The alternating minimization scheme used in [3] also facilitates sparse model-based regularization of the phase retrieval problem, and is particularly useful in conjunction with the 1-norm robust data-fit model. This work extends the alternating minimization for the primal-dual framework in [8] to real-valued signals and nonnegativity constraints, via additional variable splitting and direct enforcement of the real-valued constraint in the frequency domain.

II. Theory

The original inverse problem, as written in (2), is nonconvex and very difficult to solve, as all the elements of the vector $\mathbf{x}$ are involved in the penalty for every measurement. Introducing an auxiliary variable $\mathbf{u} = \boldsymbol{\mathcal{F}}\mathbf{x}$ can decouple the data-fit penalty. Furthermore, an additional auxiliary variable can be employed for analysis-form sparsity to simplify solution of the subproblems described in this section.

A. Analysis-Form Compressive Phase Retrieval

To begin, let us review the ADMM-based approach taken in [8], without nonnegativity or real-valued constraints on $\mathbf{x}$. In the real-valued analysis-form sparsity case, $R(\mathbf{x}) = \|\mathbf{C}\mathbf{x}\|_p^p$, with $p \le 1$. Substituting $\mathbf{u} = \boldsymbol{\mathcal{F}}\mathbf{x}$, and introducing the auxiliary vector $\mathbf{c} = \mathbf{C}\boldsymbol{\mathcal{F}}'\mathbf{u}$, the data-fit term for the mth measurement (in $\mathcal{M}$) is $|y_m - |u_m|^2|$, and the regularizer becomes $\|\mathbf{c}\|_p^p$, yielding the constrained optimization problem

$$\arg\min_{\mathbf{u},\mathbf{c}} \sum_{m\in\mathcal{M}} |y_m - |u_m|^2| + \beta\|\mathbf{c}\|_p^p, \quad \text{s.t.} \quad \mathbf{c} = \mathbf{C}\boldsymbol{\mathcal{F}}'\mathbf{u} \tag{3}$$

In many settings, including total variation, and undecimated wavelets, the real-valued analysis transform $\mathbf{C}$ can be represented as one or more circulant convolutions, meaning that $\mathbf{C}$ can be diagonalized as $\boldsymbol{\mathcal{F}}'\mathbf{D}_C\boldsymbol{\mathcal{F}}$[1]. Thus, $\mathbf{c} = \boldsymbol{\mathcal{F}}'\mathbf{D}_C\mathbf{u}$.

The Lagrange form of this contrained problem introduces the Lagrange dual vector $\mathbf{b}_1$. However, direct application of the Karush-Kuhn-Tucker (KKT) conditions only provides a necessary condition, not a sufficient one, for optimality in the nonconvex phase retrieval setting. Instead, following [8], the augmented Lagrangian form, with a quadratic penalty and parameter $\mu_1 > 0$, is used:

$$\begin{aligned}\mathcal{L}_A(\mathbf{u},\mathbf{c};\mathbf{b}_1) = &\sum_{m\in\mathcal{M}} |y_m - |u_m|^2| + \beta\|\mathbf{c}\|_p^p \\ &+ \tfrac{\mu_1}{2}\|\boldsymbol{\mathcal{F}}'\mathbf{D}_C\mathbf{u} - \mathbf{c} + \mathbf{b}_1/\mu\|^2.\end{aligned} \tag{4}$$

This technique involves solving for the primal vectors $\mathbf{u}^{(i)}$ and $\mathbf{c}^{(i)}$, and updates of the dual vector $\mathbf{b}_1^{(i)}$ as $\mathbf{b}_1^{(i)} \leftarrow \mathbf{b}_1^{(i-1)} + \mu_1(\mathbf{C}\boldsymbol{\mathcal{F}}'\mathbf{u}^{(i)} - \mathbf{c}^{(i-1)})$.

While approximate solutions are possible for $\mathbf{u}$ and $\mathbf{c}$, the ADMM method used instead alternates solving for each vector, holding the other variables fixed. The resulting subproblems take advantage of variable splitting to produce computationally simpler solutions than the joint optimization in (4).

$$\begin{aligned}\mathbf{u}^{(i)} = &\arg\min_{\mathbf{u}} \sum_{m=1}^{M} |y_m - |u_m|^2| \\ &+ \tfrac{\mu_1}{2}\|\mathbf{D}_C\mathbf{u} - \boldsymbol{\mathcal{F}}(\mathbf{c}^{(i-1)} - \tfrac{\mathbf{b}_1^{(i-1)}}{\mu_1})\|^2,\end{aligned} \tag{5}$$

$$\mathbf{c}^{(i)} = \arg\min_{\mathbf{c}} \beta\|\mathbf{c}\|_p^p + \tfrac{\mu_1}{2}\|\mathbf{c} - (\mathbf{C}\boldsymbol{\mathcal{F}}'\mathbf{u}^{(i)} + \tfrac{\mathbf{b}_1^{(i-1)}}{\mu_1})\|^2, \tag{6}$$

$$\mathbf{b}_1^{(i)} = \mathbf{b}_1^{(i-1)} + \mu_1(\mathbf{C}\boldsymbol{\mathcal{F}}'\mathbf{u}^{(i)} - \mathbf{c}^{(i)}). \tag{7}$$

In the $\mathbf{u}$-update, the unitary property of the DFT is used to simplify the problem, turning it into a separable problem in $\mathbf{u}$. The solution to this separable problem is described in [8]. The update for $\mathbf{c}$ is already separable in its usual form, which admits solution via a proximal operator such as soft-thresholding for the $p = 1$ case:

$$c_k^{(i)} = \text{soft}([\mathbf{C}\mathbf{x}^{(i)} + \tfrac{\boldsymbol{\lambda}^{(i-1)}}{\mu_2}]_k; \tfrac{\beta}{\mu_2}), \tag{8}$$

where the soft-thresholding operator is defined as

$$\text{soft}(c;\tau) = \tfrac{c}{|c|}\max\{0, |c| - \tau\}. \tag{9}$$

[1]The left inverse discrete Fourier transform is actually a set of multiple such transforms in the general case, and $\mathbf{D}_C$ is a stack of diagonal matrices.

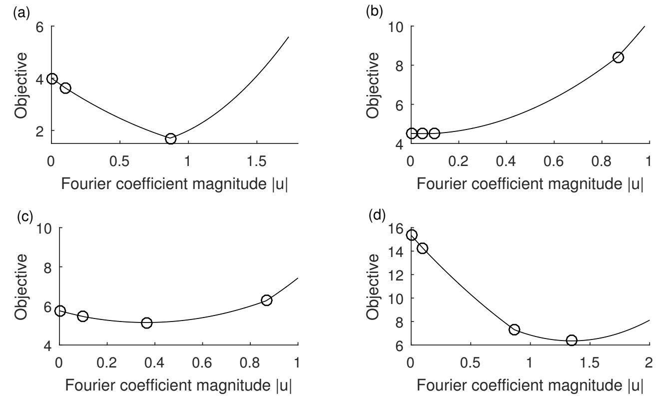

Fig. 1. Different cases for possible forms of the objective function being minimized in (12), for distinct y_m and $y_{[-m]}$: (a) only critical points 0, $\sqrt{y_m}$, and $\sqrt{y_{[-m]}}$ (circles), and additional critical points (b) smaller than the y's, (c) between the y's, and (d) greater than the y's. Note it is not possible for multiple additional critical points to occur simultaneously.

Closed forms also exist for the $p = 0$ (hard-thresholding) and other choices of p including $p = 1/2$.

B. Real-Valued Optimization and Nonnegativity

As mentioned in the introduction, this paper concerns solution of phase retrieval for nonnegative-valued images. This nonnegativity constraint is best enforced in conjunction with a real-valued constraint on the signal $\mathbf{x}$. Thus, this paper first describes modifying the analysis-form sparse formulation to enforce real-valued $\mathbf{x}$, and then introduces a nonnegativity constraint. The key observation to enforcing real-valued $\mathbf{x}$ is the one-to-one equivalency between $\mathbf{x}$ being real-valued and the Fourier transform $\mathbf{u}$ being Hermitian, or conjugate, symmetric. Introducing the notation $[-m]$ to indicate the Fourier transform coefficient index corresponding to "negative" frequency, this constraint means that $u_m = u_{[-m]}^*$, where $(\cdot)^*$ is the complex conjugate. Recalling the $\mathbf{u}$-update subproblem in (5) Let us consider two groups of frequencies: those where $m = [-m]$ (e.g., DC), and those where $m \neq [-m]$. Where the indexes are equal, our u_m-update becomes either

$$u_m^{(i)} = \arg\min_u |y_m - |u|^2| + \tfrac{\mu_1}{2}\|\mathbf{d}_m u - \mathbf{f}_m\|^2, \quad \text{or} \tag{10}$$

$$u_m^{(i)} = \arg\min_u \tfrac{\mu_1}{2}\|\mathbf{d}_m u - \mathbf{f}_m\|^2, \tag{11}$$

where $\mathbf{d}_m = [\mathbf{D}_C]_m$ and $\mathbf{f}_m = [\boldsymbol{\mathcal{F}}(\mathbf{c}^{(i-1)} - \frac{\mathbf{b}_1^{(i-1)}}{\mu_1})]_m$ are the vectors of values in each part of these matrices corresponding to the mth coefficient of $\mathbf{u}$. These are multi-element vectors in cases like 2D total variation where $\mathbf{C}$ corresponds to circulant convolutions with multiple filters.

In the case where the indexes are not equal, there are three cases: where both m and $[-m]$ are in $\mathcal{M}$, when just one is sampled (without loss of generality, call it m), or when neither is sampled. The updates for these three cases are

$$\begin{aligned}u_m^{(i)} = &\arg\min_u |y_m - |u|^2| + |y_{[-m]} - |u|^2| \\ &+ \tfrac{\mu_1}{2}\|\mathbf{d}_m u - \mathbf{f}_m\|^2 + \tfrac{\mu_1}{2}\|\mathbf{d}_{[-m]}u^* - \mathbf{f}_{[-m]}\|^2,\end{aligned} \tag{12}$$

$$u_m^{(i)} = \arg\min_u |y_m - |u|^2|$$

$$+\tfrac{\mu_1}{2}\|\mathbf{d}_m u-\mathbf{f}_m\|^2+\tfrac{\mu_1}{2}\|\mathbf{d}_{[-m]}u^*-\mathbf{f}_{[-m]}\|^2, \qquad (13)$$

and

$$u_m^{(i)}=\arg\min_u \|\mathbf{d}_m u-\mathbf{f}_m\|^2+\|\mathbf{d}_{[-m]}u^*-\mathbf{f}_{[-m]}\|^2. \quad (14)$$

For real-valued images, the analysis transform $\mathbf{c}$ and the dual vector $\mathbf{b}_1$ will also be real-valued, so their Fourier transform is also conjugate-symmetric; the same is true of the spectrum of the analysis transform spectrum coefficients $\mathbf{D}_C$. Thus, $\mathbf{d}_m=\mathbf{d}^*_{[-m]}$, and $\mathbf{f}_m=\mathbf{f}^*_{[-m]}$, and $\|\mathbf{d}_{[-m]}u^*-\mathbf{f}_{[-m]}\|=\|\mathbf{d}_m u-\mathbf{f}_m\|$. This observation simplifies the three last cases a bit.

Solving these single-variable u_m updates is fairly straightforward. The solution for the second and fifth cases are least-squares solutions $u_m^{(i)}=(\mathbf{d}'_m\mathbf{f}_m)/\|\mathbf{d}_m\|^2$. The minimizing angle $\angle u_m$ for all five cases is $\angle u_m=\angle(\mathbf{d}'_m\mathbf{f}_m)$. The remaining optimizations for the more complicated cases are over just the magnitudes. The optimization for the first case is described in [8]; $u_m^{(i)}$ is the minimizer of the four critical points 0, $\sqrt{y_m}$, $\frac{|\mu_1\mathbf{d}'_m\mathbf{f}_m|}{\mu_1\|\mathbf{d}_m\|^2-2}$, and $\frac{|\mu_1\mathbf{d}'_m\mathbf{f}_m|}{\mu_1\|\mathbf{d}_m\|^2+2}$, where the third critical point is only valid when both $\mu_1|\mathbf{d}_m|^2>2$, and $\mu_1|\mathbf{d}'_m\mathbf{f}_m|<\sqrt{y_m}(\mu_1\|\mathbf{d}_m\|^2-2)$, and the fourth critical point is only valid when $\mu_1|\mathbf{d}'_m\mathbf{f}_m|>\sqrt{y_m}(\mu_1\|\mathbf{d}_m\|^2+2)$.

The fourth case is very similar to the first case, but the four critical points are 0, $\sqrt{y_m}$, $\frac{|\mu_1\mathbf{d}'_m\mathbf{f}_m|}{\mu_1\|\mathbf{d}_m\|^2-1}$, and $\frac{|\mu_1\mathbf{d}'_m\mathbf{f}_m|}{\mu_1\|\mathbf{d}_m\|^2+1}$, where the third critical point is only valid when both $\mu_1|\mathbf{d}_m|^2>1$, and $\mu_1|\mathbf{d}'_m\mathbf{f}_m|<\sqrt{y_m}(\mu_1\|\mathbf{d}_m\|^2-1)$, and the fourth critical point is only valid when $\mu_1|\mathbf{d}'_m\mathbf{f}_m|>\sqrt{y_m}(\mu_1\|\mathbf{d}_m\|^2+1)$.

The fifth case is somewhat more complex (at least when $y_m\neq y_{[-m]}$). While the objective functions in the first and fourth cases essentially involve two ranges of $|u_m|$ with different quadratic functions for each, the fifth case has three ranges of $|u_m|$, each with a different quadratic function. Denote y_1 to be the smaller of y_m and $y_{[-m]}$, and y_2 to be the other. Then, the six critical points are 0, $\sqrt{y_m}$, $\sqrt{y_{[-m]}}$, $\frac{|\mu_1\mathbf{d}'_m\mathbf{f}_m|}{\mu_1\|\mathbf{d}_m\|^2-2}$, $\frac{|\mu_1\mathbf{d}'_m\mathbf{f}_m|}{\mu_1\|\mathbf{d}_m\|^2}$, and $\frac{|\mu_1\mathbf{d}'_m\mathbf{f}_m|}{\mu_1\|\mathbf{d}_m\|^2+2}$. The fourth point is valid only when both $\mu_1\|\mathbf{d}_m\|^2>2$ and $\mu_1|\mathbf{d}'_m\mathbf{f}_m|<\sqrt{y_1}(\mu_1\|\mathbf{d}_m\|^2-2)$. The fifth point is only valid when both $\mu_1\|\mathbf{d}_m\|^2>0$, and $\sqrt{y_1}(\mu_1\|\mathbf{d}_m\|^2)<\mu_1|\mathbf{d}'_m\mathbf{f}_m|<\sqrt{y_2}(\mu_1\|\mathbf{d}_m\|^2)$. The sixth critical point is valid when $\mu_1|\mathbf{d}'_m\mathbf{f}_m|>\sqrt{y_2}(\mu_1\|\mathbf{d}_m\|^2+2)$.

All told, the update to $\mathbf{u}$ can be performed efficiently by considering all the pairs of indexes $m,[-m]$, and identifying for each which of the five cases and optimization problems ((10)-(14)) to solve for each. Furthermore, all these computations can be performed in parallel.

Adding a nonnegative constraint can be performed as is described in [16]. In that work, $\mathbf{x}$ is treated as an additional auxiliary variable, and the ADMM framework is extended to include the additional subproblem

$$\mathbf{x}^{(i)}=\arg\min_{\mathbf{x}\geq 0}\|\mathbf{x}-(\boldsymbol{\mathcal{F}}'\mathbf{u}^{(i)}+\tfrac{\mathbf{b}_2^{(i-1)}}{\mu_2})\|^2, \qquad (15)$$

with additional dual vector $\mathbf{b}_2$ and AL penalty parameter $\mu_2>0$, for the constraint $\mathbf{x}=\boldsymbol{\mathcal{F}}'\mathbf{u}$. Also, the u_m-update would include the additional quadratic penalty term $\frac{\mu_2}{2}|u-[\boldsymbol{\mathcal{F}}(\mathbf{x}^{(i-1)}-\frac{\mathbf{b}_2^{(i-1)}}{\mu_2})]_m|^2$ (and for those with $m\neq[-m]$, the term corresponding to index $[-m]$). The second Lagrange dual vector update is

$$\mathbf{b}_2^{(i)}=\mathbf{b}_2^{(i-1)}+\mu_2(\boldsymbol{\mathcal{F}}'\mathbf{u}^{(i)}-\mathbf{x}^{(i)}). \qquad (16)$$

The $\mathbf{x}$-update is a closed-form projection onto the nonnegative orthant, which can be performed pixel-by-pixel:

$$x_n^{(i)}=\max\left\{0,[\boldsymbol{\mathcal{F}}'(\mathbf{u}^{(i)}-\tfrac{\mathbf{b}_2^{(i-1)}}{\mu_2})]_n\right\},\ \forall n. \qquad (17)$$

The updates to u_m described above are nearly identical to the real-valued case: the term $\mu_2[\boldsymbol{\mathcal{F}}(\mathbf{x}^{(i-1)}-\frac{\mathbf{b}_2^{(i-1)}}{\mu_2})]_m$ is added to every appearance of $\mu_1\mathbf{d}'_m\mathbf{f}_m$, and μ_2 is added to every appearance of $\mu_1\|\mathbf{d}_m\|^2$. Otherwise, the $\mathbf{u}$-update procedure is identical in the nonnegatively-constrained case.

An optional modification concerns the DC frequency component u_0: when $\mathbf{x}\geq 0$, the DC frequency component is also always nonnegative. Therefore, when initializing u_0 at the beginning of the iterative algorithm as described in [3], [8], $\angle u_0$ should not be randomized, but set to zero.

III. EXPERIMENTS

The 64×64-pixel Shepp-Logan phantom, as a prototypically sparse two-dimensional image, is evaluated here to evaluate the benefits of nonnegativity-constrained phase retrieval. Because the phantom produced by the MATLAB `phantom()` function is not purely nonnegative, having some pixels set to a small-in-magnitude, but negative value, those negative values are set to their absolute value to produce a truly nonnegative version of the phantom. The circulant total variation penalty is used, with periodic boundary conditions; in practice, when the object may extend beyond the field of view, the circulant total variation can be modified to avoid those periodic boundary values using a non-scalar regularization weighting parameter β that is equal to zero for just the periodic boundaries of $\mathbf{C}\mathbf{x}$. The $1,2$-norm is employed on the x- and y-directions, which are isotropically combined when computing the joint sparsity $\|\mathbf{C}\mathbf{x}\|_{1,2}=\sum_{r,c}\sqrt{(x_{r,c}-x_{r-1,c})^2+(x_{r,c}-x_{r,c-1})^2}$, summing over rows r and columns c. Thus, $p=1$, and the $\mathbf{c}$-update can be solved in closed-form using the vectorial extension of soft-thresholding.

The intended application of a "robust" phase retrieval algorithm like the proposed nonnegativity-constrained analysis-sparse reconstruction is the setting of high-resolution imaging in a setting with limited signal-to-noise ratio (SNR). Thus, the Shepp Logan phantom measurements will be simulated with varying levels of noise and/or outliers to explore the robustness of the algorithm. From the squared-magnitude measurements $|[\boldsymbol{\mathcal{F}}\mathbf{x}]_m|^2$, for each m, noise including additive white Gaussian noise (AWGN) and random uniform outliers are added to the squared magnitudes to form y_m. Since measured magnitudes are always nonzero, AWGN values that produce negative y_m are rounded up to zero. In the case of the Shepp Logan phantom, 60 dB SNR AWGN noise are added to all measurements, except for the noisefree reconstruction shown for reference. The outliers are randomly selected over the whole set of measurements $\mathcal{M}$, and each outlier is set

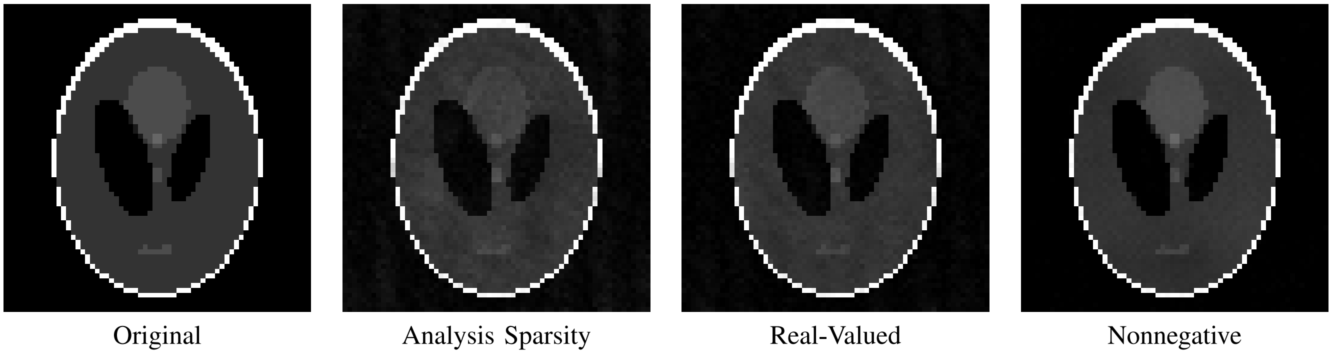

Fig. 2. The Shepp-Logan phantom (left) is measured in the Fourier domain, with 60 dB AWGN measurement noise and 1% outliers. The analysis (total variation) sparse reconstructions are shown on the right, where the first involves no constraints, the second constrains the result to be real-valued, and the third constrains the result to be nonnegative. The nonnegative constrained reconstruction yields the best result.

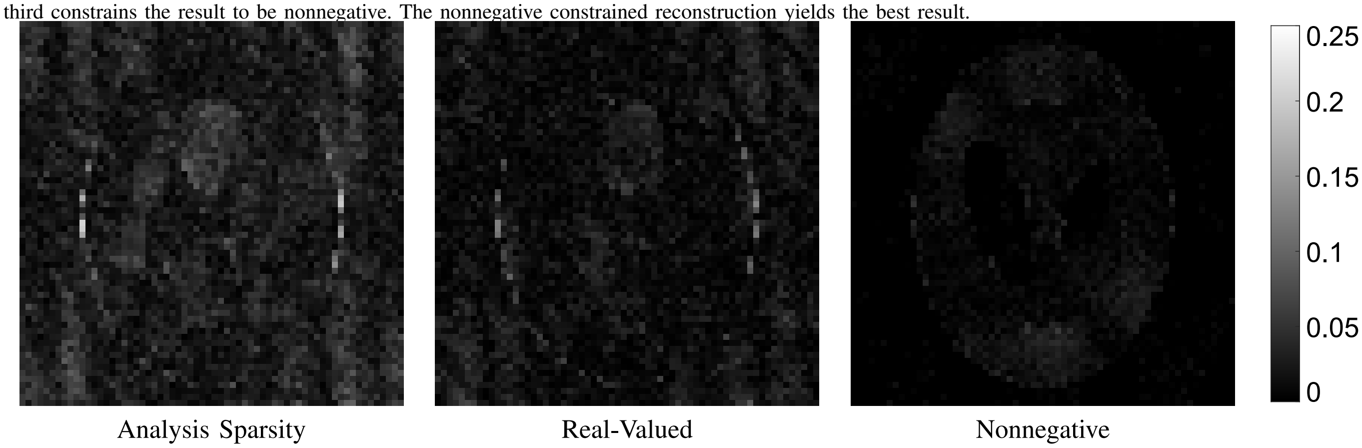

Fig. 3. The difference images for the reconstructions shown in Fig. 2 are shown above, on the scale of $[0, 0.25]$. The differences between these images are apparent; while the nonnegative-constrained result has small imperfections inside the "head", the others produce noticeable incoherent noise-like artifacts throughout the image.

to a random magnitude between 0 and twice the noisy (non-outlier) value. Different ranges of outliers are included, from 0 outliers (AWGN only) to 1 percent of all measurements (41 out of 4096 frequencies).

The phase retrieval reconstruction follows the approach described in [8], using the same β parameter set to 0.01, and using a choice of $\mu = 0.0625$ found to promote rapid convergence of the ADMM algorithm. Still to ensure incomplete convergence does not unduly influence reconstructed image quality, 1000 iterations of the ADMM algorithm are run for each initialization. Because the phase retrieval is nonconvex, 50 different initializations are used for both the proposed method and the unconstrained method described in [8]. Unlike in the existing method, the μ parameter is not regularly adjusted; in practice, self-adjustment via a heuristic like the comparison between primal and dual residuals used in [3], [8], [12]. Out of the 50 initializations, the best of the 50 reconstructions is determined automatically by choosing the one that minimizes the objective function in (2). When evaluating the error of this reconstruction versus the ground truth, the same measures are taken as in [8] to account for image flipping and shifting, and global phase changes, indeterminacies of phase retrieval that are not resolved by sparse modeling. Furthermore, the indeterminacy of the DC frequency measurement (since the DC measurement lies in the null space of the total variation operator $\mathbf{Cx}$) is resolved by ensuring that the DC frequency is included in $\mathcal{M}$. After resolving indeterminacies, the peak signal-to-noise ratio (PSNR)

$$\text{PSNR} = 10\log_{10}\frac{N}{\|\mathbf{x}-\mathbf{x}_{\text{true}}\|_2^2}.$$

and mean structural similarity image index (MSSIM) [17] values are reported for each reconstruction.

TABLE I
PSNR VALUES (IN DB) FOR SHEPP-LOGAN PHANTOM.

Noise Level	Analysis Sparsity [8]	Real-Valued	Nonnegative Constrained
None	111.3	138.3	147.0
AWGN (100 dB)	76.16	63.89	76.64
AWGN (80 dB)	53.96	50.36	58.50
AWGN (60 dB)	32.01	32.94	41.78
AWGN (60 dB) & 0.1% outliers	32.34	34.31	40.18
AWGN (60 dB) & 0.2% outliers	30.60	34.43	41.08
AWGN (60 dB) & 0.5% outliers	30.24	35.50	41.52
AWGN (60 dB) & 1% outliers	29.04	33.58	40.31

Tables I and II list the different noise levels and reconstruction PSNR and MSSIM values for each level, for the unconstrained analysis TV method, and the proposed method

TABLE II
MSSIM VALUES FOR SHEPP-LOGAN PHANTOM.

Noise Level	Analysis Sparsity [8]	Real-Valued	Nonnegative Constrained
None	1.000	1.000	1.000
AWGN (100 dB)	1.000	1.000	1.000
AWGN (80 dB)	0.996	0.989	1.000
AWGN (60 dB)	0.792	0.830	0.988
AWGN (60 dB) & 0.1% outliers	0.796	0.829	0.988
AWGN (60 dB) & 0.2% outliers	0.776	0.820	0.986
AWGN (60 dB) & 0.5% outliers	0.768	0.849	0.989
AWGN (60 dB) & 1% outliers	0.741	0.810	0.986

with real value and nonnegativity constraints. The Shepp-Logan phantom and reconstructed images are shown in Fig. 2 for the case of 60 dB AWGN noise and 1% measurement outliers. The difference images in Fig. 3, relative to the ground truth, illustrate the qualitative differences between these techniques. While the methods perform very similarly in the high-SNR regime, the nonnegative-constrained reconstruction is noticeably more robust to outliers than the unconstrained or real-value constrained methods. The MSSIM values remain over 0.98 for the nonnegative-constrained images, while the other reconstructions have significantly reduced MSSIM values, signifying the loss of meaningful structure in the reconstructed images. This loss of structure can be traced back visually to the structure present in the difference images in the figure, where one can clearly see the outline of the phantom is more distinct in the unconstrained analysis sparsity and real-valued reconstructions, and much less so in the nonnegative constrained reconstruction.

IV. DISCUSSION

The reconstructions of the Shepp-Logan phantom provide an example of the benefit of nonnegativity constraints when performing phase retrieval on real images. While introducing another auxiliary variable and modifying the optimization to include it does not significantly alter the computation involved, it does yield significantly improved images when outliers are present. The reasons for this improvement are likely twofold, but identifying the dominant factor requires further analysis. One, a nonnegativity constraint further improves upon the ill-posedness of the conventional phase retrieval problem. Secondly, the constraint stabilizes the optimization somewhat and reduces the effects of initialization, since the DC frequency, which contains most of the image energy in many cases, would have a fixed phase.

Further development of this and related phase retrieval methods for realistic application to large scale imaging will require innovations in a number of areas. First, the total variation model used is convenient for ensuring $\mathbf{C}$ is composed of circulant transforms, and it is sufficient for the Shepp-Logan phantom, but it likely would be less suitable for real images acquired by a microscope or other imaging system, where edges may be more diffuse. Computationally, iterative methods like this algorithm would scale to large images, but the memory and processing requirements would still increase substantially. Also, tuning parameters such as β, μ_1, and μ_2 need to be selected in an automatic but robust fashion, taking into account image content complexity and noise level. These issues must be taken into account when considering whether to apply this or any other iterative phase retrieval algorithm.

ACKNOWLEDGMENT

The author would like to thank Marios Pattichis for inviting this work to be presented at this conference, and the conference attendees for their valuable feedback.

REFERENCES

[1] Y. Shechtman, A. Beck, and Y. C. Eldar, "GESPAR: efficient phase retrieval of sparse signals," *IEEE Trans. Sig. Proc.*, vol. 62, no. 4, pp. 928–38, Feb. 2014.

[2] P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate message passing," in *Proc. 50th Allerton Conf. on Comm., Control, and Computing*, 2012, pp. 815–822.

[3] D. S. Weller, A. Pnueli, G. Divon, O. Radzyner, Y. C. Eldar, and J. A. Fessler, "Undersampled phase retrieval with outliers," *IEEE Trans. Comput. Imag.*, vol. 1, no. 4, pp. 247–258, Dec. 2015.

[4] P. Hand, "PhaseLift is robust to a constant fraction of arbitrary errors," *Applied and Computational Harmonic Analysis*, 2016, in press.

[5] J. R. Fienup, "Phase retrieval algorithms: a comparison," *Appl. Optics*, vol. 21, no. 15, pp. 2758–69, Aug. 1982.

[6] D. L. Donoho and J. Tanner, "Sparse nonnegative solution of underdetermined linear equations by linear programming," *Proc. Natl. Acad. Sci.*, vol. 102, no. 27, pp. 9446–9451, Jul. 2005.

[7] A. M. Bruckstein, M. Elad, and M. Zibulevsky, "On the uniqueness of nonnegative sparse solutions to underdetermined systems of equations," *IEEE Trans. Info. Theory*, vol. 54, no. 11, pp. 4813–4820, Nov. 2008.

[8] D. S. Weller, "Analysis-form sparse phase retrieval using variable-splitting," in *IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI)*, Mar. 2016, pp. 61–64.

[9] R. Glowinski and A. Marrocco, "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*, vol. 9, no. R2, pp. 41–76, 1975.

[10] D. Gabay and B. Mercier, "A dual algorithm for the solution of nonlinear variational problems via finite-element approximations," *Comput. Math. Appl.*, vol. 2, no. 1, pp. 17–40, 1976.

[11] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, no. 1-3, pp. 293–318, Apr. 1992.

[12] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. & Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2010.

[13] S. Mukherjee and C. S. Seelamantula, "Fienup algorithm with sparsity constraints: Application to frequency-domain optical-coherence tomography," *IEEE Trans. Sig. Proc.*, vol. 62, no. 18, pp. 4659–72, Sep. 2014.

[14] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM J. Imaging Sci.*, vol. 6, no. 1, pp. 199–225, 2013.

[15] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," *Mathematical Programming*, pp. 1–35, Dec. 2013.

[16] S. Ramani and J. A. Fessler, "A splitting-based iterative algorithm for accelerated statistical X-ray CT reconstruction," *IEEE Trans. Med. Imag.*, vol. 31, no. 3, pp. 677–88, Mar. 2012.

[17] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Im. Proc.*, vol. 13, no. 4, pp. 600–12, Apr. 2004.