

# SIGNAL RECONSTRUCTION FROM MODULO OBSERVATIONS

Viraj Shah and Chinmay Hegde

ECpE Department, Iowa State University, Ames, IA, 50010

## ABSTRACT

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

**Index Terms**— Sparse recovery, nonlinear observation models, modulo sensors, imaging applications.

## 1. INTRODUCTION

### 1.1. Motivation

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

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

$$y_i = \mod(\langle \mathbf{a}_i \cdot \mathbf{x}^* \rangle, R) \quad i = \{1, 2, \dots, m\}, \quad (1.1)$$

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

Recently, the use of a novel imaging sensor that wraps the data in a periodical manner has been shown to overcome certain hardware limitations of typical imaging systems [1, 2, 3, 4]. Many image acquisition systems suffer from the problem of limited dynamic range; however, real-world signals can contain a large range of intensity levels, and if tuned incorrectly, most intensity levels can lie in the saturation region of the sensors, causing loss of information through

This work was supported by grants CCF-1566281 and CAREER CCF-1750920 from the National Science Foundation, and a faculty fellowship from the Black and Veatch Foundation, and a GPU grant from the NVIDIA Corporation. We thank Praneeth Narayanamurthy, Gauri Jagatap, and Thanh Nguyen for helpful comments.

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.

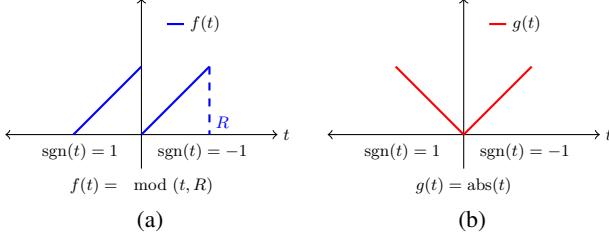
The standard solution to this issue is to improve sensor dynamic range via enhanced hardware; this, of course, can be expensive. An intriguing alternative is to deploy special digital *modulo* sensors [5, 6, 7, 8]. As the name suggests, such a sensor wraps each signal measurement around a scalar parameter  $R$  that reflects the dynamic range. However, this also makes the forward model (1.1) highly nonlinear and the reconstruction problem highly ill-posed. The approach of [1, 2] relies on the assumption that the underlying signal is bandlimited and time-continuous, deeming them unfit for being used in compressive sensing regime. Moreover, it assumes *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 (1.1) 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: one in the positive half and one in the negative half as shown in the Fig. 1(a). Such a simplified version of the modulo function already inherits much of the challenging aspects of the original function. Intuitively, this simplification requires that the norm of the target signal is not too large.

### 1.2. Our contributions

In this paper, we propose a recovery algorithm for exact reconstruction of signals from modulo measurements of the form (1.1). 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 attracted renewed interest of late with many solutions available in literature [9, 10, 11, 12, 13], including the cases where underlying signal is sparse [14, 15, 16, 17, 18, 19, 20]. In phase retrieval, we are given magnitude measurements of  $\langle \mathbf{a}_i \cdot \mathbf{x}^* \rangle$  and are tasked with reconstructing  $\mathbf{x}^*$ . While these two different class of problems appear different at face value, the common theme is the need of undoing the effect of a piecewise linear transfer function applied to the observations. See Fig. 1 for a comparison.



**Fig. 1:** Comparison between (a) modulo function ( $f(t) = \text{mod}(t, R)$ ); and (b) absolute value function ( $g(t) = \text{abs}(t)$ ).

However, several essential differences between the two problems restrict us from using algorithms for phase retrieval as-is for the modulo reconstruction problem. The absolute value function is a multiplicative transfer function (multiplication with the sign of the measurements), so in phase retrieval, a wrongly estimated phase induces an error that increases *linearly* with the magnitude of each measurement. On the other hand, the modulo function adds a constant value ( $R$ ) to negative inputs, thus 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. Apart from experimental results and comparison with existing algorithm, we also provide an analytical guarantee suggesting that such a recovery can be performed using an (essentially) optimal number of observations, given the initial estimate that lies relatively close to the underlying signal. We provide an initialization method to obtain such initial estimate, and present experiments depict the effectiveness of our approach. To the best of our knowledge we are the first to pursue this type of approach for modulo recovery problems, distinguishing us from previous work [2, 1].

### 1.3. Techniques

The basic approach in MoRAM is similar to several recent non-convex phase retrieval algorithms [20, 21, 22, 23, 24, 19]. In the first step we identify a good initial guess  $\mathbf{x}^0$  for our signal that lies relatively close to the true vector  $\mathbf{x}^*$ . *Spectral initialization*, a commonly used technique for phase retrieval [18], does not work in our case due to the markedly different nature of the modulo function. We introduce a novel approach of measurement *correction* by comparing 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 a initial vector that is close to the true signal.

In the second step, we follow an alternating minimization approach (e.g. [18, 25]) that estimates the signal and the measurement signs 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 appealing to a *robust* form of alternating minimization (specifically, the Justice Pursuit algorithm [26]). We prove that AltMin with 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.

### 1.4. Prior work

Due to space, we defer a more thorough discussion of prior work to an extended version of this paper. 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 bandlimitedness 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 sampling complexity. The only setup that allows the non-uniform sampling grid is multishot UHDR [2], under some reasonable modifications. We provide experiments comparing our algorithm with multishot UHDR in experiments section. In recent works, [27] proposed unlimited sampling algorithm for sparse signals. Similar to [1], 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 [28], modulo recovery from Gaussian random measurements is considered, however, it assumes the true signal to be distributed as mixed Bernoulli-Gaussian distribution, which is impractical in real world imaging scenarios.

## 2. MATHEMATICAL MODEL

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.). The cardinality of set  $S$  is given by  $\text{card}(S)$ . The signum function is defined as  $\text{sgn}(x) := \frac{x}{|x|}, \forall x \in \mathbb{R}, x \neq 0$ , with  $\text{sgn}(0) = 1$ . The projection of vector  $\mathbf{x} \in \mathbb{R}^n$  onto a set of coordinates  $S$  is represented as  $\mathbf{x}_S \in \mathbb{R}^n$ ,  $\mathbf{x}_{S_j} = \mathbf{x}_j$  for  $j \in S$ , and 0 elsewhere.

As depicted in Fig. 1(a), we consider the modulo operation within 2 periods. If we write the modulo function of Fig. 1(a) in terms of a signum function, then the measurement model of Eq. 1.1 becomes,

$$y_i = \langle \mathbf{a}_i \cdot \mathbf{x}^* \rangle + \left( \frac{1 - \text{sgn}(\langle \mathbf{a}_i \cdot \mathbf{x}^* \rangle)}{2} \right) R, \quad i = \{1, \dots, m\}. \quad (2.1)$$

If we divide the number line in 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  $\text{sgn}(t) = -1$ . We denote such bin-index vector as  $\mathbf{p} \in \mathbb{R}^m$ . Each element of the true bin-index vector  $\mathbf{p}^*$  is given as,

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

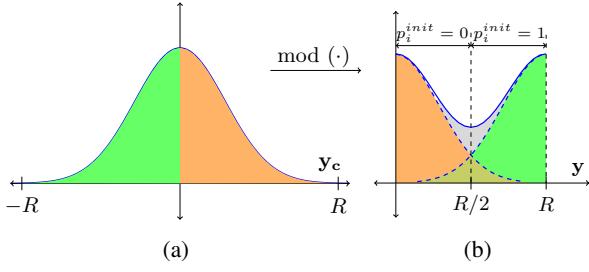
If we ignore the presence of modulo operation in above formulation, then it reduces to a standard compressive sensing problem. In that case, the compressed measurements  $y_{c_i}$  would just be equal to  $\langle \mathbf{a}_i \cdot \mathbf{x}^* \rangle$ . While we have access only to the compressed 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 [29] or basis-pursuit [30, 31, 32].

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

	Unlimited Sampling [1]	OLS Method [4]	multishot UHDR [2]	<b>MoRAM (our approach)</b>
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



**Fig. 2:** Density plots of (a)  $y_c = Ax^*$ ; (b)  $y = \text{mod}(Ax^*)$ .

### 3. ALGORITHM AND MAIN RESULTS

Given  $y, A, s, R$ , our approach recovers  $x^*$  and  $p^*$  in two steps: (i) an initialization step, and (ii) descent step via Alt-Min.

#### 3.1. Initialization by re-calculating the measurements

Similar to other non-convex approaches, MoRAM also requires an initial estimate  $x^0$  that is close to the true signal  $x^*$ .

To do so, we observe the density plots of the  $y_c = Ax^*$  and  $y = \text{mod}(Ax^*)$  as shown in Fig. 2 (a) and (b) respectively. Note that the compressed measurements  $y_c$  follow the standard normal distribution, as  $A$  is Gaussian random matrix. These plots essentially depict the distribution of our observations *before* and *after* the modulo operation.

With reference to Fig. 2(a), we divide the compressed observations  $y_c$  in two sets:  $y_{c,+}$  contains all the non-negative observations (orange) with bin-index= 0, while  $y_{c,-}$  contains all the negative ones (green) with bin-index= 1.

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

$$p_i^{\text{init}} = \begin{cases} 0, & \text{if } 0 \leq y_i < R/2 \\ 1, & \text{if } R/2 \leq y_i \leq R \end{cases} \quad (3.1)$$

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

Once we identify the initial values of bin-index for the modulo

#### Algorithm 1 MoRAM

```

Inputs:  $y, A, s, R$ ; Output:  $x^T$ 
 $m, n \leftarrow \text{size}(A)$ 
Initialization
for  $i = 0 : m$  do
    Calculate  $p_i^{\text{init}}$  according to Eq. 3.1.
end for
Calculate  $y_c^{\text{init}}$  according to Eq. 3.2.
 $x^0 \leftarrow H_s \left( \frac{1}{N} \sum_{i=1}^N y_{c,i}^{\text{init}} a_i \right)$ 
Alternating Minimization
for  $t = 0 : T$  do
     $p^t \leftarrow \frac{1 - \text{sgn}(\langle A \cdot x^t \rangle)}{2}$ 
     $y_c^t \leftarrow y - p^t R$ 
     $x^{t+1} \leftarrow JP \left( \frac{1}{\sqrt{m}} [A \quad I], \frac{1}{\sqrt{m}} y_c^t, [x^t \quad p^t]^\top \right)$ .
end for

```

measurements, we can calculate corrected measurements as,

$$y_c^{\text{init}} = y + p^{\text{init}} R. \quad (3.2)$$

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

$$x^0 = H_s \left( \frac{1}{N} \sum_{i=1}^N y_{c,i}^{\text{init}} a_i \right), \quad (3.3)$$

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.

#### 3.2. Alternating minimization

In the descent step, starting with  $x^0$ , we calculate the estimates of  $p$  and  $x$  in alternating fashion to converge to the original signal  $x^*$ . At each iteration of our Alternating Minimization, we use the current estimate of the signal  $x^t$  to get the value of the bin-index vector  $p^t$  as following:

$$p^t = \frac{1 - \text{sgn}(\langle A \cdot x^t \rangle)}{2}. \quad (3.4)$$

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

However, even the small error  $d^t = p^t - p^*$  would reflect heavily in the calculation of  $x^t$ , as each incorrect bin-index would add a noise of the magnitude  $R$  in  $y_c^t$ . The typical sparse recovery algorithms are not robust enough to cope up with such gross errors in  $y_c^t$  [26]. To tackle this issue, we employ the justice pursuit based

formulation which is specifically robust towards grossly corrupted measurements. We consider the fact that the nature of 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  $\mathbf{y}_c^t$ . Thus we solve the augmented optimization problem:

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

However, the sparsity of  $\mathbf{d}^t(s_{dt})$  is unknown, thus we employ basis pursuit [30, 31, 32] that doesn't rely on sparsity. The robust formulation of basis pursuit is referred as Justice Pursuit (JP) [26] in the literature, specified in Eq. 3.5.

$$\mathbf{x}^{t+1} = JP\left(\frac{1}{\sqrt{m}} [\mathbf{A} \ \mathbf{I}], \frac{1}{\sqrt{m}} \mathbf{y}_c^t, [\mathbf{x}^t \ \mathbf{p}^t]^\top\right). \quad (3.5)$$

We repeat the steps of bin-index calculation (as in Eq. 3.4) and sparse recovery (Eq. 3.5) alternatively for  $T$  iterations.

### 3.3. Theoretical results

We now provide theoretical guarantee for convergence of alternating minimization in our algorithm. For brevity, here we omit the proof, which is discussed in detail in an extended version [33].

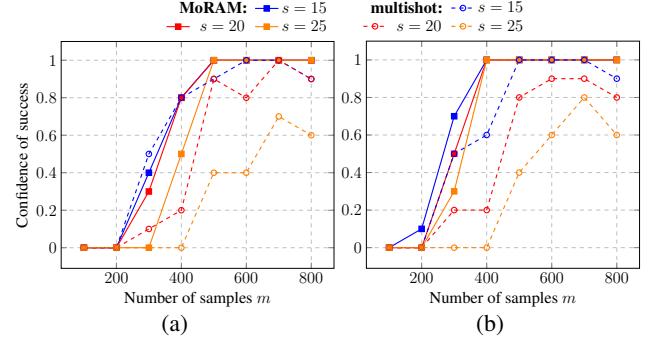
We assume the availability of an initial estimate  $\mathbf{x}^0$  that is close to  $\mathbf{x}^*$ , i.e.  $\|\mathbf{x}^0 - \mathbf{x}^*\|_2 \leq \delta \|\mathbf{x}^*\|_2$ . In our case, our initialization step (in Alg. 1) provide such  $\mathbf{x}^0$ . For simplicity, we limit our analysis of the descent to only one iteration of Alternating Minimization in our algorithm. In fact, as proven in our analysis, theoretically only one iteration of AltMin with JP is required for exact signal recovery. Contrast to that, in practice we observe that our algorithm requires more than one AltMin iterations to converge to the optimum solution. We build our proofs on the results from [34, 35, 36].

**Theorem 3.1** (Guarantee for Descent). *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 alternating minimization in Algorithm 1 is exactly equal to the true signal  $\mathbf{x}^*$  with probability at least  $1 - K \exp(-cm) - \eta$ , with  $\gamma$  being a positive fraction,  $K$  and  $c$  being numerical constants.*

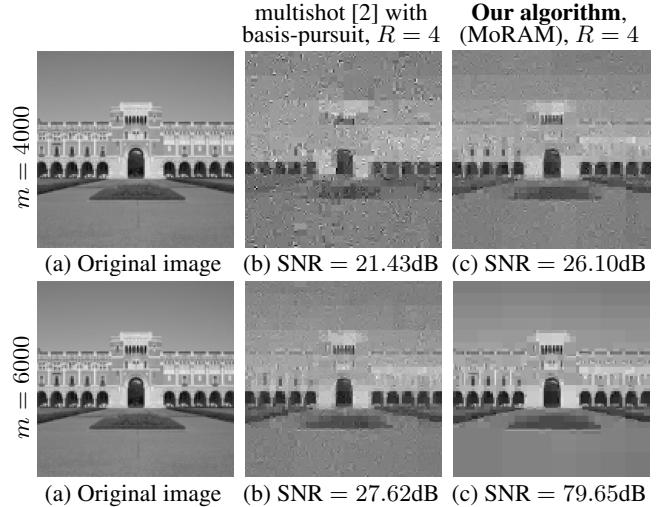
## 4. NUMERICAL EXPERIMENTS

We numerically validate the MoRAM algorithm on both synthetic data and real image data. We also provide comparisons with multishot UHDR [2], as it is the only approach (to our knowledge) that is amenable to reconstruction from random compressive observations. We generate a synthetic sparse signal  $\mathbf{x}^* \in \mathbb{R}^n$  with  $n = 1000$ . The non-zero elements of  $\mathbf{x}^*$  are generated from  $\mathcal{N}(0, 1)$  and normalized such that  $\|\mathbf{x}^*\|_2 = 1$ . The number of measurements  $m$  is varied from  $m = 100$  to  $m = 800$  in steps of 100.

We first obtain the compressed modulo measurements  $\mathbf{y} = \text{mod } \mathbf{y}_c = \mathbf{A}\mathbf{x}^*$  using the forward model described in Eq. 2.1. We reconstruct the signal from  $\mathbf{y}$  using MoRAM algorithm with  $T = 15$  and plot the confidence of perfect recovery (ratio of successful trials) across 10 such independent trials vs. number of measurements  $m$ . In comparison, we use multishot UHDR algorithm to undo the effect of modulo operation on  $\mathbf{y}$  to obtain  $\mathbf{y}_c = \mathbf{A}\mathbf{x}^*$ , and recover  $\mathbf{x}^*$  using basis-pursuit. In both cases, we keep the total number of measurements same, as multishot recovery requires oversampled ( $2 \times$  in this



**Fig. 3:** Comparison of confidence of success (ratio of successful trials) vs no. of measurements ( $m$ ) between MoRAM and multishot UHDR [2] with  $n = 1000$ , and (a)  $R = 4$ ; (b)  $R = 4.5$ .



**Fig. 4:** Sparse reconstructions ( $s = 800$ ) of original Lovett Hall image ( $n = 16,384$ ); with  $m = 4000$  and  $m = 6000$  modulo measurements for  $R = 4$  (b) using multishot UHDR followed by basis pursuit; (c) using MoRAM. MoRAM provides superior results with higher PSNR.

case) measurements. In Fig. 3 we depict such plots for 2 values of  $R = \{4, 4.5\}$  and 3 values of  $s = \{15, 20, 25\}$ . It is evident that for each combination of  $R$  and  $s$ , our algorithm achieves confidence of 1 to give the exact recovery of the true signal (zero relative error) consistently in all trials provided enough number of measurements, while multishot fails to maintain high confidence of recovery even with higher  $m$ . In all such cases, the minimum number of measurements required for exact recovery for MoRAM are well below the ambient dimension ( $n$ ) of the underlying signal.

Further, for experiments on real data, we obtain a sparse representation of the  $128 \times 128$  ( $n = 16384$ ) Lovett Hall image (Fig. 4(a)) using the wavelet transform (with Haar wavelets). We set  $s = 800$ . We reconstruct the image with MoRAM using  $m = 4000$  and  $m = 6000$  compressed modulo measurements, for  $R = 4$ . In comparison, in Fig. 4(b) we depict the images reconstructed using multishot followed by basis pursuit keeping the total measurements,  $m$ , the same. Images reconstructed using MoRAM exhibit superior performance compared to the multishot UHDR setup. As shown in Fig. 4(c, bottom), for  $m = 6000$ , our algorithm produces perfect recovery with high recovery PSNR.

## 5. REFERENCES

- [1] A. Bhandari, F. Krahmer, and R. Raskar, “On unlimited sampling,” *Proc. Sampling Theory and Applications (SampTA)*, pp. 31–35, 2017.
- [2] H. Zhao, B. Shi, C. Fernandez-Cull, S. Yeung, and R. Raskar, “Unbounded high dynamic range photography using a modulo camera,” in *Intl. Conf. on Comp. Photography (ICCP)*, 2015.
- [3] V. Shah, M. Soltani, and C. Hegde, “Reconstruction from periodic nonlinearities, with applications to hdr imaging,” in *Proc. Asilomar Conf. Signals, Systems, and Computers*. IEEE, 2017, pp. 863–867.
- [4] M. Cucuringu and H. Tyagi, “On denoising modulo 1 samples of a function,” in *Proc. Int. Conf. Art. Intell. Stat. (AISTATS)*, 2018.
- [5] J. Rhee and Y. Joo, “Wide dynamic range cmos image sensor with pixel level adc,” *Electron. Lett.*, vol. 39, pp. 360–361, 2010.
- [6] S. Kavusi and A. El Gamal, “Quantitative study of high-dynamic-range image sensor architectures,” in *Sensors and Camera Systems for Sci., Indust., and Digi. Photography Applications V*. Intl. Soc. for Optics and Photonics, 2004, vol. 5301, pp. 264–276.
- [7] K. Sasagawa, T. Yamaguchi, M. Haruta, Y. Sunaga, H. Takehara, H. Takehara, T. Noda, T. Tokuda, and J. Ohta, “An implantable cmos image sensor with self-reset pixels for functional brain imaging,” *IEEE Trans. on Electron Devices*, vol. 63, no. 1, pp. 215–222, 2016.
- [8] T. Yamaguchi, H. Takehara, Y. Sunaga, M. Haruta, M. Motoyama, Y. Ohta, T. Noda, K. Sasagawa, T. Tokuda, and J. Ohta, “Implantable self-reset cmos image sensor and its application to hemodynamic response detection in living mouse brain,” *Japanese J. of Appl. Physics*, vol. 55, no. 04EM02, pp. 04EM02, 2016.
- [9] E. Candes, T. Strohmer, and V. Voroninski, “Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming,” *Comm. Pure Appl. Math.*, vol. 66, no. 8, pp. 1241–1274, 2013.
- [10] D. Gross, F. Krahmer, and R. Kueng, “Improved recovery guarantees for phase retrieval from coded diffraction patterns,” *Appl. Comput. Harmon. Anal.*, vol. 42, no. 1, pp. 37–64, 2017.
- [11] E. Candes, X. Li, and M. Soltanolkotabi, “Phase retrieval from coded diffraction patterns,” *Appl. Comput. Harmon. Anal.*, vol. 39, no. 2, pp. 277–299, 2015.
- [12] S. Bahmani and J. Romberg, “Phase retrieval meets statistical learning theory: A flexible convex relaxation,” in *Proc. Int. Conf. Art. Intell. Stat. (AISTATS)*, 2016.
- [13] T. Goldstein and C. Studer, “Phasemax: Convex phase retrieval via basis pursuit,” *IEEE Trans. Inform. Theory*, vol. 64, pp. 2675–2689, 2018.
- [14] H. Ohlsson, A. Yang, R. Dong, and S. Sastry, “Cprl—an extension of compressive sensing to the phase retrieval problem,” in *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, 2012, pp. 1367–1375.
- [15] X. Li and V. Voroninski, “Sparse signal recovery from quadratic measurements via convex programming,” *SIAM J. on Math. Analysis*, vol. 45, no. 5, pp. 3019–3033, 2013.
- [16] S. Bahmani and J. Romberg, “Efficient compressive phase retrieval with constrained sensing vectors,” in *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, 2015, pp. 523–531.
- [17] K. Jaganathan, S. Oymak, and B. Hassibi, “Recovery of sparse 1-d signals from the magnitudes of their fourier transform,” in *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*. IEEE, 2012, pp. 1473–1477.
- [18] P. Netrapalli, P. Jain, and S. Sanghavi, “Phase retrieval using alternating minimization,” in *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, 2013, pp. 2796–2804.
- [19] T. Cai, X. Li, Z. Ma, et al., “Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow,” *Ann. Stat.*, vol. 44, no. 5, pp. 2221–2251, 2016.
- [20] G. Wang, L. Zhang, G. Giannakis, M. Akçakaya, and J. Chen, “Sparse phase retrieval via truncated amplitude flow,” *arXiv preprint arXiv:1611.07641*, 2016.
- [21] G. Wang and G. Giannakis, “Solving random systems of quadratic equations via truncated generalized gradient flow,” in *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, 2016, pp. 568–576.
- [22] E. Candes, X. Li, and M. Soltanolkotabi, “Phase retrieval via wirtinger flow: theory and algorithms,” *IEEE Trans. Inform. Theory*, vol. 61, no. 4, pp. 1985–2007, 2015.
- [23] H. Zhang and Y. Liang, “Reshaped wirtinger flow for solving quadratic system of equations,” in *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, 2016, pp. 2622–2630.
- [24] Y. Chen and E. Candes, “Solving random quadratic systems of equations is nearly as easy as solving linear systems,” in *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, 2015, pp. 739–747.
- [25] G. Jagatap and C. Hegde, “Fast, sample-efficient algorithms for structured phase retrieval,” in *Proc. Adv. in Neural Inf. Proc. Sys. (NIPS)*, 2017.
- [26] J. Laska, M. Davenport, and R. Baraniuk, “Exact signal recovery from sparsely corrupted measurements through the pursuit of justice,” in *Proc. Asilomar Conf. Signals, Systems, and Computers*, 2009, pp. 1556–1560.
- [27] A. Bhandari, F. Krahmer, and R. Raskar, “Unlimited sampling of sparse signals,” *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, pp. 4569–4573, 2018.
- [28] O. Musa, P. Jung, and N. Goertz, “Generalized approximate message passing for unlimited sampling of sparse signals,” *Proc. of IEEE Global Conf. on Signal and Inform. Processing (GlobalSIP)*, pp. 336–340, 2018.
- [29] D. Needell and J. Tropp, “Cosamp: iterative signal recovery from incomplete and inaccurate samples,” *Comm. of the ACM*, vol. 53, no. 12, pp. 93–100, 2010.
- [30] S. Chen, D. Donoho, and M. Saunders, “Atomic decomposition by basis pursuit,” *SIAM review*, vol. 43, no. 1, pp. 129–159, 2001.
- [31] E. van den Berg and M. P. Friedlander, “SPGL1: A solver for large-scale sparse reconstruction,” June 2007, <http://www.cs.ubc.ca/labs/scl/spgl1>.
- [32] E. van den Berg and M. P. Friedlander, “Probing the pareto frontier for basis pursuit solutions,” *SIAM J. on Sci. Computing*, vol. 31, no. 2, pp. 890–912, 2008.
- [33] V. Shah and C. Hegde, “Signal reconstruction from modulo observations,” *arXiv preprint arXiv:1812.00557*, 2018.
- [34] R. Vershynin, “Introduction to the non-convex analysis of random matrices,” *arXiv preprint arXiv:1011.3027*, 2010.
- [35] X. Li, “Compressed sensing and matrix completion with constant proportion of corruptions,” *Const. Approx.*, vol. 37, no. 1, pp. 73–99, 2013.
- [36] L. Jacques, J. Laska, P. Boufounos, and R. Baraniuk, “Robust 1-Bit compressive sensing via binary stable embeddings of sparse vectors,” *IEEE Trans. Inform. Theory*, vol. 59, no. 4, pp. 2082–2102, 2013.