PHASESPLIT: A VARIABLE SPLITTING FRAMEWORK FOR PHASE RETRIEVAL

Subhadip Mukherjee, Suprosanna Shit, and Chandra Sekhar Seelamantula

Department of Electrical Engineering, Indian Institute of Science, Bangalore 560012, India
Emails: {subhadipm, suprosannas}@iisc.ac.in, chandra.sekhar@ieee.org

ABSTRACT

We develop two techniques based on *alternating minimization* and *alternating directions method of multipliers* for phase retrieval (PR) by employing a variable-splitting approach in a maximum likelihood estimation framework. This leads to an additional equality constraint, which is incorporated in the optimization framework using a quadratic penalty. Both algorithms are iterative, wherein the updates are computed in closed-form. Experimental results show that: (i) the proposed techniques converge faster than the state-of-the-art PR algorithms; (ii) the complexity is comparable to the state of the art; and (iii) the performance does not depend critically on the choice of the penalty parameter. We also show how sparsity can be incorporated within the variable splitting framework and demonstrate concrete applications to image reconstruction in *frequency-domain optical-coherence tomography*.

***Index Terms*—** Phase retrieval, Sparsity, ADMM, Alternating Minimization, FDOCT

1. INTRODUCTION

Phase retrieval (PR) is a quadratic inverse problem that arises in several imaging applications, such as X-ray crystallography [1], holography [2], microscopy [3], etc. The diffraction pattern of the object in X-ray crystallography is its Fourier transform, and the sensors record only its intensity. The phase contains structural information about the object, but it is not directly available. Thus, it becomes imperative to recover the phase, starting from the intensity measurements. The objective of PR is to solve this ill-posed inverse problem by taking into account priors about the object, such as non-negativity, compact support, sparsity, etc., or by acquiring oversampled measurements. More generally, PR refers to the problem of reconstruction from quadratic measurements that do not necessarily correspond to the Fourier intensity.

The early contributions go by the name of Fienup iterations and Gerchberg-Saxton error reduction algorithms [4–8], which bounce back-and-forth between the object and the measurement domains, and incorporate respective priors. There are also non-iterative techniques, which rely on a Hilbert integral relation between the log-magnitude and phase of the Fourier transform of minimum-phase signals [9]. The two-dimensional counterpart and exact reconstruction guarantee was proposed in [10] for digital holography. Recently, we developed model-based non-iterative PR techniques [11–13].

The problem of PR addressed within the realm of sparsity [14, 15] received attention because of its wide applicability: Signals encountered in a number of applications admit a sparse representation in a suitably chosen basis. A greedy local search-based algorithm for sparse PR, referred to as GESPAR, was proposed by Schechtman et al. [18]. We developed the *sparse Fienup algorithm* [16, 17], which is an adaptation of the classical Fienup algorithm to take sparsity into account. Netrapalli et al. [19] developed an analytical convergence guarantee for the alternating minimization framework (Alt. Min.) for PR subject to *spectral initialization*, with and without sparsity, referred to as *AltMinPR* and *SparseAltMinPR*, respectively. Vaswani et al. [20] recently developed an Alt. Min. technique for low-rank matrix recovery from quadratic measurements of the columns. Other notable sparse PR techniques include dictionary learning (DOLPHIn) [21], cone programming [22], generalized message passing [23], majorization-minimization approach [24], etc. Fogel et al. showed that sparsity and positivity priors lead to faster convergence [25].

Candès et al. [26, 27] pioneered the *PhaseLift* framework, wherein one lifts a vector $\boldsymbol{x}$ to a matrix, thereby linearizing the quadratic measurements. Reconstruction is achieved by solving a tractable semi-definite program (SDP). One could consider two possibilities in this framework: (i) oversampled measurements and no signal prior; or (ii) incorporation of a signal prior such as sparsity [28]. Schechtman et al. [29] employed PhaseLift and imposed sparsity via the *log-det* penalty, and demonstrated applications to sub-wavelength imaging with partially incoherent light. Gradient-descent approaches for PR without *lifting* include the Wirtinger Flow (WF) method [30] and its truncated version (TWF) [31]. These algorithms possess better scalability than lifting and have convergence guarantees subject to the spectral initialization [19]. Waldspurger et al. proposed *PhaseCut* [32], where the PR problem is posed as a non-convex quadratic program and is solved using block-coordinate-descent with a per-iteration complexity comparable to that of the Gerchberg-Saxton-type algorithms.

Our contribution: We propose a variable-splitting approach for PR, referred to as PhaseSplit, wherein we express the quadratic measurements in one vector as equivalent bilinear measurements in two vectors. We develop two algorithms to solve the resulting constrained optimization, one based on Alt. Min. and the other based on *alternating directions method of multipliers* (ADMM). The Alt. Min. flavor turns out to be a fixed-point iteration, with the ground-truth as the fixed-point in the absence of measurement noise. Simulations show that the proposed algorithms result in a convergence behavior superior to three state-of-the-art PR techniques, and achieve a reconstruction accuracy up to machine precision when measurement noise is absent. One could incorporate sparsity in the PhaseSplit formalism – we demonstrate an application of this for signal reconstruction in *frequency-domain optical-coherence tomography* (FDOCT).

The first and second authors have contributed equally. This work was funded by the Robert-Bosch Center for Cyber-Physical Systems.

2. THE PHASESPLIT ALGORITHMS

We formulate the PR problem both with and without the sparsity prior in a variable-splitting framework and develop two reconstruction algorithms, one based on Alt. Min. and the other on ADMM.

The objective is to estimate a signal $\boldsymbol{x}^* \in \mathbb{R}^n$ from noisy quadratic measurements of the form $y_i = \left|\boldsymbol{a}_i^\top \boldsymbol{x}^*\right|^2 + \xi_i,\ i = 1:m$, where $\{\boldsymbol{a}_i\} \in \mathbb{R}^n$ are known sampling vectors, which may be random (Gaussian vectors, for instance) or deterministic (Fourier bases); $\{\xi_i\}$ denotes additive noise with $\xi_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}\left(0, \sigma_\xi^2\right)$; and $i = 1 : m$ is a shorthand notation for $i = 1, 2, \cdots, m$. The maximum-likelihood estimate of $\boldsymbol{x}^*$ is given by

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \frac{1}{2} \sum_{i=1}^{m} \left(y_i - \left|\boldsymbol{a}_i^\top \boldsymbol{x}\right|^2 \right)^2. \quad (1)$$

The principal idea in PhaseSplit is to express the quadratic $\left|\boldsymbol{a}_i^\top \boldsymbol{x}\right|^2$ as a bilinear form $\boldsymbol{u}^\top \boldsymbol{A}_i \boldsymbol{v}$ in two variables $\boldsymbol{u}$ and $\boldsymbol{v}$, where $\boldsymbol{A}_i = \boldsymbol{a}_i \boldsymbol{a}_i^\top$, and recast (1) in the following equivalent form:

$$(\hat{\boldsymbol{u}}, \hat{\boldsymbol{v}}) = \arg\min_{\boldsymbol{u},\boldsymbol{v}} \frac{1}{2} \sum_{i=1}^{m} \left(y_i - \boldsymbol{u}^\top \boldsymbol{A}_i \boldsymbol{v} \right)^2 \text{ s.t. } \boldsymbol{u} = \boldsymbol{v}. \quad (2)$$

Clearly, (1) and (2) have identical solutions, i.e., $\hat{\boldsymbol{u}} = \hat{\boldsymbol{v}} = \hat{\boldsymbol{x}}$. Both (1) and (2) are nonconvex optimization problems. The motivation for coming up with the formulation in (2) is that the cost function is a *convex quadratic* in $\boldsymbol{u}$ for a fixed $\boldsymbol{v}$ and vice versa — this makes it an ideal candidate for Alt.Min. type approaches, as shown later. We refer to (2) as *PhaseSplit*.

2.1. PhaseSplit-Alt. Min. (PS-AM) Algorithm

The equality constraint in (2) can be incorporated in the cost function using a quadratic penalty as follows:

$$(\hat{\boldsymbol{u}}, \hat{\boldsymbol{v}}) = \arg\min_{\boldsymbol{u},\boldsymbol{v}} \frac{1}{2} \sum_{i=1}^{m} \left(y_i - \boldsymbol{u}^\top \boldsymbol{A}_i \boldsymbol{v} \right)^2 + \frac{\lambda}{2} \|\boldsymbol{u} - \boldsymbol{v}\|_2^2, \quad (3)$$

where $\lambda > 0$. If λ becomes arbitrarily large, the solution to (3) would tend to the solution to (2). The cost function in (3) is a convex quadratic in $\boldsymbol{u}$ for a fixed $\boldsymbol{v}$ and vice versa, which motivates the use of an Alt.Min. strategy. The cost in (3), when viewed as a function of $\boldsymbol{u}$ for a fixed $\boldsymbol{v}$, can be expressed as

$$q(\boldsymbol{u}) = \frac{1}{2} \boldsymbol{u}^\top \boldsymbol{C}_{\boldsymbol{v}} \boldsymbol{u} - \boldsymbol{d}_{\boldsymbol{v}}{}^\top \boldsymbol{u} + (\text{terms independent of } \boldsymbol{u}),$$

where $\boldsymbol{C}_{\boldsymbol{v}}$ and $\boldsymbol{d}_{\boldsymbol{v}}$ are defined as follows:

$$\boldsymbol{C}_{\boldsymbol{v}} = \lambda \boldsymbol{I} + \sum_{i=1}^{m} \left(\boldsymbol{a}_i^\top \boldsymbol{v} \right)^2 \boldsymbol{A}_i, \quad \boldsymbol{d}_{\boldsymbol{v}} = \lambda \boldsymbol{v} + \sum_{i=1}^{m} \left(\boldsymbol{a}_i^\top \boldsymbol{v} \right) y_i \boldsymbol{a}_i. \quad (4)$$

The optimal $\boldsymbol{u}$ is then given by $\boldsymbol{u}_{\text{opt}} = \boldsymbol{C}_{\boldsymbol{v}}^{-1} \boldsymbol{d}_{\boldsymbol{v}}$. Considering the fact that the formulation is symmetric in $\boldsymbol{u}$ and $\boldsymbol{v}$, we have $\boldsymbol{v}_{\text{opt}} = \boldsymbol{C}_{\boldsymbol{u}}^{-1} \boldsymbol{d}_{\boldsymbol{u}}$, when the cost is optimized w.r.to $\boldsymbol{v}$ for a fixed $\boldsymbol{u}$. These two steps lead to an Alt.Min. strategy with the following updates:

$$\boldsymbol{u}^{t+1} = \boldsymbol{C}_{\boldsymbol{v}^t}^{-1} \boldsymbol{d}_{\boldsymbol{v}^t} \text{ and } \boldsymbol{v}^{t+1} = \boldsymbol{C}_{\boldsymbol{u}^{t+1}}^{-1} \boldsymbol{d}_{\boldsymbol{u}^{t+1}}. \quad (5)$$

The initialization for $\boldsymbol{v}^0$ in case of random sampling vectors can be made following the spectral initialization strategy developed by Netrapalli et al. [19]. This algorithm results in an alternating sequence of estimates of $\boldsymbol{x}^*$ of the form $\{\boldsymbol{u}^1, \boldsymbol{v}^1, \boldsymbol{u}^2, \boldsymbol{v}^2, \cdots\}$. The matrix inversion in (5) can be efficiently performed using the conjugate-gradient (CG) method, using the current estimate $\boldsymbol{v}^t$ as the initialization in the CG method to obtain $\boldsymbol{u}^{t+1}$ and so on. As the sequence of estimates stabilize, progressively less number of CG iterations would be needed to compute the updates, leading to a substantial reduction in computation. The steps of the PhaseSplit Alt. Min. scheme are listed in Algorithm 1.

Algorithm 1 : PhaseSplit-Alt. Min. algorithm for phase retrieval.

1. Input: Measurements $\{y_i\}_{i=1}^m$, the sampling vectors $\{\boldsymbol{a}_i\}_{i=1}^m$, maximum number of iterations N_{iter}, the desired number of non-zeros s in the estimate if the target signal is sparse, and λ.

1. Initialization: Set $t = 0$ and $\boldsymbol{v}^t = \boldsymbol{v}_{\max}$, the eigenvector corresponding to the largest eigenvalue of $\boldsymbol{S} = \sum_{i=1}^m y_i \boldsymbol{a}_i \boldsymbol{a}_i^\top$.

2. For $t = 1 : N_{\text{iter}}$ **do**:

1. $\boldsymbol{u}^t = \boldsymbol{C}_{\boldsymbol{v}^{t-1}}^{-1} \boldsymbol{d}_{\boldsymbol{v}^{t-1}}$, where $\boldsymbol{C}_{\boldsymbol{v}}$ and $\boldsymbol{d}_{\boldsymbol{v}}$ are as in (4),
2. $\boldsymbol{v}^t = \boldsymbol{C}_{\boldsymbol{u}^t}^{-1} \boldsymbol{d}_{\boldsymbol{u}^t}$, and

3. Output: the current estimates $\boldsymbol{u}^t$ or $\boldsymbol{v}^t$.

2.1.1. *A Fixed-Point Interpretation of PhaseSplit-Alt. Min.*

The symmetry (in $\boldsymbol{u}$ and $\boldsymbol{v}$) of the update rules in (5) leads to a more insightful interpretation. Redefining $\boldsymbol{v}^t = \boldsymbol{x}^{2t}$ and $\boldsymbol{u}^{t+1} = \boldsymbol{x}^{2t+1}$, for $t = 0, 1, 2, \cdots$, the update rules in (5) can be expressed succinctly as $\boldsymbol{x}^{t+1} = \boldsymbol{C}_{\boldsymbol{x}^t}^{-1} \boldsymbol{d}_{\boldsymbol{x}^t}$. In the absence of noise, we have $y_i = \left|\boldsymbol{a}_i^\top \boldsymbol{x}^*\right|^2$, leading to $\boldsymbol{d}_{\boldsymbol{u}} = (\lambda \boldsymbol{I} + \sum_{i=1}^m y_i \boldsymbol{A}_i)\boldsymbol{u} = \boldsymbol{C}_{\boldsymbol{x}^*} \boldsymbol{u}$. Consequently, the update rule becomes $\boldsymbol{x}^{t+1} = \boldsymbol{C}_{\boldsymbol{x}^t}^{-1} \boldsymbol{C}_{\boldsymbol{x}^*} \boldsymbol{x}^t$, which can be interpreted as a fixed-point iteration of the form $\boldsymbol{x}^{t+1} = h\left(\boldsymbol{x}^t\right)$, where $h(\boldsymbol{x}) = \boldsymbol{C}_{\boldsymbol{x}}^{-1} \boldsymbol{C}_{\boldsymbol{x}^*} \boldsymbol{x}$, which has the desirable property that $\boldsymbol{x}^*$ is a fixed-point of $h(\boldsymbol{x})$. This guarantees that if $\boldsymbol{x}^t = \boldsymbol{x}^*$ for some t, all subsequent updates produce $\boldsymbol{x}^*$ as the estimate. Therefore, the PhaseSplit-Alt.Min. algorithm can be interpreted as a fixed-point iterative algorithm having $\boldsymbol{x}^*$ as its fixed-point in the absence of noise.

2.2. PhaseSplit-ADMM (PS-ADMM) Algorithm

The augmented Lagrangian function defined as

$$\mathcal{L}(\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{\mu}) = \frac{1}{2} \sum_{i=1}^{m} \left[y_i - \boldsymbol{u}^\top \boldsymbol{A}_i \boldsymbol{v} \right]^2 + \frac{\lambda}{2} \|\boldsymbol{u} - \boldsymbol{v}\|_2^2 + \boldsymbol{\mu}^\top [\boldsymbol{u} - \boldsymbol{v}],$$

where $\boldsymbol{\mu}$ is the Lagrange multiplier corresponding to the equality constraint in (2), lies at the center of the development of the ADMM algorithm for solving (2). The resulting algorithm consists of the following update rules [33]:
$\boldsymbol{u}^{t+1} = \arg\min_{\boldsymbol{u}} \mathcal{L}\left(\boldsymbol{u}, \boldsymbol{v}^t, \boldsymbol{\mu}^t\right)$, $\boldsymbol{v}^{t+1} = \arg\min_{\boldsymbol{v}} \mathcal{L}\left(\boldsymbol{u}^{t+1}, \boldsymbol{v}, \boldsymbol{\mu}^t\right)$, and $\boldsymbol{\mu}^{t+1} = \boldsymbol{\mu}^t + \lambda\left(\boldsymbol{u}^{t+1} - \boldsymbol{v}^{t+1}\right)$, beginning with appropriate initializations $\boldsymbol{v}^0$ and $\boldsymbol{\mu}^0$. The updates for $\boldsymbol{u}$ and $\boldsymbol{v}$ can be computed in closed-form. The steps of the PhaseSplit-ADMM procedure are listed in Algorithm 2.

2.3. Sparse PhaseSplit Algorithms

The PhaseSplit formulation can be adapted to take into account the case where $\boldsymbol{x}^*$ is at most s-sparse in a dictionary $\boldsymbol{\Psi} \in \mathbb{R}^{n \times n}$, leading

Algorithm 2 : PhaseSplit-ADMM for phase retrieval.

1. Input: Measurements $\{y_i\}_{i=1}^m$, the sampling vectors $\{\boldsymbol{a}_i\}_{i=1}^m$, maximum number of iterations N_{iter}, and λ.

1. Initialization: Set $t = 0$, initialize $\boldsymbol{v}^t$ and $\boldsymbol{\mu}^t$.

2. For $t = 1 : N_{\text{iter}}$ **do**

1. $\boldsymbol{u}^{t+1} = \boldsymbol{C}_{\boldsymbol{v}^t}^{-1}\left(\boldsymbol{d}_{\boldsymbol{v}^t} - \boldsymbol{\mu}^t\right)$,
2. $\boldsymbol{v}^{t+1} = \boldsymbol{C}_{\boldsymbol{u}^{t+1}}^{-1}\left(\boldsymbol{d}_{\boldsymbol{u}^{t+1}} + \boldsymbol{\mu}^t\right)$, and
3. $\boldsymbol{\mu}^{t+1} = \boldsymbol{\mu}^t + \lambda\left(\boldsymbol{u}^{t+1} - \boldsymbol{v}^{t+1}\right)$.

3. Output: the current estimate $\boldsymbol{u}^t$ or $\boldsymbol{v}^t$.

to $\boldsymbol{x}^* = \boldsymbol{\Psi}\boldsymbol{\alpha}^*$, where $\|\boldsymbol{\alpha}^*\|_0 \leq s \ll n$. Redefining $\boldsymbol{a}_i$ as $\boldsymbol{\Psi}^\top \boldsymbol{a}_i$, the signal reconstruction problem can be expressed as

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \frac{1}{2}\sum_{i=1}^m \left(y_i - \left|\boldsymbol{a}_i^\top \boldsymbol{x}\right|^2\right)^2 \text{ s.t. } \|\boldsymbol{x}\|_0 \leq s. \quad (6)$$

The sparse counterpart of (2), which we shall refer to as *Sparse PhaseSplit*, can then be formulated as follows:

$$\min_{\boldsymbol{u},\boldsymbol{v}} \frac{1}{2}\sum_{i=1}^m \left(y_i - \boldsymbol{u}^\top \boldsymbol{A}_i \boldsymbol{v}\right)^2 \text{ s.t. } \boldsymbol{u} = \boldsymbol{v} \text{ and } \|\boldsymbol{u}\|_0 \leq s. \quad (7)$$

The sparse counterparts of PS-AM and PS-ADMM can be readily obtained by inserting the additional step of obtaining the best s-sparse approximation of the updates in every iteration given by an operator $\mathcal{P}_s$: $\boldsymbol{v}^t \leftarrow \mathcal{P}_s\left(\boldsymbol{v}^t\right)$.

3. SIMULATION RESULTS

The performance of PS-AM and PS-ADMM without the sparsity prior are validated on synthetic signals and compared with three state-of-the-art PR algorithms: (i) PhaseLift, (ii) TWF, and (iii) AltMinPR. We evaluate the reconstruction performance for both noiseless and noisy measurements. The effects of the regularization parameter λ and the oversampling factor $\frac{m}{n}$ are also studied. Since the underlying signal is real, the accuracy of recovery is measured in terms of the global sign-independent reconstruction signal-to-noise ratio (SNR), defined as SNR $= \max\limits_{\alpha\in\{-1,+1\}} \frac{\|\boldsymbol{x}^*\|_2^2}{\|\alpha\,\hat{\boldsymbol{x}} - \boldsymbol{x}^*\|_2^2}$, where $\hat{\boldsymbol{x}}$ is an estimate of the ground-truth $\boldsymbol{x}^*$.

First, we consider the task of reconstructing a signal $\boldsymbol{x}^*$ of dimension $n = 64$, drawn uniformly at random from the surface of the unit sphere. The sampling vectors $\boldsymbol{a}_i$ are drawn independently from the $\mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$ distribution. The performance evaluation is carried out for four different values of $\frac{m}{n}$ and the results are shown in Figure 1. We observe that for all values of $\frac{m}{n}$, PS-AM exhibits the fastest convergence, followed by PS-ADMM. As $\frac{m}{n}$ increases, the number of iterations consumed by all algorithms for convergence reduces. After sufficiently many iterations, the reconstruction SNR reaches approximately 300 dB, indicating that the ground-truth is recovered to machine precision. The convergence speed of PS-AM and PS-ADMM is higher than the competing algorithms for smaller values of $\frac{m}{n}$. For higher oversampling factors, the improvement in convergence speed over the state-of-the-art diminishes.

Subsequently, we proceed to examine the effect of measurement noise on reconstruction accuracy. The oversampling factor is kept fixed at $\frac{m}{n} = 6$ and experiments are conducted for four different values of the input SNR, defined as SNR$_{\text{in}} = \frac{1}{m\sigma_\xi^2}\sum_{i=1}^m \left|\boldsymbol{a}_i^\top \boldsymbol{x}^*\right|^4$, where σ_ξ^2 is the noise variance. The reconstruction SNR values are averaged over 20 independent noise realizations. We observe from Figure 2 that PS-AM and PS-ADMM continue to converge faster than the competing algorithms even in the presence of noise. The reconstruction SNR varies from 22 to 35 dB, depending on the input SNR level. PS-AM and PS-ADMM take about 10 iterations to converge to a fairly high value of reconstruction SNR, whereas their best performing competitor PhaseLift requires approximately 30 to 40 iterations. Since PS-AM and PS-ADMM have the same asymptotic complexity per iteration ($\mathcal{O}(n^3)$) as PhaseLift, faster convergence leads to a substantial reduction in overall computations.

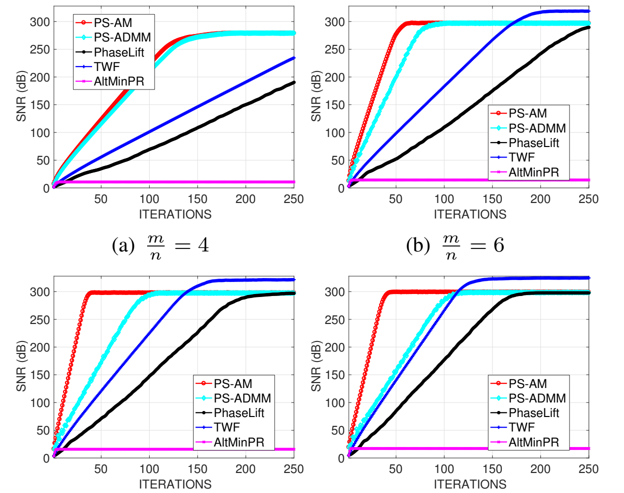

(a) $\frac{m}{n} = 4$ (b) $\frac{m}{n} = 6$ (c) $\frac{m}{n} = 8$ (d) $\frac{m}{n} = 10$

Fig. 1. (Color online) Effect of the oversampling factor $\frac{m}{n}$ on the reconstruction. The measurements are noise-free. The PhaseSplit algorithms attain a reconstruction SNR value close to 100 dB or higher within 20 to 30 iterations.

The reconstruction SNR, plotted as a function of λ in Figure 3(a) indicates that the performance of PS-AM and PS-ADMM does not depend critically on the choice of λ. The SNR of reconstruction is above 100 dB, which corresponds to near-accurate estimation, over a fairly large range of values of λ (nearly 10^2 to 10^4). The average run-times per-iteration for different algorithms shown in Figure 3(b) with respect to the signal dimension n indicate that PS-AM and PS-ADMM have better scalability than PhaseLift as n grows. The average run-times of the proposed algorithms are in the same range as that of AltMinPR. Scalability with respect to the signal dimension turns out to be the best for TWF.

The performance of sparse PS-AM is compared with GESPAR and SparseAltMinPR (cf. Figure 4). The total number of swaps in GESPAR is taken as 1000. Both SparseAltMinPR and sparse PS-AM are iterated 100 times. The reconstruction SNR values, averaged over 20 independent trials, are reported corresponding to $m = 4n$ and $\lambda = 50$, for four different values of the sparsity level $\rho = \frac{s}{n}$. The input SNR value is taken to be 25 dB. The support indices of the ground-truth are drawn uniformly at random. We observe that sparse PS-AM far outperforms SparseAltMinPR and yields competitive reconstruction performance as compared with GESPAR.

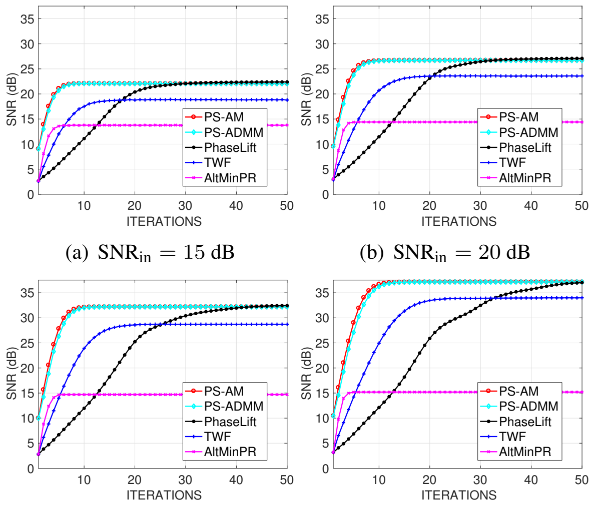

(a) $\text{SNR}_{\text{in}} = 15$ dB (b) $\text{SNR}_{\text{in}} = 20$ dB

(c) $\text{SNR}_{\text{in}} = 25$ dB (d) $\text{SNR}_{\text{in}} = 30$ dB

Fig. 2. (Color online) Effect of measurement noise corresponding to an oversampling rate of $\frac{m}{n} = 6$. The output SNR values are averaged over 20 independent trials.

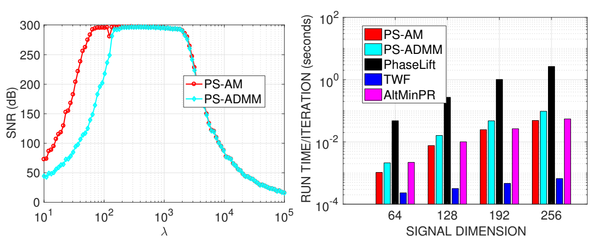

(a) Reconstruction SNR versus λ (b) Run-time versus n

Fig. 3. (Color online) Effect of λ on the reconstruction SNR, and the comparison of run-times for $\frac{m}{n} = 6$.

4. SPARSE PHASESPLIT AND APPLICATION TO FDOCT

We consider signal reconstruction in FDOCT as an application of *Sparse PhaseSplit*. FDOCT is a non-invasive imaging technique used to obtain structural details of biological specimens. The Fourier intensity of the interference between the reference and object-arm signals is recorded by the spectrometer. The phase has to be estimated starting from the intensity measurements in order to reconstruct the axial backscattering function — we achieve this using *Sparse PhaseSplit*. Since the reflected wave exhibits a strong peak only when there is a significant change of refractive index in the specimen, the assumption of sparsity is indeed apt.

Since the sampling vectors $\{\boldsymbol{a}_i\}_{i=1}^m$ corresponding to Fourier modulus measurements are complex, we write

$$y_i = \left|\boldsymbol{a}_{i\text{Re}}^\top \boldsymbol{x}^*\right|^2 + \left|\boldsymbol{a}_{i\text{Im}}^\top \boldsymbol{x}^*\right|^2 + \xi_i = \boldsymbol{x}^{*\top} \boldsymbol{A}_i \boldsymbol{x}^* + \xi_i,$$

where $\boldsymbol{A}_i = \boldsymbol{a}_{i\text{Re}}\boldsymbol{a}_{i\text{Re}}^\top + \boldsymbol{a}_{i\text{Im}}\boldsymbol{a}_{i\text{Im}}^\top$, where Re and Im indicate the real and imaginary parts, respectively. For complex $\boldsymbol{a}_i$, one has to derive again the algorithm starting from (7). The derivation will be reported separately. The performance of PhaseSplit and Sparse PhaseSplit is shown in Figure 5 for reconstruction of two specimens, namely glass and onion-peel[1]. For benchmarking, we show the reconstruction obtained using the *standard Fourier inversion* technique, which corresponds to a zero-phase inversion. We observe that sparsity helps eliminate the so-called *autocorrelation artifacts* (which are typical of the standard inversion method) and also gives better signal-to-background noise contrast. In the case of the glass specimen, the reconstruction obtained using Sparse PhaseSplit indicates three clear interfaces, and the background noise is significantly suppressed.

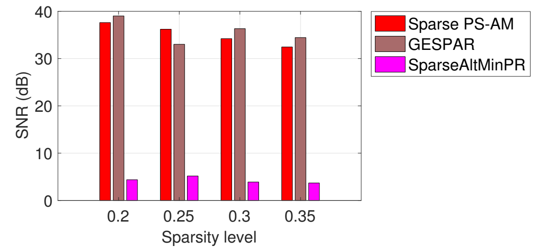

Fig. 4. Sparse PhaseSplit vis-à-vis other sparse PR algorithms for different sparsity levels. Sparse PS-AM performs on par with GESPAR for denser signals and slightly better for sparse signals.

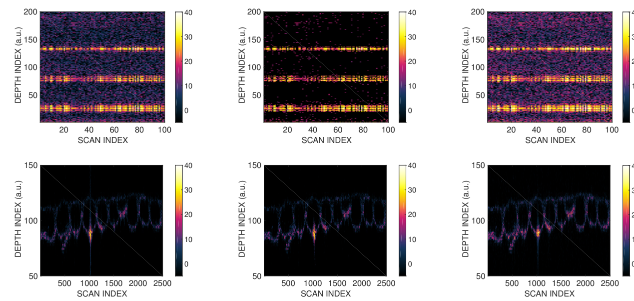

PS-AM Sparse PS-AM Direct inversion

Fig. 5. (Color online) The reconstruction of glass (first row) and onion-peel (second row) specimens using phasesplit, sparse phasesplit, and the direct Fourier inversion technique. The sparsity prior enables suppression of coherent artifacts and background noise.

5. CONCLUSIONS

We formulated the problem of phase retrieval in a variable-splitting framework and developed two algorithms, one based on Alt.Min, and the other based on ADMM. The advantage of the variable-splitting approach is that it converts a quadratic expression in one vector into an equivalent bilinear form in two vectors. It is possible to incorporate sparsity into PhaseSplit by applying a hard-thresholding operator that prunes insignificant entries to produce a sparse estimate. The matrix inversion in the update steps is done efficiently by employing the CG technique. Experimental comparisons show that the proposed algorithms converge significantly faster than the state-of-the-art PR algorithms both in the presence and absence of measurement noise. We also demonstrated successful application of sparse PhaseSplit to perform signal reconstruction in FDOCT, which establishes the practical impact of the proposed framework.

[1]The glass and onion-peel data are courtesy of Prof. Rainer A. Leitgeb, Medical University of Vienna, Austria.

6. REFERENCES

[1] R. P. Millane, "Phase retrieval in crystallography and optics," *J. Opt. Soc. Amer. A*, vol. 7, no. 3, pp. 394–411, Mar. 1990.

[2] A. Szoke, "Holographic microscopy with a complicated reference," *J. Imag. Sci. Technol.*, vol. 41, pp. 332–341, 1997.

[3] A. J. J. Drenth, A. Huiser, and H. Ferwerda, "The problem of phase retrieval in light and electron microscopy of strong objects," *Optica Acta*, vol. 22, pp. 615–628, 1975.

[4] J. R. Fienup, "Phase retrieval algorithms: A comparison," *Appl. Opt.*, vol. 21, pp. 2758–2769, 1982.

[5] J. R. Fienup, "Phase retrieval algorithms: A personal tour," *Appl. Opt.*, vol. 52, pp. 45–56, 2013.

[6] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik*, vol. 35, pp. 237–246, 1972.

[7] H. H. Bauschke, P. L. Combettes, and D. Luke, "Phase retrieval, error reduction algorithm, and Fienup variants: A view from convex optimization," *J. Opt. Soc. Amer. A*, vol. 19, pp. 1334–1345, 2002.

[8] T. F. Quatieri and A. V. Oppenheim, "Iterative techniques for minimum-phase signal reconstruction from phase or magnitude," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, pp. 1187–1193, 1981.

[9] B.Yegnanarayana and A. Dhayalan, "Noniterative techniques for minimum phase signal reconstruction from phase or magnitude," in *Proc. IEEE Intl. Conf. Acoust., Speech, and Signal Process.*, vol. 8, pp. 639–642, Apr. 1983.

[10] C. S. Seelamantula, N. Pavillon, C. Depeursinge, and M. Unser, "Exact complex-wave reconstruction in digital holography," *J. Opt. Soc. Amer. A*, vol. 28, no. 6, pp. 983–992, Jun. 2011.

[11] B. A. Shenoy, S. Mukherjee, and C. S. Seelamantula, "Phase retrieval for a class of 2-D signals characterized by first-order difference equations," in *Proc. IEEE Intl. Conf. on Image Process.*, pp. 325–329, 2013.

[12] B. A. Shenoy and C. S. Seelamantula, "Exact phase retrieval for a class of 2-D parametric signals," *IEEE Trans. Signal Process.*, vol. 63, no. 1, pp. 90–103, 2015.

[13] B. A. Shenoy, S. Mulleti, and C. S. Seelamantula, "Exact phase retrieval in principal shift-invariant spaces," *IEEE Trans. Signal Process.*, vol. 64, no. 2, pp. 406–416, 2016.

[14] M. L. Moravec, J. K. Romberg, and R. G. Baraniuk, "Compressive phase retrieval," in *Proc. Wavelets XII SPIE Int. Symp. Opt. Sci. Technol.*, Aug. 2007.

[15] Y. M. Yu and M. Vetterli, "Sparse spectral factorization: Unicity and reconstruction algorithms," in *Proc. IEEE Intl. Conf. Acoust. Speech, Signal Process.*, pp. 5976–5979, 2011.

[16] S. Mukherjee and C. S. Seelamantula, "An iterative algorithm for phase retrieval with sparsity constraints: Application to frequency domain optical coherence tomography," in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Process.*, pp. 553-556, Mar. 2012.

[17] S. Mukherjee and C. S. Seelamantula, "Fienup algorithm with sparsity constraints: Application to frequency-domain optical-coherence tomography," *IEEE Trans. Signal Process.*, vol. 62, no. 18, pp. 4659–4672, Sep. 2014.

[18] Y. Shechtman, A. Beck, and Y. C. Eldar, "GESPAR: Efficient phase retrieval of sparse signals," *IEEE Trans. Signal Process.*, vol. 62, no. 4, pp. 928–938, Feb. 2014.

[19] P. Netrapalli, P. Jain, and S. Sanghavi, "Phase retrieval using alternating minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 18, pp. 4814–4826, Sep. 2015.

[20] N. Vaswani, S. Nayer, and Y. C. Eldar, "Low rank phase retrieval," *IEEE Trans. Signal Process.*, vol. 65, no. 15, pp. 4059–4074, Aug. 2017.

[21] A. M. Tillmann, Y. C. Eldar, and J. Mairal, "DOLPHIn–Dictionary learning for phase retrieval," *arXiv:1602.02263v1*, Feb. 2016.

[22] S. Becker, E. J. Candès, and M. Grant, "Templates for convex cone problems with applications to sparse signal recovery," *Technical Report*, Department of Statistics, Stanford University, 2010.

[23] P. Schniter and S. Rangan, "Compressive phase retrieval via generalized approximate massage passing," *IEEE Trans. Signal Process.*, vol. 63, no. 4, pp. 1043–1055, Feb. 2015.

[24] T. Qiu and D. P. Palomar, "Undersampled sparse phase retrieval via majorization-minimization," *IEEE Trans. Signal Process.*, vol. 65, no. 22, pp. 5957–5969, Nov. 2017.

[25] F. Fogel, I. Waldspurger, and A. d'Aspremont, "Phase retrieval for imaging problems," *Math. Prog. Comp.*, vol. 8, issue 3, pp. 311–335, Sep. 2016.

[26] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Comm. on Pure and Appl. Math.*, vol. 66, issue 8, pp. 1241–1274, Aug. 2013.

[27] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM J. Imag. Sci.*, vol. 6, issue 1, pp. 199–224, Feb. 2013.

[28] H. Ohlsson, A. Y. Yang, R. Dong, and S. S. Sastry, "Compressive phase retrieval from squared output measurements via semidefinite programming," *arXiv:1111.6323v3*, Mar. 2012.

[29] Y. Shechtman, Y. C. Eldar, A. Szameit, and M. Segev, "Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing," *Opt. Exp.*, vol. 19, no. 16, pp. 14807–14822, 2011.

[30] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval via Wirtinger flow: Theory and algorithms," *IEEE Trans. Info. Theory*, vol. 61, no. 4, pp. 1985–2007, Apr. 2015.

[31] Y. Chen and E. J. Candès, "Solving random quadratic systems of equations is nearly as easy as solving linear systems," in *Proc. Advances in Neural Info. Process. Systems*, 2015.

[32] I. Waldspurger, A. dAspremont, and S. Mallat, "Phase recovery, maxcut and complex semidefinite programming," *Math. Program.*, vol. 149, no. 1, pp. 47–81, 2015.

[33] Y. Wang, W. Yin, and J. Zeng, "Global convergence of ADMM in nonconvex nonsmooth optimization," *arXiv:1511.06324v4*, Nov. 2016.