Sparse Signal Recovery from Modulo Observations

Viraj Shah

University of Illinois at Urbana-Champaign

Chinmay Hegde (✉ chinmay.h@nyu.edu)

New York University Tandon School of Engineering <https://orcid.org/0000-0003-4574-8066>

Research

Keywords: Sparse recovery, High dynamic range imaging, Modulo sensing

Posted Date: July 24th, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-42731/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at EURASIP Journal on Advances in Signal Processing on April 6th, 2021. See the published version at <https://doi.org/10.1186/s13634-021-00722-w>.

RESEARCH

Sparse Signal Recovery from Modulo Observations

Viraj Shah¹ and Chinmay Hegde^{2*}

A conference version of this manuscript appeared in IEEE GlobalSIP 2019 [1].

*Correspondence:

chinmay.h@nyu.edu

²Tandon School of Engineering,
New York University, Brooklyn,
NY, USA

Full list of author information is
available at the end of the article

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.

Keywords: Sparse recovery; High dynamic range imaging; Modulo sensing

Introduction

The problem of reconstructing a signal (or image) from (possibly) nonlinear observations is a principal challenge in 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:

$$y_i = \text{mod}(\langle \mathbf{a}_i, \mathbf{x}^* \rangle, R) \quad i = \{1, 2, \dots, m\}. \quad (1)$$

The task is to recover $\mathbf{x}^*$ using the modulo measurements $\mathbf{y}$ and knowledge of the measurement matrix $\mathbf{A} = [\mathbf{a}_1 \ \mathbf{a}_2 \ \dots \ \mathbf{a}_m]^\top$.

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 [2–5]. 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 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 [6–9]. 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 [2, 3] 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. 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., $\langle \mathbf{a}_i \cdot \mathbf{x}^* \rangle \in [-R, R] \forall i \in \{1, 2, \dots, m\}$.

Our contributions

In this paper, we propose a recovery algorithm for exact reconstruction of sparse 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:

$$y_i = |\langle \mathbf{a}_i, \mathbf{x}^* \rangle|, \quad i = 1, 2, \dots, m,$$

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.

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 equivalent to retrieving the phase ($\text{sign}(y_i)$) corresponding to each measurement y_i . We assume the real-valued case where 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.

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 [2, 3].

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 [10]. 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 [10]) 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 [11]). 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.

Prior work

At a high level, our algorithmic development follows two (hitherto disconnected) streams of work in the signal processing literature. Since this is a very large area of study, our review of prior work will unfortunately be incomplete.

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 [12] usually consist of solving a constrained optimization problem after lifting the true signal $\mathbf{x}^*$ in higher dimensional space. Typical non-convex approaches involve finding a good initialization, followed by iterative minimization of a loss function. Approaches based on Wirtinger Flow [13–16] and Amplitude flow [17, 18] 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 [19–22]. Similarly, non-convex approaches for sparse phase retrieval include [10, 16, 17]. 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 [23].

Modulo recovery

The modulo recovery problem is also known in the classical signal processing literature as phase unwrapping. The algorithm proposed in [24] 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 [3] 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 [25]. 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 [4], we proposed a different extension based on [3, 26] 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, [2, 27] provide a stable algorithm for signal recovery and also proves sufficiency conditions that guarantees the recovery. [5] formulates and solves a 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 rely 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, [28, 29] proposed unlimited sampling algorithm for sparse signals and images. Similar to [2], it also exploits the bandlimitedness by considering the low-pass filtered version of the sparse signal, and thus differs from our random measurements setup. In [30], modulo recovery from Gaussian random measurements is considered. However, it assumes the true signal to be distributed as a mixed Bernoulli-Gaussian distribution which is not a standard assumption.

For a qualitative comparison of our MoRAM method with existing approaches, refer Table 1. 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.

Methods

Preliminaries

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 $\text{card}(S)$. We define the signum function as $\text{sgn}(x) := \frac{x}{|x|}$ for every $x \in \mathbb{R}, x \neq 0$, with the convention that $\text{sgn}(0) = 1$. 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.

Mathematical model

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:

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

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 $\text{sgn}(t) = 1$, or 1 when $\text{sgn}(t) = -1$. Inserting the definition of f in the measurement model of Eq. 1 gives,

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

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,

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

If we ignore the presence of the modulo operation in the 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 [31] or basis-pursuit [32–34].

Sparse signal recovery

Our 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 $\mathbf{y}, \mathbf{A}, \mathbf{s}, \mathbf{R}$. Our algorithm *MoRAM* (*Modulo Reconstruction with Alternating Minimization*) comprises of two stages: (i) an initialization stage, and (ii) descent stage via alternating minimization.

Initialization

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.

Effect of the modulo transfer function

To provide some intuition, let us first examine the relation between the distributions of $\mathbf{A}\mathbf{x}^*$ and $\text{mod}(\mathbf{A}\mathbf{x}^*)$. It is easy to see that the compressed measurements $\mathbf{y}_c$ follow the standard normal distribution, as $\mathbf{A}$ is a Gaussian random matrix.

Algorithm 1 MORAM-INITIALIZATION

Inputs: $\mathbf{y}, \mathbf{A}, s, R$
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.
 $\mathbf{x}^0 \leftarrow H_s \left(\frac{1}{N} \sum_{i=1}^N y_{c,i}^{init} a_i \right)$

We can now 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. 2, after the 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 by observing their magnitudes relative to the midpoint $R/2$. This leads us to the following estimator for bin-indices ($\mathbf{p}$):

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

The vector $\mathbf{p}^{init}$ obtained with the 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 as,

$$\mathbf{y}_c^{init} = \mathbf{y} + \mathbf{p}^{init} \mathbf{R}. \quad (4)$$

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

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

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.

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:

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

Algorithm 2 MoRAM-DESCENT

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

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,

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

$$\mathbf{x}^{t+1} = \arg \min_{\mathbf{x} \in \mathcal{M}_s} \|\mathbf{A} \mathbf{x} - \mathbf{y}_c^t\|_2^2,$$

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 [35].

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 [11]. 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,

$$\mathbf{x}^{t+1} = \arg \min_{[\mathbf{x} \quad \mathbf{d}]^\top \in \mathcal{M}_{s+s_{dt}}} \left\| \begin{bmatrix} \mathbf{A} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{d} \end{bmatrix} - \mathbf{y}_c^t \right\|_2^2,$$

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 [36] which does not rely on sparsity. The robust formulation of basis pursuit is referred as Justice Pursuit (JP) [11], specified in Eq. 7.

$$\Rightarrow \mathbf{x}^{t+1} = JP\left(\frac{1}{\sqrt{m}} [\mathbf{A} \quad \mathbf{I}], \frac{1}{\sqrt{m}} \mathbf{y}_c^t, [\mathbf{x}^t \quad \mathbf{p}^t]^\top\right). \quad (7)$$

Proceeding this way, we repeat the steps of bin-index calculation (as in Eq. 6) and sparse recovery (Eq. 7) alternately for T iterations. Our algorithm is able to

achieve convergence to the true underlying signal, as supported by the results in the experiments section.

Mathematical Analysis

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 (AltMin) 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 $\mathbf{x}^0$.

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

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 [37]. 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 $F : \mathbb{R}^n \rightarrow \mathcal{B}^m$ is a binary ϵ -stable embedding (BeSE) of order s for sparse vectors if:*

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

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

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

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

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 $\|\mathbf{x}^* - \mathbf{x}^0\|_2 \leq \delta \|\mathbf{x}^*\|_2$. Let $\eta \in [0, 1]$, $\epsilon > 0$. If the number of measurements*

$m \geq \frac{2}{\epsilon^2} \left(s \log(n) + 2s \log\left(\frac{35}{\epsilon}\right) + \log\left(\frac{2}{\eta}\right) \right)$, then, the following is true with probability exceeding $1 - \eta$:

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

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 $m \geq \frac{2}{\epsilon^2} \left(s \log(n) + 2s \log\left(\frac{35}{\epsilon}\right) + \log\left(\frac{2}{\eta}\right) \right)$, using Theorem 3 from [37] we conclude that $F(\cdot)$ is a BcSE for s -sparse vectors. Thus for sparse vectors $\mathbf{x}^*, \mathbf{x}^0$:

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

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,

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

Thus,

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

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

$$\begin{aligned} d_H(\text{sgn}(\mathbf{A}\mathbf{x}^*), \text{sgn}(\mathbf{A}\mathbf{x}^0)) &\leq \frac{1}{2} \|\mathbf{x}^* - \mathbf{x}^0\| + \epsilon, \quad \text{i.e.,} \\ d_H(\text{sgn}(\mathbf{A}\mathbf{x}^*), \text{sgn}(\mathbf{A}\mathbf{x}^0)) &\leq \frac{\delta}{2} + \epsilon. \end{aligned} \quad (10)$$

We use this to obtain:

Theorem 3 Given an initialization $\mathbf{x}^0$ satisfying $\|\mathbf{x}^* - \mathbf{x}^0\|_2 \leq \delta \|\mathbf{x}^*\|_2$, for $0 < \delta < 1, \eta \in [0, 1], \epsilon > 0$, if we have number of (Gaussian) measurements satisfying $m \geq \frac{2}{\epsilon^2} \left(s \log(n) + 2s \log\left(\frac{35}{\epsilon}\right) + \log\left(\frac{2}{\eta}\right) \right)$ and $s \leq \gamma m / (\log(n/m) + 1)$, 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. 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:

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

Here, the ℓ_0 -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:

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

Expanding further,

$$\begin{aligned} \|\mathbf{d}\|_0 &= \frac{1 - \text{sgn}(\langle \mathbf{A} \cdot \mathbf{x}^0 \rangle)}{2} - \frac{1 - \text{sgn}(\langle \mathbf{A} \cdot \mathbf{x}^* \rangle)}{2} \\ &= \frac{\text{sgn}(\mathbf{A}\mathbf{x}^*) - \text{sgn}(\mathbf{A}\mathbf{x}^0)}{2} \\ &= \frac{F(\mathbf{x}^*) - F(\mathbf{x}^0)}{2} \\ &= d_H(F(\mathbf{x}^*), F(\mathbf{x}^0)), \\ &\leq \frac{\delta}{2} + \epsilon = \gamma m. \end{aligned}$$

Algorithm 2 is essentially the Justice Pursuit (JP) formulation as described in [11]. 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 [38] 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 [38] to complete the proof. $\square$

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. 3 (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.

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. 4(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. 4(bottom), for $m = 6000$, The algorithm produces near-perfect recovery for all 3 values of R with high PSNR.

Conclusions

In this paper, we presented a novel algorithmic approach for sparse signal recovery from compressed modulo measurements, inspired by techniques from phase retrieval. We also support our proposed algorithm via mathematical analysis and several experimental results. Our work points the way to a few directions for further research. While in this paper we considered only two modulo periods, extending the proposed approach for more periods (up to a theoretically infinite number) 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 [39–41] is an additional direction. Moreover, our analysis is limited to the case of Gaussian measurements schemes, which may or may not be physically realizable. Extending our results to more practical measurement schemes such as Fourier-based sampling or ptychography [42] can be an interesting problem for future study.

Declarations

Availability of data and materials

Reproducible source code and datasets supporting the conclusions of this paper can be accessed via the Github repository <http://github.com/shahviraj/MoRAM>. This is made accessible under the MIT License.

Competing interests

The authors declare that they have no competing interests.

Funding

This work was supported by grants CCF-1566281 and CAREER CCF-1750920 from the National Science Foundation, a faculty fellowship grant from the Black and Veatch Foundation, and a GPU grant from the NVIDIA Corporation.

Author's contributions

VS contributed to the theory developments and the experiments; CH contributed to the problem formulation, theory, and overall project guidance. All authors read and approved the final manuscript.

Acknowledgements

This work was conducted when both VS and CH were at the ECE Department at Iowa State University, Ames IA, USA. The authors thank Praneeth Narayanamurthy, Gauri Jagatap, and Thanh Nguyen for helpful comments.

Author details

¹ECE Department, University of Illinois Urbana-Champaign, Urbana, IL, USA. ²Tandon School of Engineering, New York University, Brooklyn, NY, USA.

References

- Shah, V., Hegde, C.: Signal reconstruction from modulo observations. In: 2019 IEEE Global Conference on Signal and Information Processing (GlobalSIP), pp. 1–5 (2019)
- Bhandari, A., Krahmer, F., Raskar, R.: On unlimited sampling. *Proc. Sampling Theory and Applications (SampTA)*, 31–35 (2017)
- Zhao, H., Shi, B., Fernandez-Cull, C., Yeung, S., Raskar, R.: Unbounded high dynamic range photography using a modulo camera. In: *Intl. Conf. on Comp. Photography (ICCP)* (2015)
- Shah, V., Soltani, M., Hegde, C.: Reconstruction from periodic nonlinearities, with applications to hdr imaging. In: *Proc. Asilomar Conf. Signals, Systems, and Computers*, pp. 863–867 (2017). IEEE
- Cucuringu, M., Tyagi, H.: On denoising modulo 1 samples of a function. In: *Proc. Int. Conf. Art. Intell. Stat. (AISTATS)* (2018)
- Rhee, J., Joo, Y.: Wide dynamic range cmos image sensor with pixel level adc. *Electron. Lett.* **39**, 360–361 (2010)
- Kavusi, S., El Gamal, A.: Quantitative study of high-dynamic-range image sensor architectures. In: *Sensors and Camera Systems for Sci., Indust., and Digi. Photography Applications V*, vol. 5301, pp. 264–276 (2004). Intl. Soc. for Optics and Photonics
- Sasagawa, K., Yamaguchi, T., Haruta, M., Sunaga, Y., Takehara, H., Takehara, H., Noda, T., Tokuda, T., Ohta, J.: An implantable cmos image sensor with self-reset pixels for functional brain imaging. *IEEE Trans. on Electron Devices* **63**(1), 215–222 (2016)
- Yamaguchi, T., Takehara, H., Sunaga, Y., Haruta, M., Motoyama, M., Ohta, Y., Noda, T., Sasagawa, K., Tokuda, T., Ohta, J.: Implantable self-reset cmos image sensor and its application to hemodynamic response detection in living mouse brain. *Japanese J. of Appl. Physics* **55**(4S), 04–02 (2016)
- Netrapalli, P., Jain, P., Sanghavi, S.: Phase retrieval using alternating minimization. In: *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pp. 2796–2804 (2013)
- Laska, J., Davenport, M., Baraniuk, R.: Exact signal recovery from sparsely corrupted measurements through the pursuit of justice. In: *Proc. Asilomar Conf. Signals, Systems, and Computers*, pp. 1556–1560
- Candes, E., Strohmer, T., Vershynin, V.: Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Comm. Pure Appl. Math.* **66**(8), 1241–1274 (2013)
- Candes, E., Li, X., Soltanolkotabi, M.: Phase retrieval via wirtinger flow: theory and algorithms. *IEEE Trans. Inform. Theory* **61**(4), 1985–2007 (2015)
- Zhang, H., Liang, Y.: Reshaped wirtinger flow for solving quadratic system of equations. In: *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pp. 2622–2630 (2016)
- Chen, Y., Candes, E.: Solving random quadratic systems of equations is nearly as easy as solving linear systems. In: *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pp. 739–747 (2015)
- Cai, T., Li, X., Ma, Z., et al.: Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow. *Ann. Stat.* **44**(5), 2221–2251 (2016)
- Wang, G., Zhang, L., Giannakis, G., Akçakaya, M., Chen, J.: Sparse phase retrieval via truncated amplitude flow. *arXiv preprint arXiv:1611.07641* (2016)
- Wang, G., Giannakis, G.: Solving random systems of quadratic equations via truncated generalized gradient flow. In: *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pp. 568–576 (2016)
- Ohlsson, H., Yang, A., Dong, R., Sastry, S.: Cprl—an extension of compressive sensing to the phase retrieval problem. In: *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pp. 1367–1375 (2012)
- Li, X., Vershynin, V.: Sparse signal recovery from quadratic measurements via convex programming. *SIAM J. on Math. Analysis* **45**(5), 3019–3033 (2013)
- Bahmani, S., Romberg, J.: Efficient compressive phase retrieval with constrained sensing vectors. In: *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, pp. 523–531 (2015)
- Jaganathan, K., Oymak, S., Hassibi, B.: Recovery of sparse 1-d signals from the magnitudes of their fourier transform. In: *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, pp. 1473–1477 (2012). IEEE
- Jagatap, G., Hegde, C.: Fast, sample-efficient algorithms for structured phase retrieval. In: *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)* (2017)
- Bioucas-Dias, J., Valadao, G.: Phase unwrapping via graph cuts. *IEEE Trans. Image Proc.* **16**(3), 698–709 (2007)

25. Lang, F., Plötz, T., Roth, S.: 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), vol. 10496 LNCS, pp. 78–89 (2017). [1707.01317](https://arxiv.org/pdf/1707.01317v1.pdf). <https://arxiv.org/pdf/1707.01317v1.pdf>
26. Soltani, M., Hegde, C.: Stable recovery of sparse vectors from random sinusoidal feature maps. *icassp*, 6384–6388 (2017)
27. Ordentlich, O., Tabak, G., Hanumolu, P.K., Singer, A.C., Wornell, G.W.: A modulo-based architecture for analog-to-digital conversion. *IEEE Journal of Selected Topics in Signal Processing* **12**(5), 825–840 (2018)
28. Bhandari, A., Krahmer, F., Raskar, R.: Unlimited sampling of sparse signals. *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, 4569–4573 (2018)
29. Rudresh, S., Adiga, A., Shenoy, B.A., Seelamantula, C.S.: Wavelet-based reconstruction for unlimited sampling. In: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 4584–4588 (2018). IEEE
30. Musa, O., Jung, P., Goertz, N.: Generalized approximate message passing for unlimited sampling of sparse signals. *Proc. of IEEE Global Conf. on Signal and Inform. Processing (GlobalSIP)*, 336–340 (2018)
31. Needell, D., Tropp, J.: Cosamp: iterative signal recovery from incomplete and inaccurate samples. *Comm. of the ACM* **53**(12), 93–100 (2010)
32. Chen, S., Donoho, D., Saunders, M.: Atomic decomposition by basis pursuit. *SIAM review* **43**(1), 129–159 (2001)
33. van den Berg, E., Friedlander, M.P.: SPGL1: A solver for large-scale sparse reconstruction. <http://www.cs.ubc.ca/labs/scl/spgl1> (2007)
34. van den Berg, E., Friedlander, M.P.: Probing the pareto frontier for basis pursuit solutions. *SIAM J. on Sci. Computing* **31**(2), 890–912 (2008)
35. Baraniuk, R., Cevher, V., Duarte, M., Hegde, C.: Model-based compressive sensing. *IEEE Trans. Inform. Theory* **56**, 1982–2001 (2010)
36. Candès, E.J., et al.: Compressive sampling. In: *Proceedings of the International Congress of Mathematicians*, vol. 3, pp. 1433–1452 (2006). Madrid, Spain
37. Jacques, L., Laska, J., Boufounos, P., Baraniuk, R.: Robust 1-Bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Trans. Inform. Theory* **59**(4), 2082–2102 (2013). [1104.3160](https://arxiv.org/abs/1104.3160)
38. Li, X.: Compressed sensing and matrix completion with constant proportion of corruptions. *Const. Approx.* **37**(1), 73–99 (2013)
39. Bora, A., Jalal, A., Price, E., Dimakis, A.: Compressed sensing using generative models. *Proc. Int. Conf. Machine Learning* (2017)
40. Shah, V., Hegde, C.: Solving linear inverse problems using gan priors: An algorithm with provable guarantees. *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)* (2018)
41. Jagatap, G., Hegde, C.: High dynamic range imaging using deep image priors. In: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 9289–9293 (2020)
42. Jagatap, G., Chen, Z., Nayer, S., Hegde, C., Vaswani, N.: Sample efficient fourier ptychography for structured data. *IEEE Transactions on Computational Imaging* **6**, 344–357 (2020)

Figures

Figure 1: The modulo transfer function.

Figure 2: *Density plot of $\text{mod}(\mathbf{Ax}^*)$* . Best viewed in color.

Figure 3: Mean relative reconstruction error vs no. of measurements (m) for MoRAM with $\|\mathbf{x}^*\|_2 = 1$, $n = 1000$, and (a) $R = 4$; (b) $R = 4.5$.

Figure 4: **Real image recovery results.** (a) Original Lovett Hall image; sparse reconstructions using $m = 4000$ (top) and $m = 6000$ (bottom) measurements for (b) $R = 4$, (c) $R = 4.25$, (d) $R = 4.5$.

Tables

Table 1: Comparison of MoRAM with existing modulo recovery methods.

	Unlimited Sampling [2]	OLS Method [5]	multishot UHDR [3]	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	$2R$

$m = 4000$

(a) Original image

(b) $R = 4$, SNR = 26.10dB (c) $R = 4.25$, SNR = 26.96dB (d) $R = 4.5$, SNR = 27.28dB $m = 6000$

(a) Original image

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

Figures

Figure 1

The modulo transfer function.

Figure 2

Density plot of $\text{mod}(Ax^*)$. Best viewed in color.

Figure 3

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

Figure 4

Real image recovery results. (a) Original Lovett Hall image; sparse reconstructions using $m = 4000$ (top) and $m = 6000$ (bottom) measurements for (b) $R = 4$, (c) $R = 4.25$, (d) $R = 4.5$.