Signal Reconstruction from Modulo Observations

Viraj Shah and Chinmay Hegde

Electrical and Computer Engineering Department

Iowa State University

Ames, IA, USA 50010

Abstract—We consider the problem of reconstructing a signal from under-determined *modulo* observations (or measurements). This observation model is inspired by a (relatively) less well-known imaging mechanism called modulo imaging, which can be used to extend the dynamic range of imaging systems; variations of this model have also been studied under the category of *phase unwrapping*. Signal reconstruction in the under-determined regime with modulo observations is a challenging ill-posed problem, and existing reconstruction methods cannot be used directly. In this paper, we propose a novel approach to solving the inverse problem limited to two modulo periods, inspired by recent advances in algorithms for *phase retrieval under sparsity constraints*. We show that given a sufficient number of measurements, our algorithm perfectly recovers the underlying signal and provides improved performance over other existing algorithms. We also provide experiments validating our approach on both synthetic and real data to depict its superior performance.

I. INTRODUCTION

A. Motivation

The problem of reconstructing a signal (or image) from (possibly) nonlinear observations is widely encountered in standard signal acquisition and imaging systems. Our focus in this paper is the problem of signal reconstruction from *modulo* measurements, where the modulo operation with respect to a positive real valued parameter R returns the (fractional) remainder after division by R. See Fig. 1 for an illustration.

Formally, we consider a high dimensional signal (or image) $\mathbf{x}^* \in \mathbb{R}^n$. We are given modulo measurements of $\mathbf{x}^*$, that is, for each measurement vector $\mathbf{a_i} \in \mathbb{R}^n$, we observe:

$$y_i = \mod (\langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle, R) \quad i = \{1, 2, ..., m\}, \tag{1}$$

The task is to recover $\mathbf{x}^*$ using the modulo measurements $\mathbf{y}$ and knowledge the measurement matrix $\mathbf{A} = [\mathbf{a_1} \ \mathbf{a_2} \ ... \ \mathbf{a_m}]^\top$.

This specific form of signal recovery is gaining rapid interest in recent times. Recently, the use of a novel imaging sensor that wraps the data in a periodical manner has been shown to overcome certain hardware limitations of typical imaging systems [1]–[4]. Many image acquisition systems suffer from the problem of limited dynamic range; however, real-world signals can contain a large range of intensity levels, and if tuned incorrectly, most intensity levels can lie in the saturation region of the sensors, causing loss of information through signal clipping. The problem gets amplified in the case of multiplexed
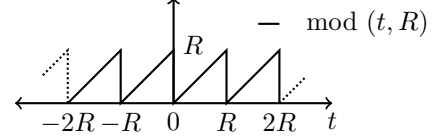
Fig. 1: *The modulo transfer function.*

linear imaging systems (such as compressive cameras or coded aperture systems), where required dynamic range is very high because of the fact that each linear measurement is a weighted aggregation of the original image intensity values.

The standard solution to this issue is to improve sensor dynamic range via enhanced hardware; this, of course, can be expensive. An intriguing alternative is to deploy special digital *modulo* sensors [5]–[8]. As the name suggests, such a sensor wraps each signal measurement around a scalar parameter R that reflects the dynamic range. However, this also makes the forward model (1) highly nonlinear and the reconstruction problem highly ill-posed. The approach of [1], [2] resolves this problem by assuming *overcomplete* observations, meaning that the number of measurements m is higher than the ambient dimension n of the signal itself. For the cases where m and n are large, this requirement puts a heavy burden on computation and storage.

In contrast, our focus is on solving the the inverse problem (1) with very few number of samples, i.e., we are interested in the case $m \ll n$. While this makes the problem even more ill-posed, we show that such a barrier can be avoided if we assume that the underlying signal obeys a certain low-dimensional structure. In this paper, we focus on the *sparsity* assumption on the underlying signal, but our techniques could be extended to other signal structures. Further, for simplicity, we assume that our forward model is limited to only two modulo periods, as shown in the Fig. 2(a). Such a simplified variation of the modulo function already inherits much of the challenging aspects of the original recovery problem. Intuitively, this simplification requires that the value of dynamic range parameter R should be large enough so that all the measurements $\langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle$ can be covered within the domain of operation of the modulo function, *i.e.*, $\langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle \in [-R, R] \ \forall i \in \{1, 2, .., m\}$.

B. Our contributions

In this paper, we propose a recovery algorithm for exact reconstruction of signals from modulo measurements of the form (1). We refer our algorithm as *MoRAM*, short for *Modulo Recovery using Alternating Minimization*. The key idea in our

approach is to identify and draw parallels between modulo recovery and the problem of *phase retrieval*. Indeed, this connection enables us to bring in algorithmic ideas from classical phase retrieval, which also helps in our analysis.

Phase retrieval has its roots in several classical imaging problems, but has attracted renewed interest of late. There, we are given observations of the form:

$$y_i = |\langle \mathbf{a_i}, \mathbf{x}^* \rangle|, \quad i = 1, 2, \ldots, m,$$

and are tasked with reconstructing $\mathbf{x}^*$. While these two different class of problems appear different at face value, the common theme is the need of undoing the effect of a piecewise linear transfer function applied to the observations. See Fig. 2 for a comparison. Both the functions are identical to the identity function in the positive half, but differ significantly in the negative half. Solving the phase retrieval problem is essentially equivalent to retrieving the phase ($\text{sign}(y_i)$) corresponding to each measurement y_i. However, the phase can take only two values: 1 if $t \geq 0$, or -1 if $t < 0$. Along the same lines, for modulo recovery case, the challenge is to identify the bin-index for each measurement. Estimating the bin-index correctly lets us "unravel" the modulo transfer function, thereby enabling signal recovery.
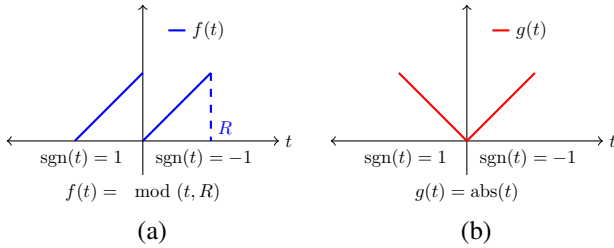

Fig. 2: *Comparison between (a) modulo function ($f(t) = \text{mod}(t, R)$); and (b) absolute value function ($g(t) = \text{abs}(t)$).*

At the same time, several essential differences between the two problems restrict us from using phase retrieval algorithms as-is for the modulo reconstruction problem. The absolute value function can be represented as a *multiplicative* transfer function (with the multiplying factors being the signs of the linear measurements), while the modulo function adds a constant value (R) to negative inputs. Therefore, the estimation procedures propagate very differently in the two cases. In the case of phase retrieval, a wrongly estimated phase induces an error that increases *linearly* with the magnitude of each measurement. On the other hand, for modulo recovery problem, the error induced by an incorrect bin-index is R (or larger), irrespective of the measurement. Therefore, existing algorithms for phase retrieval perform rather poorly for our problem (both in theory and practice).

We resolve this issue by making non-trivial modifications to existing phase retrieval algorithms that better exploit the structure of modulo reconstruction. We also provide analytical proofs for recovering the underlying signal using our algorithm, and show that such a recovery can be performed using an (essentially) optimal number of observations, provided certain standard assumptions are met. To the best of our knowledge we

are the first to pursue this type of approach for modulo recovery problems with *generic* linear measurements, distinguishing us from previous work [1], [2].

C. Techniques

The basic approach in our proposed (MoRAM) algorithm is similar to several recent non-convex phase retrieval approaches. We pursue two stages.

In the first stage, we identify a good initial estimated signal $\mathbf{x^0}$ that that lies (relatively) close to the true signal $\mathbf{x}^*$. A commonly used initialization technique for phase retrieval is *spectral initialization* as described in [9]. However, that does not seem to succeed in our case, due to markedly different behavior of the modulo transfer function. Instead, we introduce a novel approach of measurement *correction* by comparing our observed measurements with typical density plots of Gaussian observations. Given access to such corrected measurements, $\mathbf{x^0}$ can be calculated simply by using a first-order estimator. This method is intuitive, yet provides a provable guarantee for getting an initial vector that is close to the true signal.

In the second stage, we refine this coarse initial estimate to recover the true underlying signal. Again, we follow an alternating-minimization (AltMin) approach inspired from phase retrieval algorithms (such as [9]) that estimates the signal and the measurement bin-indices alternatively. However, as mentioned above, any estimation errors incurred in the first step induces fairly large additive errors (proportional to the dynamic range parameter R.) We resolve this issue by using a *robust* form of alternating-minimization (specifically, the Justice Pursuit algorithm [10]). We prove that AltMin, based on Justice Pursuit, succeeds provided the number of wrongly estimated bin-indices in the beginning is a small fraction of the total number of measurements. This gives us a natural radius for initialization, and also leads to provable sample-complexity upper bounds.

D. Paper organization

The reminder of this paper is organized as follows. In Section II, we briefly discuss the prior work. Section III contains notation and mathematical model used for our analysis. In Section IV, we introduce the MoRAM algorithm and provide a theoretical analysis of its performance. We demonstrate the performance of our algorithm by providing series of numerical experiments in Section VI. Section VII provides concluding remarks.

II. PRIOR WORK

We now provide a brief overview of related prior work. At a high level, our algorithmic development follows two (hitherto disconnected) streams of work in the signal processing literature.

Phase retrieval: As stated earlier, in this paper we borrow algorithmic ideas from previously proposed solutions for phase retrieval to solve the modulo recovery problem. Being a classical problem with a variety of applications, phase retrieval has been studied significantly in past few years. Approaches to

solve this problem can be broadly classified into two categories: convex and non-convex.

Convex approaches usually consist of solving a constrained optimization problem after lifting the true signal $\mathbf{x}^*$ in higher dimensional space. The seminal PhaseLift formulation [11] and its variations [12], [13] come under this category. Recently, [14], [15] proposed a convex optimization based algorithm for solving the phase retrieval problem that doesn't lift the underlying signal and hence does not square the number of variables. The precise analysis of it has been derived in [16], [17]. Typical non-convex approaches involve finding a good initialization, followed by iterative minimization of a loss function. Approaches based on Wirtinger Flow [18]–[21] and Amplitude flow [22], [23] come under this category.

In recent works, extending phase retrieval algorithms to situations where the underlying signal exhibits a sparse representation in some known basis has attracted interest. Convex approaches for sparse phase retrieval include [24]–[27]. Similarly, non-convex approaches for sparse phase retrieval include [9], [21], [22]. Our approach in this paper towards solving the modulo recovery problem can be viewed as a complement to the non-convex sparse phase retrieval framework advocated in [28].

Modulo recovery: The modulo recovery problem is also known in the classical signal processing literature as phase unwrapping. The algorithm proposed in [29] is specialized to images, and employs graph cuts for phase unwrapping from a single modulo measurement per pixel. However, the inherent assumption there is that the input image has very few sharp discontinuities, and this makes it unsuitable for practical situations with textured images. Our work is motivated by the recent work of [2] on high dynamic range (HDR) imaging using a modulo camera sensor. For image reconstruction using multiple measurements, they propose the multi-shot UHDR recovery algorithm, with follow-ups developed further in [30]. However, the multi-shot approach depends on carefully designed camera exposures, while our approach succeeds for non-designed (generic) linear observations; moreover, they do not include sparsity in their model reconstructions. In our previous work [3], we proposed a different extension based on [2], [31] for signal recovery from quantized modulo measurements, which can also be adapted for sparse measurements, but there too the measurements need to be carefully designed.

In the literature, several authors have attempted to theoretically understand the modulo recovery problem. Given modulo-transformed time-domain samples of a band-limited function, [1] provides a stable algorithm for perfect recovery of the signal and also proves sufficiency conditions that guarantees the perfect recovery. [4] formulates and solves an QCQP problem with non-convex constraints for denoising the modulo-1 samples of the unknown function along with providing a least-square based modulo recovery algorithm.. However, both these methods relay on the smoothness of the band-limited function as a prior structure on the signal, and as such it is unclear how to extend their use to more complex modeling priors (such as sparsity in a given basis).

In recent works, [32] proposed unlimited sampling algorithm for sparse signals. Similar to [1], it also exploits the bandlimit-edness by considering the low-pass filtered version of the sparse signal, and thus differs from our random measurements setup. In [33], modulo recovery from Gaussian random measurements is considered, however, it assumes the true signal to be distributed as mixed Bernoulli-Gaussian distribution, which is impractical in real world imaging scenarios.

For a qualitative comparison of our MoRAM method with existing approaches, refer Table I. The table suggests that the previous approaches varied from the Nyquist-Shannon sampling setup only along the amplitude dimension, as they rely on band-limitedness of the signal and uniform sampling grid. We vary the sampling setup along both the amplitude and time dimensions by incorporating sparsity in our model, which enables us to work with non-uniform sampling grid (random measurements) and achieve a provable sub-Nyquist sample complexity.

III. PRELIMINARIES

A. Notation

Let us introduce some notation. We denote matrices using bold capital-case letters ($\mathbf{A}, \mathbf{B}$), column vectors using bold-small case letters ($\mathbf{x}, \mathbf{y}, \mathbf{z}$ etc.) and scalars using non-bold letters (R, m etc.). We use letters C and c to represent constants that are large enough and small enough respectively. We use $\mathbf{x}^\top, \mathbf{A}^\top$ to denote the transpose of the vector $\mathbf{x}$ and matrix $\mathbf{A}$ respectively. The cardinality of set S is denoted by $\mathrm{card}(S)$. We define the signum function as $\mathrm{sgn}(x) := \frac{x}{|x|}$ for every $x \in \mathbb{R}, x \neq 0$, with the convention that $\mathrm{sgn}(0) = 1$. The i^{th} element of the vector $\mathbf{x} \in \mathbb{R}^n$ is denoted by x_i. Similarly, i^{th} row of the matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is denoted by $\mathbf{a_i}$, while the element of $\mathbf{A}$ in the i^{th} row and j^{th} column is denoted as a_{ij}. The projection of $\mathbf{x} \in \mathbb{R}^n$ onto a set of coordinates S is represented as $\mathbf{x}_S \in \mathbb{R}^n$, i.e., $\mathbf{x}_{S_j} = \mathbf{x}_j$ for $j \in S$, and 0 elsewhere.

B. Mathematical model

As depicted in Fig. 2(a), we consider the modulo operation within 2 periods (one in the positive half and one in the negative half). We assume that the value of dynamic range parameter R is large enough so that all the measurements $\langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle$ are covered within the domain of operation of modulo function. Rewriting in terms of the signum function, the (variation of) modulo function under consideration can be defined as:

$$f(t) := t + \left(\frac{1 - \mathrm{sgn}(t)}{2} \right) R.$$

One can easily notice that the modulo operation in this case is nothing but an addition of scalar R if the input is negative, while the non-negative inputs remain unaffected by it. If we divide the number line in these two bins, then the coefficient of R in above equation can be seen as a bin-index, a binary variable which takes value 0 when $\mathrm{sgn}(t) = 1$, or 1 when $\mathrm{sgn}(t) = -1$. Inserting the definition of f in the measurement model of Eq. 1 gives,

$$y_i = \langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle + \left(\frac{1 - \mathrm{sgn}(\langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle)}{2} \right) R, \quad i = \{1,..,m\}. \quad (2)$$

TABLE I: Comparison of MoRAM with existing modulo recovery methods.

	Unlimited Sampling [1]	OLS Method [4]	multishot UHDR [2]	**MoRAM (our approach)**
Assumption on structure of signal	Bandlimited	Bandlimited	No assumptions	Sparsity
Sampling scheme	uniform grid	uniform grid	(carefully chosen) linear measurements	random linear measurements
Sample complexity	oversampled, $\mathcal{O}(n)$	–	oversampled, $\mathcal{O}(n)$	undersampled, $\mathcal{O}(s \log(n))$
Provides sample complexity bounds?	Yes	–	No	Yes
Leverages Sparsity?	No	No	No	Yes
(Theoretical) bound on dynamic range	Unbounded	Unbounded	Unbounded	$2R$

We can rewrite Eq. 2 using a bin-index vector $\mathbf{p} \in \{0,1\}^m$. Each element of the true bin-index vector $\mathbf{p}^*$ is given as,

$$p_i^* = \frac{1 - \text{sgn}(\langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle)}{2}, \quad i = \{1, .., m\}.$$

If we ignore the presence of modulo operation in above formulation, then it reduces to a standard compressive sensing reconstruction problem. In that case, the compressed measurements y_{c_i} would just be equal to $\langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle$.

While we have access only to the compressed modulo measurements $\mathbf{y}$, it is useful to write $\mathbf{y}$ in terms of true compressed measurements $\mathbf{y}_c$. Thus,

$$y_i = \langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle + p_i^* R = y_{c_i} + p_i^* R.$$

It is evident that if we can recover $\mathbf{p}^*$ successfully, we can calculate the true compressed measurements $\langle \mathbf{a_i} \cdot \mathbf{x}^* \rangle$ and use them to reconstruct $\mathbf{x}^*$ with any sparse recovery algorithm such as CoSaMP [34] or basis-pursuit [35]–[37].

IV. SPARSE SIGNAL RECOVERY

Of course, the major challenge is that we do *not* know the bin-index vector. In this section, we describe our algorithm to recover both $\mathbf{x}^*$ and $\mathbf{p}^*$, given $\mathbf{y}, \mathbf{A}, s, R$. Our algorithm *MoRAM (Modulo Reconstruction with Alternating Minimization)* comprises of two stages: (i) an initialization stage, and (ii) descent stage via alternating minimization.

A. Initialization by re-calculating the measurements

Similar to other non-convex approaches, MoRAM also requires an initial estimate $\mathbf{x}^0$ that is close to the true signal $\mathbf{x}^*$. We have several initialization techniques available; in phase retrieval, techniques such as spectral initialization are often used. However, the nature of the problem in our case is fundamentally different due to the non-linear *additive* behavior of the modulo transfer function. To overcome this issue, we propose a method to re-calculate the true Gaussian measurements ($\mathbf{y}_c = \mathbf{A}\mathbf{x}^*$) from the available modulo measurements.

The high level idea is to undo the nonlinear effect of modulo operation in a significant fraction of the total available measurements. To understand the method for such re-calculation, we will first try to understand the effect of modulo operation on the linear measurements.
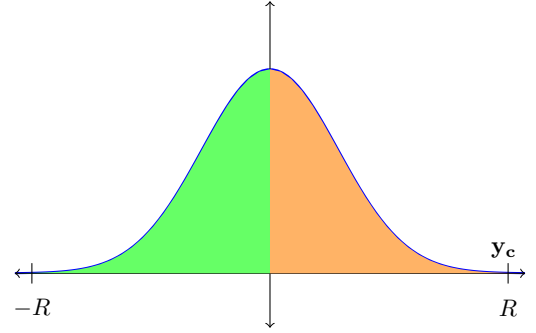
Fig. 3: *Density plot of* $\mathbf{A}\mathbf{x}^*$

1) Effect of the modulo transfer function: To provide some intuition, let us first examine the distribution of the $\mathbf{A}\mathbf{x}^*$(Fig. 3) and $\text{mod}(\mathbf{A}\mathbf{x}^*)$ (shown in Fig. 4) to understand what information can be obtained from the modulo measurements. We are particularly interested in the case where the elements of $\mathbf{A}\mathbf{x}^*$ are small compared to the modulo range parameter R.

Note that the compressed measurements $\mathbf{y}_c$ follow the standard normal distribution, as $\mathbf{A}$ is Gaussian random matrix. These plots essentially depict the distribution of our observations *before* and *after* the modulo operation.

With reference to Fig. 3, we divide the compressed observations $\mathbf{y}_c$ in two sets: $\mathbf{y}_{c,+}$ contains all the non-negative observations (orange) with bin-index$= 0$, while $\mathbf{y}_{c,-}$ contains all the negative ones (green) with bin-index$= 1$.

As shown in Fig. 4, after modulo operation, the set $\mathbf{y}_{c,-}$ (green) shifts to the right by R and gets concentrated in the right half ($[R/2, R]$); while the set $\mathbf{y}_{c,+}$ (orange) remains unaffected and concentrated in the left half ($[0, R/2]$). Thus, for some of the modulo measurements, their correct bin-index can be identified just by observing their magnitudes relative to the midpoint $R/2$. This leads us to obtain following maximum likelihood estimator for bin-indices ($\mathbf{p}$):

$$p_i^{init} = \begin{cases} 0, & \text{if } 0 \leq y_i < R/2 \\ 1, & \text{if } R/2 \leq y_i \leq R \end{cases} \tag{3}$$

The $\mathbf{p}^{\mathbf{init}}$ obtained with above method contains the correct values of bin-indices for many of the measurements, except for the ones concentrated within the ambiguous region in the center.

Once we identify the initial values of bin-index for the modulo measurements, we can calculate corrected measurements
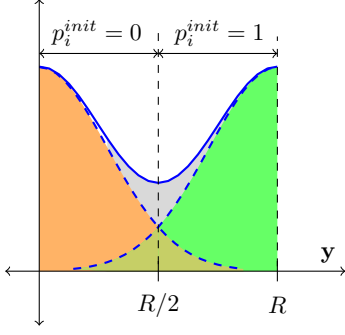
Fig. 4: *Density plot of* $\mod(\mathbf{A}\mathbf{x}^*)$. *Best viewed in color.*

Algorithm 1 MoRAM-INITIALIZATION

Inputs: y, A, s, R
Output: $\mathbf{x^0}$
for $i = 0 : m$ **do**
    Calculate p_i^{init} according to Eq. 3.
end for
Calculate $\mathbf{y_c^{init}}$ according to Eq. 4.
$\mathbf{x^0} \leftarrow H_s\left(\frac{1}{N}\sum_{i=1}^{N} y_{c,i}^{init} a_i\right)$

---

as,

$$\mathbf{y_c^{init}} = \mathbf{y} + \mathbf{p^{init}}R. \tag{4}$$

We use these corrected measurements $\mathbf{y_c^{init}}$ to calculate the initial estimate $\mathbf{x^0}$ with first order unbiased estimator.

$$\mathbf{x^0} = H_s\left(\frac{1}{N}\sum_{i=1}^{N} y_{c,i}^{init} a_i\right), \tag{5}$$

where H_s denotes the hard thresholding operator that keeps the s largest absolute entries of a vector and sets the other entries to zero.

B. Alternating Minimization

Using Eq. 5, we calculate the initial estimate of the signal $\mathbf{x^0}$ which is relatively close to the true vector $\mathbf{x}^*$. Starting with $\mathbf{x^0}$, we calculate the estimates of $\mathbf{p}$ and $\mathbf{x}$ in an alternating fashion to converge to the original signal $\mathbf{x}^*$. At each iteration of alternating-minimization, we use the current estimate of the signal $\mathbf{x^t}$ to get the value of the bin-index vector $\mathbf{p^t}$ as following:

$$\mathbf{p^t} = \frac{1 - \operatorname{sgn}(\langle \mathbf{A} \cdot \mathbf{x^t}\rangle)}{2}. \tag{6}$$

Given that $\mathbf{x^0}$ is close to $\mathbf{x}^*$, we expect that $\mathbf{p^0}$ would also be close to $\mathbf{p}^*$. Ideally, we would calculate the correct compressed measurements $\mathbf{y_c^t}$ using $\mathbf{p^t}$, and use $\mathbf{y_c^t}$ with any popular compressive recovery algorithms such as CoSaMP or basis pursuit to calculate the next estimate $\mathbf{x^{t+1}}$. Thus,

$$\mathbf{y_c^t} = \langle \mathbf{A}\mathbf{x^{t+1}}\rangle = \mathbf{y} - \mathbf{p^t}R,$$

$$\mathbf{x^{t+1}} = \underset{\mathbf{x} \in \mathcal{M}_s}{\arg\min} \|\mathbf{A}\mathbf{x} - \mathbf{y_c^t}\|_2^2,$$

Algorithm 2 MoRAM-DESCENT

Inputs: y, A, s, R
Output: $\mathbf{x^T}$
$m, n \leftarrow \operatorname{size}(\mathbf{A})$
Initialization
$\mathbf{x^0} \leftarrow \text{MoRAM-initialization}(\mathbf{y}, \mathbf{A})$
Alternating Minimization
for $t = 0 : T$ **do**
    $\mathbf{p^t} \leftarrow \frac{1 - \operatorname{sgn}(\langle \mathbf{A} \cdot \mathbf{x^t}\rangle)}{2}$
    $\mathbf{y_c^t} \leftarrow \mathbf{y} - \mathbf{p^t}R$
    $\mathbf{x^{t+1}} \leftarrow JP(\frac{1}{\sqrt{m}}[\mathbf{A} \quad \mathbf{I}], \frac{1}{\sqrt{m}}\mathbf{y_c^t}, [\mathbf{x^t} \quad \mathbf{p^t}]^\top)$.
end for

---

where $\mathcal{M}_s$ denotes the set of s-sparse vectors in $\mathbb{R}^n$. Note that sparsity is only one of several signal models that can be used here, and in principle a rather similar formulation would extend to cases where $\mathcal{M}$ denotes any other structured sparsity model [38].

However, it should be noted that the "bin" error $\mathbf{d^t} = \mathbf{p^t} - \mathbf{p}^*$, even if small, would significantly impact the correction step that constructs $\mathbf{y_c^t}$, as each incorrect bin-index would add a noise of the magnitude R in $\mathbf{y_c^t}$. Our experiments suggest that the typical sparse recovery algorithms are not robust enough to cope up with such large errors in $\mathbf{y_c^t}$. To tackle this issue, we employ an outlier-robust sparse recovery method [10]. We consider the fact that the nature of the error $\mathbf{d^t}$ is sparse with sparsity $s_{dt} = \|\mathbf{d^t}\|_0$; and each erroneous element of $\mathbf{p}$ adds a noise of the magnitude R in $\mathbf{y_c^t}$.

Rewriting in terms of Justice Pursuit, the recovery problem now becomes problem becomes,

$$\mathbf{x^{t+1}} = \underset{[\mathbf{x} \ \mathbf{d}]^\top \in \mathcal{M}_{s+s_{dt}}}{\arg\min} \|[\mathbf{A} \quad \mathbf{I}]\begin{bmatrix}\mathbf{X}\\\mathbf{d}\end{bmatrix} - \mathbf{y_c^t}\|_2^2,$$

However, the sparsity of $\mathbf{d^t}$ is unknown, suggesting that greedy sparse recovery methods cannot be directly used without an additional hyper-parameter. Thus, we employ basis pursuit [39] which does not rely on sparsity. The robust formulation of basis pursuit is referred as Justice Pursuit (JP) [10], specified in Eq. 7.

$$\implies \mathbf{x^{t+1}} = JP(\frac{1}{\sqrt{m}}[\mathbf{A} \quad \mathbf{I}], \frac{1}{\sqrt{m}}\mathbf{y_c^t}, [\mathbf{x^t} \quad \mathbf{p^t}]^\top). \tag{7}$$

Proceeding this way, we repeat the steps of bin-index calculation (as in Eq. 6) and sparse recovery (Eq. 7) altenately for T iterations. Our algorithm is able to achieve convergence to the true underlying signal, as supported by the results in the experiments section.

V. MATHEMATICAL ANALYSIS

Before presenting experimental validation of our proposed MoRAM algorithm, we now perform a theoretical analysis of the descent stage of our algorithm. We assume the availability of an initial estimate $\mathbf{x^0}$ that is close to $\mathbf{x}^*$, i.e. $\|\mathbf{x^0} - \mathbf{x}^*\|_2 \leq \delta\|\mathbf{x}^*\|_2$. In our case, our initialization step (in Alg. 2) provide such $\mathbf{x^0}$.

We perform alternating-minimization as described in 2, starting with $\mathbf{x^0}$ calculated using Alg. 1. For simplicity, we limit

our analysis of the convergence to only one AltMin iteration. In fact, according to our theoretical analysis, if initialized closely enough, one iteration of AltMin suffices for exact signal recovery with sufficiently many measurements. However, in practice we have observed that our algorithm requires more than one AltMin iterations.

The first step is to obtain the initial guess of the bin-index vector (say $\mathbf{p^0}$) using $\mathbf{x^0}$.

$$\mathbf{p^0} = \frac{1 - \mathrm{sgn}(\langle \mathbf{A} \cdot \mathbf{x^0} \rangle)}{2}.$$

If we try to undo the effect of modulo operation by adding back R for the affected measurements based on the bin-index vector $\mathbf{p^0}$, it would introduce an additive error equal to R corresponding to each of the incorrect bin-indices in $\mathbf{p^0}$.

$$\mathbf{y_c^0} = \langle \mathbf{A}\mathbf{x^0} \rangle = \mathbf{y} - \mathbf{p^0}R,$$

We show the guaranteed recovery of the true signal as the corruption in the first set of corrected measurements $\mathbf{y_c^0}$ can be modeled as sparse vector with sparsity less than or equal to λm, with c being a fraction that can be explicitly bounded.

To prove this, we first introduce the concept of *binary ϵ-stable embedding* as proposed by [40]. Let $\mathcal{B}^m$ be a Boolean cube defined as $\mathcal{B}^m := \{-1, 1\}^m$ and let $S^{n-1} := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 = 1\}$ be the unit hyper-sphere of dimension n.

Definition 1 (Binary ϵ-Stable Embedding). *A mapping $F : \mathbb{R}^n \to \mathcal{B}^m$ is a binary ϵ-stable embedding (BϵSE) of order s for sparse vectors if:*

$$d_S(\mathbf{x}, \mathbf{y}) - \epsilon \le d_H(F(\mathbf{x}), F(\mathbf{y})) \le d_S(\mathbf{x}, \mathbf{y}) + \epsilon;$$

for all $\mathbf{x}, \mathbf{y} \in S^{n-1}$ with $|supp(x) \cup supp(y)| \le s$.

In our case, let us define the mapping $F : \mathbb{R}^n \to \mathcal{B}^m$ as:

$$F(\mathbf{x}) := \mathrm{sgn}(\mathbf{A}\mathbf{x});$$

with $\mathbf{A} \sim \mathcal{N}^{m \times n}(0, 1)$. We obtain:

Lemma 2. *Let $\mathbf{A}$ be the matrix generated as $\mathbf{A} \sim \mathcal{N}^{m \times n}(0, 1)$ and suppose $\mathbf{x}^*, \mathbf{x^0} \in \mathbb{R}^n$ are $s-$sparse vectors satisfying $\|\mathbf{x}^* - \mathbf{x^0}\|_2 \le \delta \|\mathbf{x}^*\|_2$. Let $\eta \in [0, 1]$, $\epsilon > 0$. If the number of measurements $m \ge \frac{2}{\epsilon^2} \left(s \log(n) + 2s \log\left(\frac{35}{\epsilon}\right) + \log\left(\frac{2}{\eta}\right) \right)$, then, the following is true with probability exceeding $1 - \eta$:*

$$d_H(\mathrm{sgn}(\mathbf{A}\mathbf{x}^*), \mathrm{sgn}(\mathbf{A}\mathbf{x^0})) \le \frac{\delta}{2} + \epsilon;$$

where d_H is Hamming distance between binary vectors defined as:

$$d_H(\mathbf{a}, \mathbf{b}) := \frac{1}{n} \sum_{i=1}^{n} a_i \oplus b_i,$$

for $n-$ dimensional binary vectors $\mathbf{a}, \mathbf{b}$.

Proof. Given $m \ge \frac{2}{\epsilon^2} \left(s \log(n) + 2s \log\left(\frac{35}{\epsilon}\right) + \log\left(\frac{2}{\eta}\right) \right)$, using Theorem 3 from [40] we conclude that $F(\cdot)$ is a BϵSE for s-sparse vectors. Thus for sparse vectors $\mathbf{x}^*, \mathbf{x^0}$:

$$d_H(F(\mathbf{x}^*), F(\mathbf{x^0})) \le d_S(\mathbf{x}^*, \mathbf{x^0}) + \epsilon. \quad (8)$$

Here, $d_S(\cdot)$ is defined as the natural angle formed by two vectors. Specifically, for $\mathbf{p}, \mathbf{q}$ in unit norm ball,

$$d_s(\mathbf{p}, \mathbf{q}) := \frac{1}{\pi} \arccos\langle \mathbf{p}, \mathbf{q} \rangle = \frac{1}{\pi}\theta,$$

where θ is the angle between two unit norm vectors $\mathbf{p}$ and $\mathbf{q}$.

We note that,

$$\|\mathbf{p} - \mathbf{q}\|_2 = 2\sin(\frac{\theta}{2}).$$

Thus,

$$2d_s(\mathbf{p}, \mathbf{q}) \le \|\mathbf{p} - \mathbf{q}\|_2 \le \pi d_s(\mathbf{p}, \mathbf{q}) \quad (9)$$

Combining eq. 8 and eq. 9, we conclude,

$$d_H(\mathrm{sgn}(\mathbf{A}\mathbf{x}^*), \mathrm{sgn}(\mathbf{A}\mathbf{x^0})) \le \frac{1}{2}\|\mathbf{x}^* - \mathbf{x^0}\| + \epsilon, \quad \text{i.e.,}$$

$$d_H(\mathrm{sgn}(\mathbf{A}\mathbf{x}^*), \mathrm{sgn}(\mathbf{A}\mathbf{x^0})) \le \frac{\delta}{2} + \epsilon. \quad (10)$$

We use this to obtain:

Theorem 3. *Given an initialization $\mathbf{x^0}$ satisfying $\|\mathbf{x}^* - \mathbf{x^0}\|_2 \le \delta \|\mathbf{x}^*\|_2$, for $0 < \delta < 1, \eta \in [0, 1]$, $\epsilon > 0$, if we have number of (Gaussian) measurements satisfying $m \ge \frac{2}{\epsilon^2} \left(s \log(n) + 2s \log\left(\frac{35}{\epsilon}\right) + \log\left(\frac{2}{\eta}\right) \right)$ and $s \le \gamma m / (\log(n/m) + 1)$, then the estimate after the first iteration $\mathbf{x^1}$ of Algorithm 2 is exactly equal to the true signal $\mathbf{x}^*$ with probability at least $1 - K \exp(-cm) - \eta$, with K and c being numerical constants.*

Proof. In the estimation step, Algorithm 2 dubs the problem of recovering the true signal $\mathbf{x}^*$ from the modulo measurements as the special case of signal recovery from sparsely corrupted compressive measurements. As we discussed in Section IV-A1, the presence of modulo operation modifies the compressive measurements by adding a constant noise of the value R in fraction of total measurements. However, once we identify correct bin-index for some of the measurements using $\mathbf{x^0}$, the remaining noise can be modeled as sparse corruptions $\mathbf{d}$ according to the formulation:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{I_n}R(\mathbf{p^0} - \mathbf{p}^*) = \mathbf{A}\mathbf{x} + \mathbf{d}.$$

Here, the $l0$-norm of $\mathbf{d}$ gives us the number of noisy measurements in $\mathbf{y_c^0}$.

If the initial bin-index vector $\mathbf{p^0}$ is close to the true bin-index vector $\mathbf{p}^*$, then $\|\mathbf{d}\|_0$ is small enough with respect to total number of measurements m; thus, $\mathbf{d}$ can be treated as sparse corruption. If we model this corruption as a sparse noise, then we can employ JP for a guaranteed recovery of the true signal given (i) sparsity of the noise is a fraction of total number of measurements; (ii) sufficiently large number of measurements are available.

We compute $\|\mathbf{d}\|_0$ as,

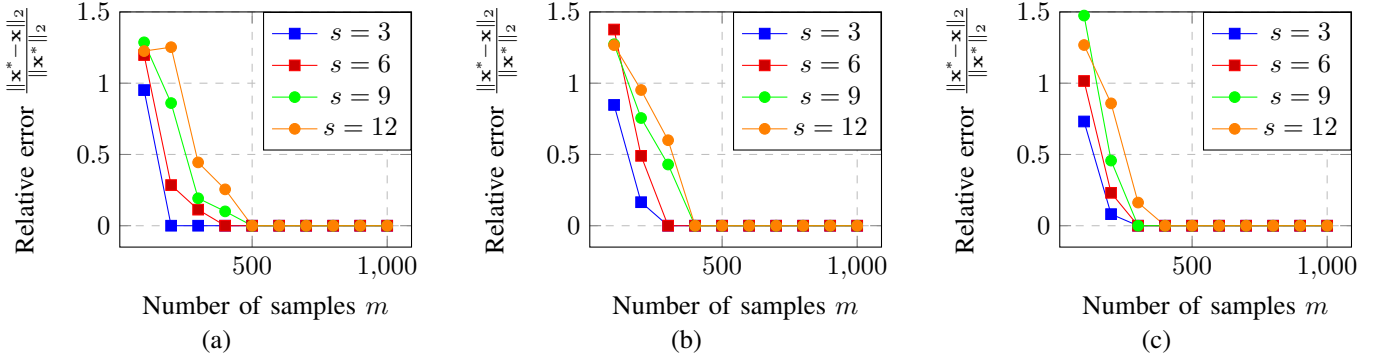$$\|\mathbf{d}\|_0 = \|(\mathbf{p}^* - \mathbf{p^0})R\|_0;$$

Fig. 5: Mean relative reconstruction error vs no. of measurements (m) for MoRAM with $\|\mathbf{x}^*\|_2 = 1, n = 1000$, and (a) $R = 4$; (b) $R = 4.25$; (c) $R = 4.5$.

(a) Original image   (b) $R = 4$, SNR $= 26.10$dB   (c) $R = 4.25$, SNR $= 26.96$dB   (d) $R = 4.5$, SNR $= 27.28$dB

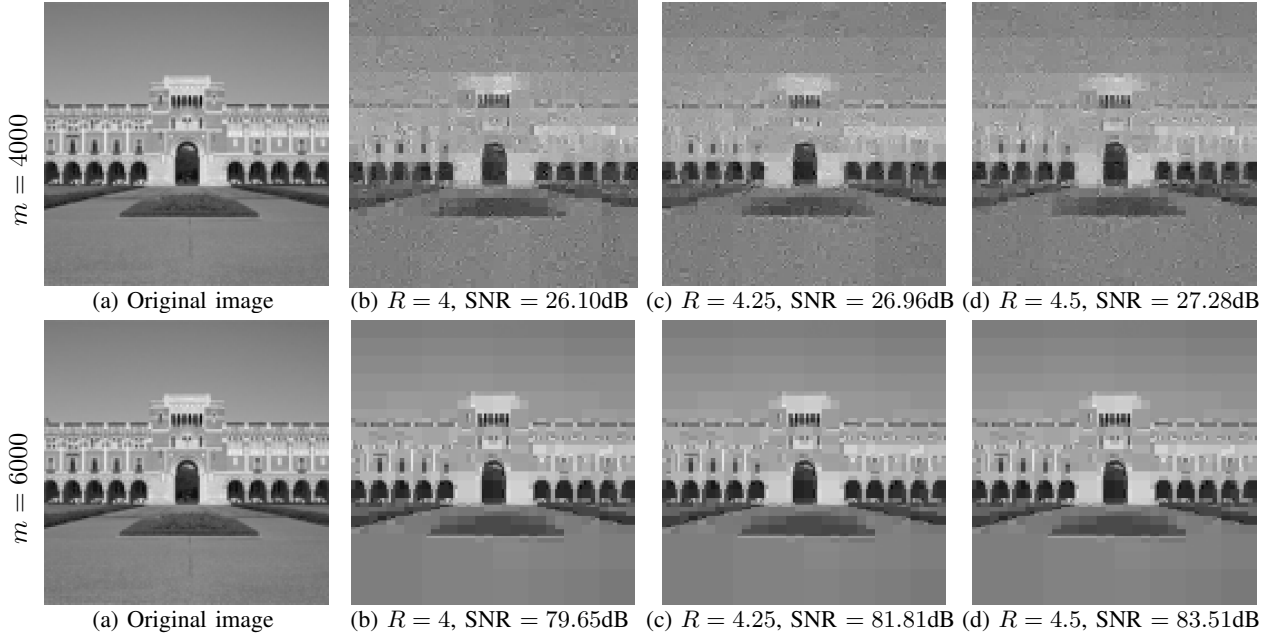(a) Original image   (b) $R = 4$, SNR $= 79.65$dB   (c) $R = 4.25$, SNR $= 81.81$dB   (d) $R = 4.5$, SNR $= 83.51$dB

Fig. 6: (a) Original Lovett Hall image $(n = 16,384)$; sparse reconstructions $(s = 800)$ using $m = 4000$ (top) and $m = 6000$ (bottom) measurements for (b) $R = 4$, (c) $R = 4.25$, (d) $R = 4.5$.

expanding further,

$$\|\mathbf{d}\|_0 = \frac{1 - \operatorname{sgn}(\langle \mathbf{A} \cdot \mathbf{x^0} \rangle)}{2} - \frac{1 - \operatorname{sgn}(\langle \mathbf{A} \cdot \mathbf{x}^* \rangle)}{2}$$
$$= \frac{\operatorname{sgn}(\mathbf{A}\mathbf{x}^*) - \operatorname{sgn}(\mathbf{A}\mathbf{x^0})}{2}$$
$$= \frac{F(\mathbf{x}^*) - F(\mathbf{x^0})}{2}$$
$$= d_H(F(\mathbf{x}^*), F(\mathbf{x^0})).$$

From eq. 10,

$$\leq \frac{\delta}{2} + \epsilon = \gamma m.$$

Algorithm 2 is essentially the Justice Pursuit (JP) formulation as described in [10]. Exact signal recovery from sparsely corrupted measurements is a well-studied domain with uniform recovery guarantees available in the existing literature. We use the guarantee proved in [41] for Gaussian random measurement matrix, which states that one can recover a sparse signal exactly by tractable ℓ_1-minimization even if a positive fraction of the measurements are arbitrarily corrupted. With $\|\mathbf{d}\|_0 \leq \gamma m$, we invoke Theorem 1.1 from [41] to complete the proof. $\qquad\square$

VI. NUMERICAL EXPERIMENTS

In this section, we present the results of simulations of signal reconstruction using our algorithm. All numerical experiments were conducted using MATLAB R2017a on a Linux system with an Intel CPU and 64GB RAM. Our experiments explores the performance of the MoRAM algorithm on both synthetic data as well as real images.

We perform experiments on a synthetic sparse signal $\mathbf{x}^* \in \mathbb{R}^n$ with $n = 1000$. The sparsity level of the signal is chosen in steps of 3 starting from 3 with a maximum value of 12. The non-zero elements of the test signal $\mathbf{x}^*$ are generated using zero-mean Gaussian distribution $\mathcal{N}(0, 1)$ and normalized such

that $\|\mathbf{x}^*\| = 1$. The elements of the Gaussian measurement matrix $\mathbf{A} \in \mathbb{R}^{m \times n}, a_{ij}$ are also generated using the standard normal distribution $\mathcal{N}(0, 1)$. The number of measurements m is varied from $m = 100$ to $m = 1000$ in steps of 100.

Using $\mathbf{A}, \mathbf{x}^*$ and $\mathbb{R}$, We first obtain the compressed modulo measurements $\mathbf{y}$ by passing the signal through forward model described by Eq. 2. We compute the initial estimate $\mathbf{x}^0$ using the algorithm 1. For reconstruction, algorithm 2 is employed. we plot the variation of the relative reconstruction error $(\frac{\|\mathbf{x}^* - \mathbf{x}^T\|}{\|\mathbf{x}^*\|})$ with number of measurements m for our AltMin based sparse recovery algorithm MoRAM.

For each combination of R, m and s, we run 10 independent Monte Carlo trials, and calculate mean of the relative reconstruction error over these trials. Fig. 5 (a), (b) and (c) illustrate the performance of our algorithm for increasing values of R respectively. It is evident that for each combination of R and s, our algorithm converges with probability 1 to give the exact recovery of the true signal (zero relative error) provided enough number of measurements. In all such cases, the minimum number of measurements required for exact recovery are well below the ambient dimension (n) of the underlying signal.

A. Experiments on real image

We also evaluated the performance of our algorithm on a real image. We obtain sparse representation of the real image by transforming the original image in the wavelet basis (db1). The image used in our experiment is 128×128 ($n = 16384$) image of Lovett Hall (fig. 6(a)), and we use the thresholded wavelet transform (with Haar wavelet) to sparsify this image with $s = 800$. We reconstruct the image with MoRAM using $m = 4000$ and $m = 6000$ compressed modulo measurements, for 3 different values of R, $4, 4.25$ and 4.5. As expected, the reconstruction performance increases with increasing value of R. As shown in Fig. 6(bottom), for $m = 6000$, The algorithm produces near-perfect recovery for all 3 values of R with high PSNR.

VII. DISCUSSION

In this paper, for signal recovery from compressed modulo measurements, we presented a novel algorithmic approach inspired from the classical phase retrieval solutions. Our mathematical and experimental analysis support our claim of exact signal recovery through proposed algorithm. Several open questions remain that can serve as the future directions of our work. While in this paper we considered only two periods within the modulo operation, extending the proposed approach for more periods (and theoretically infinite periods) is a significant and interesting research direction. Instead of relying on sparsity prior for compressed recovery, employing novel set of priors such as GAN priors [42], [43] can also be a direction to be explored. Analysis and guaretes for initialization in our algorithm can also be an interesting direction. Moreover, our analysis is limited to the case of Gaussian measurements, thus extending our results to various measurement schemes such as Fourier samples can be an interesting problem for future study.

REFERENCES

[1] A. Bhandari, F. Krahmer, and R. Raskar. On unlimited sampling. *Proc. Sampling Theory and Applications (SampTA)*, pages 31–35, 2017.

[2] H. Zhao, B. Shi, C. Fernandez-Cull, S. Yeung, and R. Raskar. Unbounded high dynamic range photography using a modulo camera. In *Intl. Conf. on Comp. Photography (ICCP)*, 2015.

[3] V. Shah, M. Soltani, and C. Hegde. Reconstruction from periodic nonlinearities, with applications to hdr imaging. In *Proc. Asilomar Conf. Signals, Systems, and Computers*, pages 863–867. IEEE, 2017.

[4] M. Cucuringu and H. Tyagi. On denoising modulo 1 samples of a function. In *Proc. Int. Conf. Art. Intell. Stat. (AISTATS)*, 2018.

[5] J. Rhee and Y. Joo. Wide dynamic range cmos image sensor with pixel level adc. *Electron. Lett.*, 39:360–361, 2010.

[6] S. Kavusi and A. El Gamal. Quantitative study of high-dynamic-range image sensor architectures. In *Sensors and Camera Systems for Sci., Indust., and Digi. Photography Applications V*, volume 5301, pages 264–276. Intl. Soc. for Optics and Photonics, 2004.

[7] K. Sasagawa, T. Yamaguchi, M. Haruta, Y. Sunaga, H. Takehara, H. Takehara, T. Noda, T. Tokuda, and J. Ohta. An implantable cmos image sensor with self-reset pixels for functional brain imaging. *IEEE Trans. on Electron Devices*, 63(1):215–222, 2016.

[8] T. Yamaguchi, H. Takehara, Y. Sunaga, M. Haruta, M. Motoyama, Y. Ohta, T. Noda, K. Sasagawa, T. Tokuda, and J. Ohta. Implantable self-reset cmos image sensor and its application to hemodynamic response detection in living mouse brain. *Japanese J. of Appl. Physics*, 55(4S):04EM02, 2016.

[9] P. Netrapalli, P. Jain, and S. Sanghavi. Phase retrieval using alternating minimization. In *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pages 2796–2804, 2013.

[10] J. Laska, M. Davenport, and R. Baraniuk. Exact signal recovery from sparsely corrupted measurements through the pursuit of justice. In *Proc. Asilomar Conf. Signals, Systems, and Computers*, pages 1556–1560, 2009.

[11] E. Candes, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Comm. Pure Appl. Math.*, 66(8):1241–1274, 2013.

[12] D. Gross, F. Krahmer, and R. Kueng. Improved recovery guarantees for phase retrieval from coded diffraction patterns. *Appl. Comput. Harmon. Anal.*, 42(1):37–64, 2017.

[13] E. Candes, X. Li, and M. Soltanolkotabi. Phase retrieval from coded diffraction patterns. *Appl. Comput. Harmon. Anal.*, 39(2):277–299, 2015.

[14] S. Bahmani and J. Romberg. Phase retrieval meets statistical learning theory: A flexible convex relaxation. In *Proc. Int. Conf. Art. Intell. Stat. (AISTATS)*, 2016.

[15] T. Goldstein and C. Studer. Phasemax: Convex phase retrieval via basis pursuit. *IEEE Trans. Inform. Theory*, 64:2675–2689, 2018.

[16] O. Dhifallah, C. Thrampoulidis, and Y. Lu. Phase retrieval via linear programming: Fundamental limits and algorithmic improvements. *Proc. Allerton Conf. Communication, Control, and Computing*, pages 1071–1077, 2017.

[17] F. Salehi, E. Abbasi, and B. Hassibi. A precise analysis of phasemax in phase retrieval. *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, pages 976–980, 2018.

[18] E. Candes, X. Li, and M. Soltanolkotabi. Phase retrieval via wirtinger flow: theory and algorithms. *IEEE Trans. Inform. Theory*, 61(4):1985–2007, 2015.

[19] H. Zhang and Y. Liang. Reshaped wirtinger flow for solving quadratic system of equations. In *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pages 2622–2630, 2016.

[20] Y. Chen and E. Candes. Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pages 739–747, 2015.

[21] T. Cai, X. Li, Z. Ma, et al. Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow. *Ann. Stat.*, 44(5):2221–2251, 2016.

[22] G. Wang, L. Zhang, G. Giannakis, M. Akçakaya, and J. Chen. Sparse phase retrieval via truncated amplitude flow. *arXiv preprint arXiv:1611.07641*, 2016.

[23] G. Wang and G. Giannakis. Solving random systems of quadratic equations via truncated generalized gradient flow. In *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pages 568–576, 2016.

[24] H. Ohlsson, A. Yang, R. Dong, and S. Sastry. Cprl–an extension of compressive sensing to the phase retrieval problem. In *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pages 1367–1375, 2012.

[25] X. Li and V. Voroninski. Sparse signal recovery from quadratic measurements via convex programming. *SIAM J. on Math. Analysis*, 45(5):3019–3033, 2013.

[26] S. Bahmani and J. Romberg. Efficient compressive phase retrieval with constrained sensing vectors. In *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pages 523–531, 2015.

[27] K. Jaganathan, S. Oymak, and B. Hassibi. Recovery of sparse 1-d signals from the magnitudes of their fourier transform. In *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, pages 1473–1477. IEEE, 2012.

[28] G. Jagatap and C. Hegde. Fast, sample-efficient algorithms for structured phase retrieval. In *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, 2017.

[29] J. Bioucas-Dias and G. Valadao. Phase unwrapping via graph cuts. *IEEE Trans. Image Proc.*, 16(3):698–709, 2007.

[30] F. Lang, T. Plötz, and S. Roth. Robust multi-image HDR reconstruction for the modulo camera. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 10496 LNCS, pages 78–89, 2017.

[31] Mohammadreza Soltani and Chinmay Hegde. Stable recovery of sparse vectors from random sinusoidal feature maps. *icassp*, pages 6384–6388, 2017.

[32] A. Bhandari, F. Krahmer, and R. Raskar. Unlimited sampling of sparse signals. *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, pages 4569–4573, 2018.

[33] O. Musa, P. Jung, and N. Goertz. Generalized approximate message passing for unlimited sampling of sparse signals. *Proc. of IEEE Global Conf. on Signal and Inform. Processing (GlobalSIP)*, pages 336–340, 2018.

[34] D. Needell and J. Tropp. Cosamp: iterative signal recovery from incomplete and inaccurate samples. *Comm. of the ACM*, 53(12):93–100, 2010.

[35] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.

[36] E. van den Berg and M. P. Friedlander. SPGL1: A solver for large-scale sparse reconstruction, June 2007. http://www.cs.ubc.ca/labs/scl/spgl1.

[37] E. van den Berg and M. P. Friedlander. Probing the pareto frontier for basis pursuit solutions. *SIAM J. on Sci. Computing*, 31(2):890–912, 2008.

[38] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE Trans. Inform. Theory*, 56:1982–2001, 2010.

[39] Emmanuel J Candès et al. Compressive sampling. In *Proceedings of the international congress of mathematicians*, volume 3, pages 1433–1452. Madrid, Spain, 2006.

[40] L. Jacques, J. Laska, P. Boufounos, and R. Baraniuk. Robust 1-Bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Trans. Inform. Theory*, 59(4):2082–2102, 2013.

[41] X. Li. Compressed sensing and matrix completion with constant proportion of corruptions. *Const. Approx.*, 37(1):73–99, 2013.

[42] A. Bora, A. Jalal, E. Price, and A. Dimakis. Compressed sensing using generative models. *Proc. Int. Conf. Machine Learning*, 2017.

[43] V. Shah and C. Hegde. Solving linear inverse problems using gan priors: An algorithm with provable guarantees. *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, 2018.