# THE USE OF UNIT NORM TIGHT MEASUREMENT MATRICES FOR ONE-BIT COMPRESSED SENSING

*Cristian Rusu, Nuria González-Prelcic*

University of Vigo
Email: {crusu,nuria}@gts.uvigo.es

*Robert W. Heath Jr.*

The University of Texas at Austin
Email: rheath@utexas.edu

**ABSTRACT**

In this paper we analyze the mean squared error (MSE) for one-bit compressed sensing schemes based on measurement matrices that correspond to unit norm tight frames. We show that, as in the unquantized case, sensing with unit norm tight frames improves the MSE in the reconstruction of sparse vectors from one-bit measurements using $\ell_1$ and thresholding algorithms. From our analytical and experimental results we conclude that when implementing one-bit compressed sensing schemes with fixed measurement matrices unit norm tight frames are the measurements of choice.

***Index Terms***— one-bit compressed sensing, unit norm tight frames

## 1. INTRODUCTION

Analog-to-digital converters (ADCs) are being pushed to their limits by applications involving new signal processing and communication systems, designed with ever higher acquisition resolutions and wider bandwidths. Developing high rate, high resolution ADCs is still a technological challenge. In addition, power consumption in the ADCs increases quadratically with sampling frequency [1].

One approach to overcoming this consumption limitation is using low resolution and even one-bit ADCs. A one-bit ADCs is built from a simple comparator, which can work at high sampling rates with a very low power consumption. Intuitively, recovering a signal from the sign of its samples has strong limitations if no other assumptions are made. One solution is to exploit structures in the problem, for example sparsity, which allow for the use of compressive sensing techniques [2]. The recovery of a sparse signal using only one-bit measurements is called one-bit compressed sensing (CS) [3]. One-bit CS techniques have proved useful in the context of optical imaging [4] and channel estimation in mmWave communications [5].

This work was partially funded by the Spanish Government and the European Regional Development Fund (ERDF) under projects TACTICA and COMPASS (TEC2013-47020-C2-1-R) and was partially supported by the National Science Foundation under Grant No. NSF-CCF-1527079.

The choice of the measurements for one-bit CS has been addressed in previous papers using mainly two different approaches: i) use fixed random measurement matrices [3] and ii) use adaptive CS, by either choosing adaptively the measurements or the threshold of the one-bit quantizer [6, 7, 8]. Since controlling the accuracy of the error estimate is difficult when using fixed transforms, the main objective of the adaptive approach is to reduce the mean square error (MSE) in the reconstruction by using the information extracted from previous measurements to iteratively construct the measurement matrix. The main limitation of this approach is that the adaptive design introduces additional computational complexity, since an optimization problem has to be solved during the acquisition process to construct the new measurements to be applied. This may be restrictive in many situations.

In this work we focus on defining one-bit CS strategies capable of reducing the MSE using fixed measurement matrices to avoid the extra complexity of adaptive approaches. In the general framework of sparse recovery from unquantized and fixed measurements, the reconstruction capability of various recovery algorithms (basis pursuit denoising, orthogonal matching pursuit and the Dantzig selector) has been analyzed in [9]. This work shows that the performance of these algorithms is close, with high probability, to the performance of the oracle estimator (i.e., solving a least squares problem on the known support). A more recent result [10] has shown that to minimize the mean squared error (MSE) of the oracle estimator the measurement matrix should be a unit norm tight frame. Coupled with the results in [9], the results in [10] also extends to the sparse recovery algorithms from unquantized measurements. Given that the reconstruction algorithms for one-bit compressed sensing are different, a natural question that arises is whether there is also a benefit from using tight frames as measurement matrices in the context of one-bit sparse recovery.

In this paper we provide a theoretical argument and then show experimentally that, in the context of one-bit sparse recovery, selecting a unit norm tight measurement matrix provides lower MSE in the reconstruction than using a random or an incoherent measurement matrix. Unlike for the adaptive methods, given that tight frames can be constructed offline there is no extra complexity associated to this approach. In the results section we show experimentally that tight frames perform better than random measurement matrices for any choice of sparsity and number of measurements. The conclusion of our work is that tight frames should be used for one-bit CS with fixed measurements.

## 2. BACKGROUND ON 1-BIT CS

Using the one-bit compressive sensing framework the available measurements are

$$\bar{\mathbf{y}} = \text{sign}(\mathbf{A}\mathbf{x}), \tag{1}$$

where sign$(z)$ is $-1$ for all $z \leq 0$ and 1 otherwise; $\mathbf{x} \in \mathbb{R}^N$ denote the sparse vector to be measured via the measurement matrix $\mathbf{A} \in \mathbb{R}^{m \times N}$ such that $m < N$ with a power constraint, i.e., a constraint on the Frobenius norm $\|\mathbf{A}\|_F$. Notice that by these binary measurements we do not obtain any information about the magnitude of $\mathbf{x}$. Thus the best we can hope to do is recover a normalized $\mathbf{x}$, i.e., placed on the unit hypersphere $\|\mathbf{x}\|_2 = 1$.

When dealing with one-bit CS, the goal is to recover the sparse signal $\mathbf{x}$ from the $m$ binary measurements by solving a non-convex optimization problem:

$$\underset{\mathbf{x}}{\text{minimize}} \ \|\mathbf{x}\|_0 \quad \text{subject to} \quad \bar{\mathbf{y}} = \text{sign}(\mathbf{A}\mathbf{x}). \tag{2}$$

Previous work has already proposed several ways to approach this problem [3, 11, 12, 7, 6, 13, 14]. In this paper we choose to proceed using the approach in [11]. Therefore, the problem in (2) is relaxed and the signal can be accurately recovered, under certain assumptions, by the following $\ell_1$ minimization convex optimization problem:

$$\begin{aligned} &\underset{\mathbf{x}}{\text{minimize}} \quad \|\mathbf{x}\|_1 \\ &\text{subject to} \quad \text{diag}(\bar{\mathbf{y}})\mathbf{A}\mathbf{x} \geq 0, \ \sum_{i=1}^{m} \bar{y}_i \mathbf{a}_i^T \mathbf{x} = m. \end{aligned} \tag{3}$$

The last constraint has the role of keeping the solution $\mathbf{x}$ away from the $\mathbf{0}_{N\times 1}$ solution and also deals with the energy ambiguity which is intrinsic to the one-bit compressive sensing problem. It has been shown that $s$-sparse signals in $\mathbb{R}^N$ can be recovered (up to normalization) from $\Omega(s\log(N/s))$ one-bit measurements [11, 12] by solving (3). In this paper, problem (3) is the standard way we will recover sparse signals from one-bit measurements.

For comparison, in the results section we also consider a binary thresholding algorithm [13] that has been proposed as an alternative to the $\ell_1$ norm recovery problem (3).

All discussions deal with the real valued case. In the case of complex variables and measurements, i.e., $\mathbf{y} = \mathbf{A}\mathbf{x}$, we equivalently have the extended real valued $\mathbf{y}^e = \mathbf{A}^e\mathbf{x}^e$:

$$\begin{bmatrix} \Re(\mathbf{y}) \\ \Im(\mathbf{y}) \end{bmatrix} = \begin{bmatrix} \Re(\mathbf{A}) & -\Im(\mathbf{A}) \\ \Im(\mathbf{A}) & \Re(\mathbf{A}) \end{bmatrix} \begin{bmatrix} \Re(\mathbf{x}) \\ \Im(\mathbf{x}) \end{bmatrix}. \tag{4}$$

Define the one-bit measurements $\bar{\mathbf{y}}^e = \text{sign}(\mathbf{y}^e)$. Thus all the previous proposed methods are solved in a $2N$–dimensional space. Most importantly, the $\ell_1$ recovery problem (3) is slightly reformulated in the complex valued case:

$$\begin{aligned} &\underset{\mathbf{x}^e}{\text{minimize}} \quad \|\mathbf{x}^e_{1:N} + j\mathbf{x}^e_{N+1:2N}\|_1 \\ &\text{subject to} \quad \text{diag}(\bar{\mathbf{y}}^e)\mathbf{A}^e\mathbf{x}^e \geq 0, \ \sum_{i=1}^{m} \bar{y}_i^e \begin{bmatrix} \Re(\mathbf{a}(\theta_i)) \\ \Im(\mathbf{a}(\theta_i)) \end{bmatrix}^T \mathbf{x}^e = m. \end{aligned} \tag{5}$$

The objective function is chosen to reflect the fact that the goal is to produce a sparse complex vector of size $N$, not a sparse real vector of size $2N$ which would be the case with the objective function $\|\mathbf{x}^e\|_1$. Notice that in the complex case each measurement provides two signs sign$(\Re(y_i))$ and sign$(\Im(y_i))$.

## 3. THE USE OF TIGHT FRAMES FOR 1-BIT CS

In this section we consider ways of designing measurements such that we reduce the MSE when using sparse signal recovery from one-bit information. We study the effects of using a measurement matrix that is a tight frame, i.e., a frame $\mathbf{A} \in \mathbb{R}^{m\times N}$ such that $\mathbf{A}\mathbf{A}^T = Nm^{-1}\mathbf{I}_m$. Tight frames have been extensively studied in the frame theory literature and have found many applications. For example reconstruction of a signal from its tight frame coefficients is numerically optimally stable while uniform tight frames have been shown to be useful for robust data transmission [15]. Unlike incoherent frames, tight frames can be easily constructed via a QR factorization applied to a random matrix and, if unit norm columns are desired, by an iterative process that exhibits low computational complexity [16] or by selecting rows of Hadamard or Fourier matrices.

We now consider the optimization problem in (3) when we assume a measurement matrix $\mathbf{A}$ that is a tight frame. We re-write problem (3) equivalently as:

$$\begin{aligned} &\underset{\mathbf{x},\mathbf{r};\ \mathbf{r}\leq 0}{\text{minimize}} \quad \|\mathbf{x}\|_1 \\ &\text{subject to} \quad \begin{bmatrix} \text{diag}(\bar{\mathbf{y}})\mathbf{A} & \mathbf{I}_m \\ m^{-1}\mathbf{1}_{1\times m}\text{diag}(\bar{\mathbf{y}})\mathbf{A} & \mathbf{0}_{1\times m} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{r} \end{bmatrix} = \begin{bmatrix} \mathbf{0}_{m\times 1} \\ 1 \end{bmatrix}, \end{aligned} \tag{6}$$

where $\mathbf{r} \in \mathbb{R}^m$ is the slack variable. We will denote

$$\mathbf{B} = \begin{bmatrix} \text{diag}(\bar{\mathbf{y}})\mathbf{A} & \mathbf{I}_m \\ m^{-1}\mathbf{1}_{1\times m}\text{diag}(\bar{\mathbf{y}})\mathbf{A} & \mathbf{0}_{1\times m} \end{bmatrix}. \tag{7}$$

We are interested in the properties of the effective measurement matrix $\mathbf{B}$ used in (6). Observe that

$$\mathbf{B}\mathbf{B}^T = \begin{bmatrix} \frac{N+m}{m}\mathbf{I}_m & \frac{N}{m^2}\mathbf{1}_{m\times 1} \\ \frac{N}{m^2}\mathbf{1}_{1\times m} & \frac{N}{m^2} \end{bmatrix} \in \mathbb{R}^{(m+1)\times(m+1)}, \tag{8}$$

has a diagonally dominant symmetric arrowhead matrix structure and is positive semi-definite. We have used the fact that

$$\text{diag}(\bar{\mathbf{y}})\mathbf{A}\mathbf{A}^T\text{diag}(\bar{\mathbf{y}})^T = \frac{N}{m}\mathbf{I}_m. \tag{9}$$

The spectral properties of general arrowhead matrices have been analyzed in [17]. In our case, there are two results that exactly describe the spectral properties of the matrix in (8).

**Result 1**.The matrix in (8) has $m-1$ eigenvalues $(N+m)m^{-1}$.

**Proof.** Let $\lambda_1, \lambda_2, \ldots, \lambda_{m+1}$ be the sorted eigenvalues of $\mathbf{BB}^T$ in increasing order. We use Cauchy's interleaving theorem ([18, page 89]) the eigenvalues $\lambda_i$ are interlaced with the sorted diagonal elements of the upper right-hand corner of $\mathbf{BB}^T$ as $0 \leq \lambda_1 \leq \frac{N+m}{m} \leq \cdots \leq \lambda_m \leq \frac{N+m}{m} \leq \lambda_{m+1}$. This shows that $\lambda_2 = \lambda_3 = \cdots = \lambda_m = (N+m)m^{-1}$ since the first $m$ diagonal elements of $\mathbf{BB}^T$ are all equal to $(N+m)m^{-1}$. ■

**Result 2**.The matrix in (8) has two eigenvalues with values $(2m^2)^{-1}(\alpha \pm \sqrt{\alpha^2 - 4m^2N})$ where $\alpha = m^2 + mN + N$.

**Proof.** Since we are looking for only two eigenvalues, $\lambda_1$ and $\lambda_{m+1}$, it is enough to find their sum and product in order to identify them. We focus on calculating tr$(\mathbf{BB}^T)$ and det$(\mathbf{BB}^T)$. The trace is trivial to compute while for the determinant we use the matrix determinant lemma det$(\mathbf{A} + \mathbf{UV}^T) = $ det$(\mathbf{I} + \mathbf{V}^T\mathbf{A}^{-1}\mathbf{U})$det$(\mathbf{A})$ on

$$\mathbf{BB}^T = \begin{bmatrix} \frac{N}{m}\mathbf{I}_m & \mathbf{0}_{m\times 1} \\ \mathbf{0}_{1\times m} & \frac{N}{m^2} - \left(\frac{N}{m^2}\right)^2 m \end{bmatrix} + \begin{bmatrix} \mathbf{I}_m \\ \frac{N}{m^2}\mathbf{1}_{1\times m} \end{bmatrix} \begin{bmatrix} \mathbf{I}_m & \frac{N}{m^2}\mathbf{1}_{m\times 1} \end{bmatrix}. \tag{10}$$

Using this result we reach the that det$(\mathbf{BB}^T) = \frac{N}{m^2}\left(\frac{N+m}{m}\right)^{m-1}$. Therefore, for the two unknown eigenvalues we have that $\lambda_1 + \lambda_{m+1} = \frac{N+m}{m} + \frac{N}{m^2}$, $\lambda_1\lambda_{m+1} = \frac{N}{m^2}$. Solving the quadratic equation for the unknowns finalizes the proof. For relative high values of $m$ (with $m < N$) we have $Nm^{-2} \approx 0$ and therefore $\lambda_1 \approx 0$ and $\lambda_{m+1} \approx (N+m)m^{-1}$. ■

These results show that simply starting from a unit norm tight measurement matrix $\mathbf{A}$ the $\ell_1$ recovery problem in (6) deals with an effective measurement matrix $\mathbf{B}$ whose properties are closely related to those of unit norm tight frames.

Given the structure of $\mathbf{B}$ these results are to be expected. In fact notice that when $\mathbf{A}$ is a unit norm tight frame so is $[\text{diag}(\bar{\mathbf{y}})\mathbf{A} \quad \mathbf{I}_m]$. Therefore, $\mathbf{B}$ operates like a tight frame in a subspace of dimension $m$. Since tight frames have been shown to improve MSE [10] for general, unquantized, sparse recovery $\ell_1$ optimization problems that are related to (6) we expect them to provide similar improvement also for the one-bit CS framework. Since the tightness of $\mathbf{B}$ depends on the number of measurements $m$ we expect that tight measurement matrices $\mathbf{A}$ will outperform random measurement matrices especially for larger $m$.

## 4. RESULTS

In this section we provide numerical evidence to highlight the benefits of using tight frames as measurement matrices for one-bit compressed sensing.

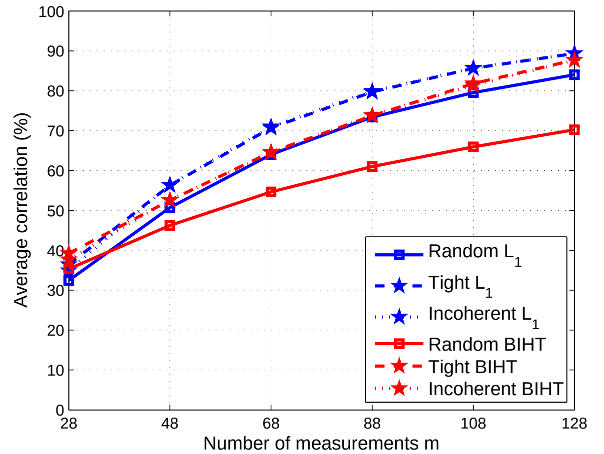


**Fig. 1**: Average correlation $|\mathbf{x}^H\mathbf{x}^\star|$ between correct and recovered target reached over 5000 experimental runs for $N = 128$ and $s = 20$. The tight and incoherent approaches almost completely overlap.

### 4.1. One-bit CS recovery

In the first experimental setting the goal is to recover a sparse signal $\mathbf{x}^\star \in \mathbb{R}^N$ generated at random using tight measurement matrices $\mathbf{A}$.

The unknown signals $\mathbf{x}^\star$ are $s$-sparse and they are generated by selecting uniformly at random the support and then drawing the entries from the standard Gaussian distribution. Finally, the signals are normalized by placing them on the unit hypersphere, i.e., $\|\mathbf{x}^\star\|_2 = 1$.

The measurement matrices are generated as follows. The random measurement matrix is generated by drawing entries from the standard Gaussian distribution. Each measurement (row of $\mathbf{A}$) is normalized separately. The tight measurement matrix is generated randomly. Both measurement matrices are normalized such that they have equal Frobenius norm.

Since the target signals $\mathbf{x}^\star$ and the recovered signals $\mathbf{x}$ are both normalized in energy we chose as performance indicator the dot product between them, i.e., $|\mathbf{x}^H\mathbf{x}^\star|$. Figure 1 shows the performance of $\ell_1$ [11] and binary iterative hard thresholding (BIHT) [13] recovery using both types of measurement matrices. Using a tight frames always yields (on average) better performance while the performance gap to the random measurements increases with larger $m$, as predicted in the previous section. In general, the $\ell_1$ based approach performs better than BIHT but the performance gap is reduced with larger $m$ and the tight measurements. We also compare against measurement matrices that are highly incoherent designed via the method in [19] that creates frames with coherence, in increasing order of $m$: 0.2099, 0.1308, 0.0910, 0.0645 and 0.0414 (for $m = 128$ we have a full orthogonal $\mathbf{A}$). As shown in [10], the incoherence does not play a crucial role in improving the MSE performance. Note that since [19] is creating highly incoherent frames (approaching

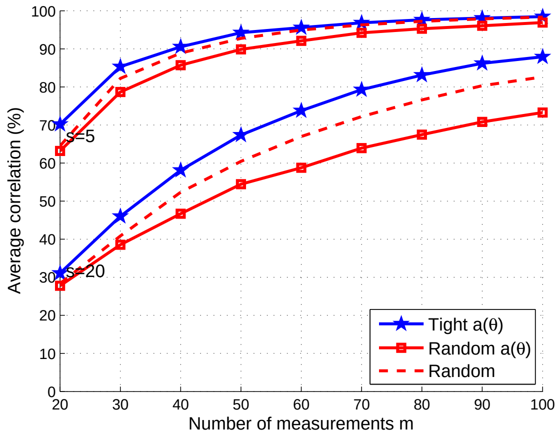


**Fig. 2**: Average correlation $|\mathbf{x}^H\mathbf{x}^\star|$ between correct and target recovered by $\ell_1$ minimization reached over 5000 experimental runs for $N = 100$, $s \in \{5, 20\}$ and measurements (12).

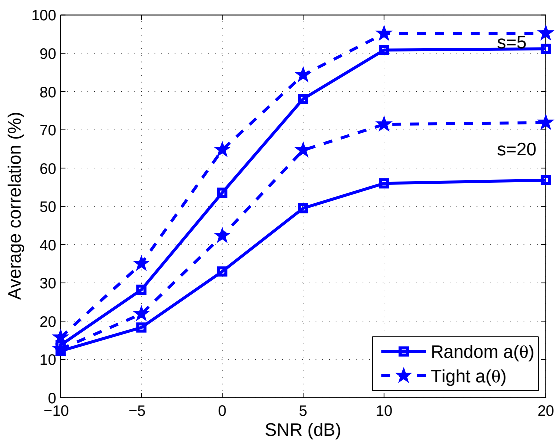


**Fig. 3**: Average correlation $|\mathbf{x}^H\mathbf{x}^\star|$ between correct and target recovered by $\ell_1$ minimization reached over 1000 experimental runs for $N = 100$, $m = 55$, $s \in \{5, 20\}$ for various SNR levels and measurements (12).

the Welch bound), these frames are very close to being tight. This explaining the small recovery differences between the tight and incoherent approaches and the performance overlap between the tight and incoherent measurement matrices in Figure 1. Indeed, incoherent frames are important when support recovery results are called upon. Unfortunately, our variable $\begin{bmatrix}\mathbf{x} & \mathbf{r}\end{bmatrix}^T$ is not sparse in general.

### 4.2. One-bit CS recovery with constrained measurements

In this section we show the effectiveness of tight frames when dealing with complex value data and with an additional constraint of selecting the measurements from a predefined set.

Consider the one-bit measurements given by:

$$\bar{\mathbf{y}} = \text{sign}(\mathbf{A}\mathbf{x} + \mathbf{n}), \tag{11}$$

where sign is the signum function applied component-wise and separately to the real and imaginary parts and $\mathbf{n} \sim \mathcal{N}(0, \sigma^2\mathbf{I})$ is Gaussian white noise. One-bit CS recovery with noisy measurements has been studied for example in [20].

Considering now that the possible measurements have the following structure:

$$\mathbf{a}^H(\theta) = \begin{bmatrix} e^{j\frac{1-N}{2}c} & e^{j\left(\frac{1-N}{2}+1\right)c} & \dots & e^{j\frac{N-1}{2}c} \end{bmatrix}, \tag{12}$$

with $c = \pi\cos(\theta)$. Our problem is to recover the transmitted signal $\mathbf{x}$ from the one-bit measurements $\bar{\mathbf{y}}$ obtained via measurement vectors like $\mathbf{a}(\theta_i), i = 1, \dots, m$.

We randomly generate instances of sparse complex valued targets $\mathbf{x}^\star \in \mathbb{C}^{N\times 1}$ whose entries are drawn from a standard Gaussian distribution and whose support is chosen uniformly at random. These targets are placed on the unit hypersphere and then try to recover them via (5) using three approaches for choose the $\mathbf{a}(\theta_i)$ which will work as measurement vectors:

- Tight $\mathbf{a}(\theta)$: the measurements are chosen randomly but such that they are orthogonal to each other, i.e., they form a tight frame.
- Random $\mathbf{a}(\theta)$: the measurements are randomly selected with an angle $\theta$ chosen uniformly from a fine grid.
- Random: the entries of the measurement matrix are sampled from a Gaussian distribution and then normalized to match the energy of the previous approaches.

The possible set of measurements is made of the vectors defined in (12) evaluated for $\theta$ on a fine grid of points $0 \le \theta_k < 2\pi$ for $k = 1, \dots, K$. In these experimental settings we chose $K = 10^4$. For $N = 100$ we show in Figure 2 the recovery results for two sparsity levels $s$: the correlation of the computed solution, after normalization, $\mathbf{x}$ with the known correct $\mathbf{x}^\star$, i.e., $|\mathbf{x}^H\mathbf{x}^\star|$, and the average error in the support recovery. Due to the normalization of $\mathbf{x}$ and $\mathbf{x}^\star$ the correlation of the two is an indication of the mean squared error. Again, the tight measurement matrices always perform better on average than the random counterparts.

In the last Figure 3 we show the effects of noise on the recovery performance. With low SNR there is little difference between the tight and random measurement matrices while with higher SNR the gap increases. Tight frames provide better performance at all SNR levels.

## 5. CONCLUSIONS

In this paper we show that unit norm tight measurement matrices provide lower reconstruction errors when dealing with sparse approximation algorithms in the context of one-bit compressed sensing. We provide both a theoretical argument and numerical example to showcase the benefits of unit norm tight measurement matrices.

## 6. REFERENCES


[1] B. Le, T. Rondeau, J. Reed, and C. Bostian, "Analog-to-digital converters," *Signal Processing Magazine, IEEE*, vol. 22, no. 6, pp. 69–77, Nov 2005.

[2] E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Comm. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, 2006.

[3] P. T. Boufounos and R. G. Baraniuk, "1-bit compressive sensing," in *Proceedings of the IEEE 42nd Annual Conference on Information Sciences and Systems (CISS)*, 2008, pp. 16–21.

[4] A. Bourquard, F. Aguet, and M. Unser, "Optical imaging using binary sensors," *Opt. Express*, vol. 18, no. 5, pp. 4876–4888, 2010.

[5] J. Mo, P. Schniter, N. González-Prelcic, and R. W. Heath Jr., "Channel estimation in millimeter wave MIMO systems with one-bit quantization," in *Proc. Asilomar Conf. on Signals, Systems and Computers*, 2014.

[6] L. Zhang, J. Yi, and R. Jin, "Efficient algorithms for robust one-bit compressive sensing," in *ICML*, 2014, pp. 820–828.

[7] R. Baraniuk, S. Foucart, D. Needell, Y. Plan, and M. Wootters, "Exponential decay of reconstruction error from binary measurements of sparse signals," *arXiv:1407.8246*, 2014.

[8] U. S. Kamilov, A. Bourquard, A. Amini, and M. Unser, "One-bit measurements with adaptive thresholds," *IEEE Sig. Proc. Letters*, vol. 19, no. 10, pp. 607–610, 2012.

[9] Z. Ben-Haim, Y. Eldar, and M. Elad, "Coherence-based performance guarantees for estimating a sparse vector under random noise," *IEEE Trans. Sig. Proc.*, vol. 58, no. 10, pp. 5030–5043, 2010.

[10] C. Wei, M. R. D. Rodrigues, and I. J. Wassell, "On the use of unit-norm tight frames to improve the average mse performance in compressive sensing applications," *IEEE Signal Processing Letters*, vol. 19, pp. 8–11, 2012.

[11] Y. Plan and R. Vershynin, "One-bit compressed sensing by linear programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1275–1297, 2013.

[12] ——, "Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach," *IEEE Trans. Inf. Theory*, vol. 59, no. 1, pp. 482–494, 2013.

[13] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, "Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors," *arXiv:1104.3160*, 2011.

[14] A. Gupta, R. Nowak, and B. Recht, "Sample complexity for 1-bit compressed sensing and sparse classification," in *Proc. IEEE Int. Symp. Information Theory*, 2010, pp. 1553–1557.

[15] V. K. Goyal, J. Kovacevic, and J. A. Kelner, "Quantized frame expansions with erasures," *Appl. Comput. Harmon. Anal.*, vol. 10, pp. 203–233, 2001.

[16] P. Davies and N. Higham, "Numerically stable generation of correlation matrices and their factors," *BIT Num. Math.*, vol. 40, no. 4, pp. 640–651, 2000.

[17] D. P. O'Leary and G. W. Stewart, "Computing the eigenvalues and eigenvectors of symmetric arrowhead matrices," *Journal of Computational Physics*, vol. 90, no. 2, pp. 497–505, 1990.

[18] J. H. Wilkinson, *The Algebraic Eigenvalue Problem*. Clarendon Press, Oxford, 1965.

[19] C. Rusu and N. González-Prelcic, "Optimized compressed sensing via incoherent frames designed by convex optimization," *arXiv:1507.02454*, 2015.

[20] M. Yan, Y. Yang, and S. Osher, "Robust 1-bit compressive sensing using adaptive outlier pursuit," *IEEE Trans. Sig. Proc.*, vol. 60, no. 7, pp. 3868–3875, 2012.