

RESEARCH

Open Access

Sparse signal recovery from modulo observations



Viraj Shah¹ and Chinmay Hegde^{2*}

*Correspondence:

chinmay.h@nyu.edu

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

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 new 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 signal recovery problem under sparsity constraints for the special case to modulo folding limited to two periods. We show that given a sufficient number of measurements, our algorithm perfectly recovers the underlying signal. We also provide experiments validating our approach on toy signal and image data and demonstrate its promising performance.

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


1 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. 1a 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 [2–5]. 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 [6–8]. Several 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, signal measurements 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 the 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 the sensor dynamic range via enhanced hardware; this, of course, can be untenably expensive. An intriguing alternative is to deploy special digital modulo sensors [2–5]. 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 [6, 7] 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 inverse problem (1) with very few 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 special case when the underlying signal is sparse, but our techniques could be extended to other signal structures. Further, for simplicity, we assume that our observation model is limited to only two modulo fold periods, one for positive- and one for negative-valued coefficients. This does not reflect practice, but we will see that such a 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 finite, but large enough that most of the measurements fall within the interval $[-R, R]$. We emphasize that such requirement does not put a hard constraint in our algorithm, and successful recovery is possible even if some measurements lie outside the interval covered by two modulo periods.

1.1 Our contributions

In this paper, we propose a recovery algorithm for exact reconstruction of sparse signals from modulo measurements of the form (1) with the modulo fold operation being limited to two periods. We refer to our algorithm as *MoRAM*, short for *Modulo Recovery using*

Alternating Minimization. We observe that the modulo operation with respect to parameter R can be seen as the addition or subtraction of an integer multiple of R from the input. We refer to this integer multiple as the bin-index $\mathbf{p}$. It is not hard to observe that a successful recovery of bin-index can lead to successful recovery of the input using standard sparse recovery methods. Concretely, the forward model in (1) can be written as follows:

$$y_i = \langle \mathbf{a}_i, \mathbf{x}^* \rangle + p_i^* R, \quad i = \{1, \dots, m\}.$$

As we restrict our modulo operation to only two periods, the bin-index can only take two values: 0 for non-negative inputs, and 1 for negative inputs. We discuss extending this to multiple modulo fold periods in Section 4.

As mentioned above, 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. However, due to the highly nonlinear nature of the forward model, each incorrect bin-index introduces a gross additive error of magnitude R irrespective of the magnitude of the measurement. Commonly used signal recovery algorithms fail to cope up with such gross corruptions of large magnitudes, and perform poorly on our problem (both in theory and practice). Thus, it becomes even more important to be able to find a good estimate of the bin-index and to employ a robust signal recovery algorithm that can recover the signal even in the presence of grossly corrupted measurements.

To this end, our proposed algorithm follows two steps. We first leverage the structure of the modulo folding operator to obtain a promising initial estimate of the bin-indices. We then plug this estimate into a robust sparse recovery formulation that gives us the final signal estimate. We provide analytical correctness proofs for both states and show that signal recovery can be performed using an (essentially) optimal number of observations provided certain assumptions are met. To the best of our knowledge, we are the first to pursue this type of approach for modulo recovery problems with Gaussian linear measurements, distinguishing us from previous work [6, 7].

1.2 Techniques

Our proposed algorithm (MoRAM) consists of two stages.

In the first stage, we identify a good initial estimated bin-index $\mathbf{p}^0$ that is relatively close to the true bin-index $\mathbf{p}^*$. To obtain this estimate, we employ a simple bin-index estimation step by comparing our observed measurements with typical density plots of Gaussian observations. This method is able to recover a large number of bin-indices correctly, and also provides a provable upper bound on the number of erroneous bin-indices.

In the second stage, we use the initial estimate of the bin-indices of the measurements to recover the true underlying signal. We follow an alternating minimization (AltMin) approach inspired from phase retrieval algorithms (such as [9]) that estimates the signal and the measurement of 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, using the Justice Pursuit algorithm [10]). We prove that given enough number of measurements, our Justice Pursuit-based AltMin approach succeeds provided the number of wrongly estimated bin-indices in the beginning is a

sufficiently small fraction of the total number of measurements. This gives us a natural radius of initialization for the initial bin-index estimate and also leads to provable sample-complexity upper bounds.

1.3 Prior work

Since signal recovery from nonlinear measurements is a very large and classical area of study, our review of prior work will unfortunately be incomplete.

1.3.1 Modulo recovery

The modulo recovery problem shares some similarities with the problem of phase unwrapping from the classical signal processing literature that would allow one to use apply phase unwrapping methods to modulo recovery; however, these methods do not provide proven guarantees. For example, the algorithm proposed in [11] 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 [7] 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 [12]. However, the multi-shot approach depends on carefully designed camera exposures, while our approach succeeds for non-designed random Gaussian linear observations; moreover, they do not include sparsity in their model reconstructions. In our previous work [8], we proposed a different extension based on [7, 13] 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, [6, 14, 15] provide a stable algorithm for signal recovery and also proves sufficiency conditions that guarantees the recovery. Ordentlich et al. [14] use non-modulo data for the initialization to exploit the statistical structure in order to undo the effects of the modulo operation. Cucuringu and Tyagi [16] formulate and solve 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). On the contrast, our approach does not rely on such smoothness assumptions.

In recent works, [17] proposed unlimited sampling algorithm for sparse signals and images. Similar to [6], it also exploits the band-limitedness by considering the low-pass filtered version of the sparse signal and thus differs from our random measurements setup. In [18], 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, we refer the reader to 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

Table 1 Comparison of MoRAM with existing modulo recovery methods

	Unlimited sampling [6]	OLS method [16]	multi-shot UHDR [7]	GAMP [18]	MoRAM (our approach)
Assumption on structure of signal	Band-limited	Band-limited	No assumptions	Sparsity, Bernoulli-Gaussian mixture distribution	Sparsity
Sampling scheme	Uniform grid	Uniform grid	(Carefully chosen) linear measurements	Random linear measurements	Random linear measurements
Sample complexity	Oversampled, $\mathcal{O}(n)$	–	Oversampled, $\mathcal{O}(n)$	Undersampled	Undersampled, $\mathcal{O}(\epsilon \log(n))$
Provides sample complexity bounds?	Yes	–	No	No	Yes
Leverages sparsity?	No	No	No	Yes	Yes
(Theoretical) bound on dynamic range	Unbounded	Unbounded	Unbounded	Unbounded	$2R$

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 a non-uniform sampling grid (random measurements) and achieve a provable sub-Nyquist sample complexity.

2 Methods

2.1 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 $\mathbf{x}^\top$ and $\mathbf{A}^\top$ 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, the 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} .

2.2 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 most of the measurements $\langle \mathbf{a}_i, \mathbf{x}^* \rangle$ 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,

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

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,

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

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 $\langle \mathbf{a}_i, \mathbf{x}^* \rangle$. 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,

$$\begin{aligned} y_i &= \langle \mathbf{a}_i, \mathbf{x}^* \rangle + p_i^* R \\ &= y_{c_i} + p_i^* R. \end{aligned}$$

It is evident that if we can recover $\mathbf{p}^*$ successfully, we can calculate the true compressed measurements $\langle \mathbf{a}_i, \mathbf{x}^* \rangle$ and use them to reconstruct $\mathbf{x}^*$ with any sparse recovery algorithm such as CoSaMP [19] or basis pursuit [20–22].

2.3 Signal recovery

The major barrier to signal recovery is that the bin-index vector is unknown. In this section, we describe our algorithm to recover both $\mathbf{x}^*$ and $\mathbf{p}^*$, given the modulo measurements $\mathbf{y}$, measurement matrix $\mathbf{A}$, sparsity of underlying signal s , and the modulo parameter R . In this work, we rely on the assumption that our signal is sparse in a known domain with sparsity being s . Our algorithm *MoRAM* (Modulo Reconstruction with Alternating Minimization) comprises two stages: (i) an bin-index initialization stage and (ii) a descent stage via alternating minimization.

2.3.1 Bin-index initialization

As stated earlier, if we recover true bin-index $\mathbf{p}^*$ successfully, $\mathbf{x}^*$ can be recovered easily using any sparse recovery algorithm as we can obtain the true compressed measurements $\langle \mathbf{a}_i, \mathbf{x}^* \rangle$ from $\mathbf{p}^*$. Thus, in the absence of $\mathbf{p}^*$, we propose to estimate a fraction of the values from the $\mathbf{p}^*$ correctly. To understand the rationale for such a procedure, we will first try to understand the effect of the modulo operation on the linear measurements.

2.3.2 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 a normal distribution.

We can now divide the compressed observations $\mathbf{y}_c$ into two sets: $\mathbf{y}_{c,+}$, which contains all the non-negative observations with bin-index = 0, and $\mathbf{y}_{c,-}$, which contains all the negative observations 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 ($\mathbf{p}$):

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

The vector $\mathbf{p}^0$ 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. We should highlight that the procedure in Eq. 3 will succeed only for the specific case of modulo fold operations limited to two periods, one for the positive and one for the negative cycle.


Once we identify the initial values of bin-index for the modulo measurements, we can calculate corrected measurements as:

$$\mathbf{y}_c^0 = \mathbf{y} - \mathbf{p}^0 R. \quad (4)$$

Algorithm 1 MORAM**Inputs:** $\mathbf{y}$, $\mathbf{A}$, R **Output:** $\mathbf{x}^T$ $m, n \leftarrow \text{size}(\mathbf{A})$ *Initialization***for** $i = 0 : m$ **do** Calculate p_i^0 according to Eq. 3**end for***Descent***for** $t = 0 : T$ **do** $\mathbf{y}_c^t \leftarrow \mathbf{y} - \mathbf{p}^t R$ $\mathbf{x}^t \leftarrow JP(\mathbf{A}, \mathbf{y}_c^t)$ $\mathbf{p}^{t+1} \leftarrow \frac{1 - \text{sgn}(\mathbf{A}\mathbf{x}^t)}{2}$ **end for****2.3.3 Alternating minimization**

Using Eq. 3, we calculate the initial estimate of the bin-index $\mathbf{p}^0$ in which significant fraction of the total values are estimated correctly. Starting with $\mathbf{p}^0$, we calculate the estimates of $\mathbf{x}$ and $\mathbf{p}$ in an alternating fashion to converge to the original signal $\mathbf{x}^*$ and true bin-index $\mathbf{p}^*$.

With $\mathbf{p}^t$ being close to $\mathbf{p}^*$, 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 signal estimate $\mathbf{x}^t$. Therefore:



$$\mathbf{y}_c^t = \mathbf{y} - \mathbf{p}^t R, \quad (5)$$

$$\mathbf{x}^t = \arg \min_{\mathbf{x} \in \mathcal{M}_s} \|\mathbf{A}\mathbf{x} - \mathbf{y}_c^t\|_2^2, \quad (6)$$

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 a rather similar formulation would extend to cases where $\mathcal{M}$ denotes any other structured sparsity model [23, 24].

However, the bin-index estimation error, $\mathbf{d}^t = \mathbf{p}^t - \mathbf{p}^*$, even if small, would significantly impact the correction step that constructs $\mathbf{y}_c^t$ since 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 known as Justice Pursuit [10].

At a high level, Justice Pursuit tackles the problem of sparse signal recovery from the measurements that are corrupted by a sparse but large (unbounded) corruptions. Justice Pursuit leverages the fact that the corruptions are also sparse, and reformulates the problem to recover both the sparse signal and sparse corruptions together in the form of a concatenated sparse vector. In our case, the error $\mathbf{d}^t$ is sparse with sparsity $s_{dt} = \|\mathbf{d}^t\|_0$, and each erroneous element of $\mathbf{p}$ adds a corruption of magnitude R in $\mathbf{y}_c^t$. Following [10], we augment the measurement matrix $\mathbf{A}$ with an identity matrix $\mathbf{I}_{m \times m}$ and introduce an intermediate vector $\mathbf{u} \in \mathbb{R}^{n+m}$ to represent our measurements at iteration t as:

$$\mathbf{A}\mathbf{x}^* + R\mathbf{I}_m\mathbf{d}^t = \begin{bmatrix} \mathbf{A} & R\mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{x}^* \\ \mathbf{d}^t \end{bmatrix} = \begin{bmatrix} \mathbf{A} & R\mathbf{I} \end{bmatrix} \mathbf{u}, \quad (7)$$

and solve for the $(s + s_{dt})$ -sparse estimate $\hat{\mathbf{u}}$:

$$\begin{bmatrix} \hat{\mathbf{x}}^t \\ \hat{\mathbf{d}}^t \end{bmatrix} = \hat{\mathbf{u}} = \arg \min_{\mathbf{u}} \|\mathbf{u}\|_1 \quad s.t. \quad \begin{bmatrix} \mathbf{A} & R\mathbf{I} \end{bmatrix} \mathbf{u} = \mathbf{y}_c^t \quad (8)$$

Here, the signal estimate $\hat{\mathbf{x}}^t$ is obtained by selecting the first n elements of $\hat{\mathbf{u}}$, while an estimate of the corruptions can be obtained by selecting the last m elements of $\hat{\mathbf{u}}$. The problem in Eq. 8 can be solved by any stable sparse recovery algorithm such as CoSaMP or IHT; however, note that the sparsity of $\mathbf{d}^t$ is unknown, suggesting that greedy sparse recovery methods cannot be directly used without an additional hyper-parameter. Therefore, we employ basis pursuit [25] which does not heavily depend on a priori knowledge of the sparsity level.

We refer to the routine that solves the program in Eq. 8 using basis pursuit as *JP*. Given $\mathbf{A}, \mathbf{y}_c^t$, JP returns $\mathbf{x}^t$. Thus,

$$\mathbf{x}^{t+1} = JP(\mathbf{A}, \mathbf{y}_c^t). \quad (9)$$

Once the signal estimate $\mathbf{x}^t$ is obtained at each iteration of alternating minimization, we use it to calculate the value of the bin-index vector $\mathbf{p}^{t+1}$ as follows:

$$\mathbf{p}^{t+1} = \frac{\mathbf{1} - \text{sgn}(\mathbf{A}\mathbf{x}^t)}{2}. \quad (10)$$

Proceeding this way, we repeat the steps of sparse recovery (Eq. 8) and bin-index calculation (Eq. 10) in alternating fashion for T iterations. Under certain conditions (described in Section 2.4 below), our algorithm is able to achieve convergence to the true underlying signal.

2.4 Mathematical analysis

In this section, we provide correctness proofs for both steps of Algorithm 1. For the first stage, we derive an upper bound on the number of incorrect estimations in $\mathbf{p}^0$ obtained in the bin-index initialization step. This upper bound essentially provides an upper bound on the permissible sparsity of $\mathbf{d}^0$. For the second stage, we calculate a sufficient number of measurements required such that the augmented matrix used in the Justice Pursuit formulation in (8) satisfies the Restricted Isometry Property (RIP), which would in turn enable a recovery guarantee.

2.5 Bin-index initialization


In this step, we initialize the bin-index vector $\mathbf{p}^0$ according to Eq. 3. We can also quantify the number of correctly estimated bin-indices by calculating the area under the curve of the density plots of the measurements before and after the modulo operation. An illustration is provided in Fig. 3.

In this analysis, our goal is to characterize the distribution of total number of measurements for which we can estimate the correct bin-index through Eq. 3. Such a random variable is denoted by M_c . From M_c , we can calculate the sparsity of $\mathbf{d}^0$ as $\|\mathbf{d}^0\|_0 = m - M_c$. The following lemma presents a bound on the sparsity of $\mathbf{d}^0$.

Lemma 1 *Let the entries of the measurement matrix be generated as $\mathbf{A}_{ij} \sim \mathcal{N}(0, 1/m)$, and $\mathbf{y}$ be the modulo measurements obtained as per Eq. 1. Let M_c the random variable denoting the number of measurements for which the correct bin-indices are identified in the initialization method provided in Eq. 3. Then, with probability at least $1 - e^{-O(m\delta^2)}$:*

$$M_c > (1 - \delta)m \left(1 - 2 \frac{\sigma^2 \phi(R/2)}{(R/2)} \right).$$

Here, $\phi(\cdot)$ is a Gaussian density with mean $\mu = 0$ and variance $\sigma^2 = \|\mathbf{x}^*\|_2^2$.



Proof Observe that each element of $\mathbf{A}$ is i.i.d. standard normal, i.e., $\mu_{A_{ij}} = 0$ and $\sigma_{A_{ij}}^2 = 1$. Recall that

$$y_{c,i} = \langle \mathbf{a}_i, \mathbf{x}^* \rangle = \sum_{j=1}^n A_{ij} x_j^*.$$

Therefore, we have

$$y_{c,i} \sim \mathcal{N} \left(\mu = \sum_{j=1}^n x_j^* \mu_{A_{ij}} = 0, \sigma^2 = \sum_{j=1}^n x_j^{*2} \sigma_{A_{ij}}^2 \right).$$

Thus, each element of $\mathbf{y}_c$ follows a zero-mean Gaussian distribution with variance σ^2 . Let E_i be the event that the random variable $y_{c,i}$ lies in the interval $[-R/2, R/2]$; this event indicates that the corresponding measurement is appropriately corrected using Eq. 4. Clearly, E_i is a Bernoulli random variable with probability $q = P[-R/2 \leq y_{c,i} \leq R/2]$. Elementary probability calculations give us:

$$q = 1 - 2Q_{0,\sigma^2}(R/2),$$

where $Q_{0,\sigma^2}(\cdot)$ is the usual Q-function. This is not calculable in closed form; however, it can be lower bounded using the following identity (where $\phi_{0,\sigma^2}(\cdot)$ is a Gaussian density function with mean zero and variance σ^2):

$$Q_{0,\sigma^2}(t) < \sigma^2 \frac{\phi_{0,\sigma^2}(t)}{t}.$$

The random variable $M_c = \sum_{i=1}^m E_i$ denotes the number of corrected measurements. By an application of the Chernoff bound,

$$P(M_c \leq (1 - \delta)\mu') \leq e^{-\mu'\delta^2/2},$$

for any $\delta \in (0, 1)$, where μ' is the mean of M_c . Plugging in $\mu' = mq$ gives the desired result. $\square$

We now perform a theoretical analysis of the descent stage of our algorithm. We assume the availability of an initial estimate of bin-index vector $\mathbf{p}^0$ that is close to $\mathbf{p}^*$. In our case, our initialization step (in Alg. 1) provides such $\mathbf{p}^0$.

We perform alternating minimization (AltMin) as described in 1, starting with $\mathbf{p}^0$ calculated using Eq. 3. For simplicity, we limit our analysis of the convergence to only one AltMin iteration. In fact, according to our theoretical analysis, if initialized well enough, one iteration of AltMin suffices for exact signal recovery with sufficiently many measurements; however, in practice, we have observed that our algorithm performs better with multiple AltMin iterations.

Theorem 2 *Given the initial estimate of bin-index $\mathbf{p}^0$ obtained using Eq. 3, if the number of modulo measurements m satisfies:*

$$m \geq C_1 (\|\mathbf{x}^*\|_0 + m(1 - U + \delta U)) \log \left(\frac{n + m}{\|\mathbf{x}^*\|_0 + m(1 - U + \delta U)} \right),$$

then the first iteration of Algorithm 1 returns the true signal $\mathbf{x}^0$ with probability exceeding $1 - e^{-O(m\delta^2)}$ with small $\delta > 0$. Here, C_1 depends only on the RIP constant for the augmented measurement matrix $[\mathbf{A} \ \mathbf{I}]$, $q = 1 - 2Q_{0,\sigma^2}(R/2)$, and $U = 1 - 2\sigma^2 \frac{\phi(R/2)}{(R/2)}$.

Proof In the estimation step, Algorithm 1 recasts the problem of recovering the true signal $\mathbf{x}^*$ as a special case of sparse 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 Eq. 3, the remaining noise can be modeled as sparse corruptions $\mathbf{d}$ according to the formulation:

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

Here, the ℓ_0 -norm of $\mathbf{d}^0$ 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 $\|\mathbf{d}^0\|_0$ is small enough with respect to total number of measurements m ; thus, $\mathbf{d}^0$ 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 sufficiently large number of measurements are available. Denote $\|\mathbf{d}^0\|_0 = m - M_c$ as number of measurements for which the bin-index estimates were incorrect. Then, using Lemma 1, with probability at least $1 - e^{-O(m\delta^2)}$,

$$\begin{aligned} \|\mathbf{d}^0\|_0 &\leq m - (1 - \delta)mU \\ &\leq m(1 - U + \delta U), \text{ with } U = \left(1 - 2\sigma^2 \frac{\phi(R/2)}{(R/2)}\right). \end{aligned}$$


Algorithm 1 is essentially the Justice Pursuit (JP) formulation as described in [10]. Exact signal recovery from sparsely corrupted measurements is a well-studied problem with uniform recovery guarantees available in the existing literature. We use the guarantee proved in [10] for Gaussian observations, which states that provided enough measurements, the augmented matrix $[\mathbf{A} \ \mathbf{I}]$ satisfies the Restricted Isometry Property. As stated in [26], one can recover a sparse signal exactly by tractable ℓ_1 -minimization if the measurement matrix is known to satisfy the RIP. Thus, provided $m \geq C(\|\mathbf{x}^*\|_0 + \|\mathbf{d}^0\|_0) \log((n + m)/(\|\mathbf{x}^*\|_0 + \|\mathbf{d}^0\|_0))$, we invoke Theorem 1.1 from [10] and replace $\|\mathbf{d}^0\|_0$ with $m(1 - U + \delta U)$ as stated above to complete the proof. $\square$

From the theorem, we see that the number of measurements required for guaranteed recovery depends on the ratio of σ (standard deviation of the measurements) and R . In practical applications, choosing a sufficiently large R such that the interval $[-R, R]$ covers multiple standard deviations on both sides of origin enables successful recovery.

3 Experiments

In this section, we present the results of simulations of signal reconstruction using our algorithm. All numerical experiments were conducted using MATLAB R2020b on a Windows system with an Intel CPU and 16GB RAM. Our experiments explores the performance of the MoRAM algorithm on both synthetic data and 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 6 starting from 6 with a maximum value of 24. The non-zero elements of the test signal $\mathbf{x}^*$ are generated using a 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 $\mathcal{N}(0, 1/m)$. The number of measurements m is varied from $m = 100$ to $m = 1000$ in steps of 100.

Using $\mathbf{A}$, $\mathbf{x}^*$, and R , we first obtain the compressed modulo measurements $\mathbf{y}$ by passing the signal through forward model described by Eq. 2. For reconstruction, algorithm 1 is employed. We plot the variation of the relative reconstruction error $\left(\frac{\|\mathbf{x}^* - \mathbf{x}^{\dagger}\|}{\|\mathbf{x}^*\|}\right)$ with the 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. Figure 4a and b illustrate the performance of our algorithm for two values of R respectively. Additionally, in Fig. 5, we also show reconstruction performance results for higher sparsity values $s = 20, 30, 40, 50$ as they vary with the number of measurements. It is evident that for each combination of R and s , our algorithm recovers the true signal (with zero relative error) provided enough measurements. In all such cases, the minimum number of measurements required for exact recovery is well below the ambient dimension (n) of the underlying signal.





From the analysis provided in the previous section, it is clear that the success of our algorithm depends on the ratio of the standard deviation σ of the Gaussian measurements and modulo period R . When the matrix $\mathbf{A}$ is correctly normalized, σ directly depends on the Euclidean norm of the signal, which in turn depends on the sparsity. Thus, as discussed earlier, if R is chosen to be large enough, success of our algorithm is guaranteed.

To put this into perspective, we provide additional results where we fix the number of measurements m and sparsity s , and vary the modulo parameter R . These values suggest that even when the R becomes small and the ratio of σ and R increases, our algorithm is able to recover the underlying signal perfectly as shown in Fig. 6. As R increases, recovery becomes easier to achieve for our algorithm. It also shows that the performance of our algorithm decays gracefully when varying R .

In practical scenarios, we may encounter cases where a few measurements exceed the modulo period R . If we define γ as $\max(|\langle \mathbf{A}, \mathbf{x}^* \rangle|)$, then the value of γ would vary across each Monte Carlo run of our experiment (due to the randomness in the measurement matrix $\mathbf{A}$), and we may encounter cases with $\gamma > R$ as well. To analyze such cases, we provide an additional experiment where we fix the number of measurements m to 400, sparsity s to 12, and modulo parameter R to 3.2 and run our recovery algorithm multiple times with different realizations of $\mathbf{A}$. We note down the values of γ during these experiments and note summary statistics as $mean_\gamma = 3.45$, $variance_\gamma = 0.2065$, $max_\gamma = 4.41$, $min_\gamma = 2.82$ across 50 random trials. In *all* these 50 cases, our algorithm recovered the true signal perfectly, i.e., resulting in zero relative error in each case. As the γ varies widely and also takes values higher than R , perfect recovery achieved in all cases shows that our algorithm is robust to measurements with magnitudes exceeding R , as far as such measurements are low in number (as guaranteed by Gaussian tail bounds on $\langle \mathbf{A}, \mathbf{x}^* \rangle$). In this sense, our algorithm follows a “graceful decay” with respect to the assumption that most of the measurements are contained in the interval $[-R, R]$ in this sense.

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 a 128×128 ($n = 16384$) image (Fig. 7a),





and we use a wavelet transform (with Haar wavelets) 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. 7 (bottom), for $m = 6000$, the algorithm produces near-perfect recovery for all 3 values of R with high PSNR. Here, let us note that the blocky artifacts appearing in the recovered image are actually due to sparsification of the original image using the Haar wavelet transform. Since we use $s = 800$ to obtain the ground truth image, the ground truth itself contains some compression artifacts as depicted in Fig. 7 (a, bottom). The PSNR values are calculated with respect to the sparse ground truth image and not with respect to the original image as our algorithm aims to recover only the former. We expect the effect of compression artifacts to decrease with a better choice of sparsity basis for the underlying image.

4 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. Also, instead of relying on a sparsity prior for compressed recovery, employing richer priors such as GANs [27–29] 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 [30] can be an interesting problem for future study.

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, and the anonymous reviewers for helpful comments.

Authors' 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.

Funding

This work was supported by grants CCF-1566281, CAREER CCF-1750920/2005804, and CCF-1815101 from the National Science Foundation; a faculty fellowship grant from the Black and Veatch Foundation; and a GPU grant from the NVIDIA Corporation.

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.

Declarations

Competing interests

The authors declare no competing interests.

Author details

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

Received: 16 July 2020 Accepted: 19 February 2021

Published online: 06 April 2021

References

1. V. Shah, C. Hegde, in *2019 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, Signal reconstruction from modulo observations, (2019), pp. 1–5
2. J. Rhee, Y. Joo, Wide dynamic range CMOS image sensor with pixel level ADC. *Electron. Lett.* **39**, 360–361 (2010)
3. S. Kavusi, A. El Gamal, in *Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications V*, Quantitative study of high-dynamic-range image sensor architectures, vol. 5301 (International Society for Optics and Photonics, 2004), pp. 264–276
4. K. Sasagawa, T. Yamaguchi, M. Haruta, Y. Sunaga, H. Takehara, H. Takehara, T. Noda, T. Tokuda, J. Ohta, An implantable CMOS image sensor with self-reset pixels for functional brain imaging. *IEEE Trans. Electron Devices*. **63**(1), 215–222 (2016)
5. T. Yamaguchi, H. Takehara, Y. Sunaga, M. Haruta, M. Motoyama, Y. Ohta, T. Noda, K. Sasagawa, T. Tokuda, J. Ohta, Implantable self-reset CMOS image sensor and its application to hemodynamic response detection in living mouse brain. *Jpn. J. Appl. Phys.* **55**(45), 04–02 (2016)
6. A. Bhandari, F. Krahmer, R. Raskar, in *2017 International Conference on Sampling Theory and Applications (SampTA)*, On unlimited sampling, (2017), pp. 31–35
7. H. Zhao, B. Shi, C. Fernandez-Cull, S. Yeung, R. Raskar, in *IEEE International Conference on Computational Photography (ICCP)*, Unbounded high dynamic range photography using a modulo camera, (2015)

8. V. Shah, M. Soltani, C. Hegde, in *Proceedings of the Asilomar Conference on Signals, Systems, and Computers*, Reconstruction from periodic nonlinearities, with applications to HDR imaging (IEEE, 2017), pp. 863–867
9. P. Netrapalli, P. Jain, S. Sanghavi, in *Proceedings of the Advances in Neural Information Processing Systems (NIPS)*, Phase retrieval using alternating minimization, (2013), pp. 2796–2804
10. J. Laska, M. Davenport, R. Baraniuk, *Exact signal recovery from sparsely corrupted measurements through the pursuit of justice*, (2009), pp. 1556–1560
11. J. Bioucas-Dias, G. Valadao, Phase unwrapping via graph cuts. *IEEE Trans. Image Proc.* **16**(3), 698–709 (2007)
12. F. Lang, T. Plötz, S. Roth, in *German Conference on Pattern Recognition*, Robust multi-image HDR reconstruction for the modulo camera, (2017), pp. 78–89. <https://arxiv.org/pdf/1707.01317v1.pdf>
13. M. Soltani, C. Hegde, in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Stable recovery of sparse vectors from random sinusoidal feature maps, (2017), pp. 6384–6388
14. O. Ordentlich, G. Tabak, P. K. Hanumolu, A. C. Singer, G. W. Wornell, A modulo-based architecture for analog-to-digital conversion. *IEEE J. Sel. Top. Signal Process.* **12**(5), 825–840 (2018)
15. S. Rudresh, A. Adiga, B. A. Shenoy, C. S. Seelamantula, in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Wavelet-based reconstruction for unlimited sampling (IEEE, 2018), pp. 4584–4588
16. M. Cucuringu, H. Tyagi, in *International Conference on Artificial Intelligence and Statistics*, On denoising modulo 1 samples of a function, (2018), pp. 1868–1876
17. A. Bhandari, F. Krahmer, R. Raskar, in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Unlimited sampling of sparse signals, (2018), pp. 4569–4573
18. O. Musa, P. Jung, N. Goertz, in *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, Generalized approximate message passing for unlimited sampling of sparse signals, (2018), pp. 336–340
19. D. Needell, J. Tropp, CoSaMP: iterative signal recovery from incomplete and inaccurate samples. *Commun. ACM.* **53**(12), 93–100 (2010)
20. S. Chen, D. Donoho, M. Saunders, Atomic decomposition by basis pursuit. *SIAM review.* **43**(1), 129–159 (2001)
21. E. van den Berg, M. P. Friedlander, SPGL1: a solver for large-scale sparse reconstruction (2007). <http://www.cs.ubc.ca/labs/scl/spgl1>
22. E. van den Berg, M. P. Friedlander, Probing the Pareto frontier for basis pursuit solutions. *SIAM J. Sci. Comput.* **31**(2), 890–912 (2008)
23. R. Baraniuk, V. Cevher, M. Duarte, C. Hegde, Model-based compressive sensing. *IEEE Trans. Inf. Theory.* **56**, 1982–2001 (2010)
24. C. Hegde, P. Indyk, L. Schmidt, Fast algorithms for structured sparsity. *Bull. EATCS.* **3**(117) (2015)
25. E. J. Candès, *et al*, in *Proceedings of the International Congress of Mathematicians*, Compressive sampling, vol. 3, (Madrid, Spain, 2006), pp. 1433–1452
26. E. J. Candès, The restricted isometry property and its implications for compressed sensing. *C. R. Acad. Sci. I.* **346**, 589–592 (2008)
27. A. Bora, A. Jalal, E. Price, A. Dimakis, in *International Conference on Machine Learning*, Compressed sensing using generative models, (2017), pp. 537–546
28. V. Shah, C. Hegde, in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, Solving linear inverse problems using GAN priors: an algorithm with provable guarantees, (2018), pp. 4609–4613
29. G. Jagatap, C. Hegde, in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, High dynamic range imaging using deep image priors, (2020), pp. 9289–9293
30. G. Jagatap, Z. Chen, S. Nayer, C. Hegde, N. Vaswani, Sample efficient fourier ptychography for structured data. *IEEE Trans. Comput. Imaging.* **6**, 344–357 (2020)

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)