One-Bit Measurements With Adaptive Thresholds

Ulugbek S. Kamilov, *Student Member, IEEE*, Aurélien Bourquard, *Student Member, IEEE*, Arash Amini, and Michael Unser, *Fellow, IEEE*

***Abstract*—We introduce a new method for adaptive one-bit quantization of linear measurements and propose an algorithm for the recovery of signals based on generalized approximate message passing (GAMP). Our method exploits the prior statistical information on the signal for estimating the minimum-mean-squared error solution from one-bit measurements. Our approach allows the one-bit quantizer to use thresholds on the real line. Given the previous measurements, each new threshold is selected so as to partition the consistent region along its centroid computed by GAMP. We demonstrate that the proposed adaptive-quantization scheme with GAMP reconstruction greatly improves the performance of signal and image recovery from one-bit measurements.**

***Index Terms*—Analog-to-digital conversion, approximate message passing, compressive sensing, one-bit quantization.**

I. Introduction

THE linear acquisition model, where an unknown signal or image $\mathbf{x} \in \mathbb{R}^n$ is represented by measurements $\mathbf{z} = \mathbf{A}\mathbf{x} \in \mathbb{R}^m$, is central to signal processing, and many practical acquisition devices can be modeled in this way. The challenge is often to recover $\mathbf{x}$ by combining the measurements with known prior information [1]–[3]. For example, compressive sensing [4], [5] has demonstrated that it is possible to exploit the sparsity of the signal when performing the nonlinear reconstruction of $\mathbf{x}$ from $\mathbf{z}$, even when $m < n$. However, the standard approaches disregard quantization.

In realistic settings, the measurements $\mathbf{z}$ are never exact and must be discretized prior to further digital processing. In this work, we are concerned with the estimation of $\mathbf{x}$ from quantized measurements of the form $\mathbf{y} = \mathrm{Q}(\mathbf{A}\mathbf{x})$, where Q is a *one-bit* scalar quantizer. This concept was introduced in the context of compressive sensing in [6]. The key advantage of one-bit quantization is its simple and cost-effective hardware implementation as a comparator. Numerous empirical results [6]–[8] and rigorous theoretical analysis [9], [10] have demonstrated that good reconstruction performance is achievable. In fact, one-bit quantization outperforms multibit quantization in some practical configurations [7]. Several sophisticated algorithms recovering signals from one-bit measurements were proposed in [6], [8], [9], [11]–[13].

Manuscript received April 16, 2012; revised June 05, 2012; accepted July 02, 2012. Date of publication July 19, 2012; date of current version July 27, 2012. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Petros Boufounos.

The authors are with the Biomedical Imaging Group, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland (e-mail: ulugbek.kamilov@epfl.ch; aurelien.bourquard@epfl.ch; arash.amini@epfl.ch; michael.unser@epfl.ch).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/LSP.2012.2209640

Although the current formulations of one-bit compressive sensing are fundamentally deterministic, by formulating the problem in Bayesian terms we are able to extend the framework to a much larger class of signals. In particular, we present a method that can incorporate arbitrary separable priors, including sparsity-inducing ones as special cases. Moreover, the availability of a statistical model allows us to adapt the discretization to the distribution of the signal and rely on various statistical estimators. In this paper, we show that, by tuning the comparator, it is possible to significantly improve the performance of the one-bit framework. The main contributions of this work are as follows:

- An adaptation of the message-passing de-quantization algorithm of [14] to the problem of reconstruction from one-bit measurements. The algorithm is based on generalized approximate message passing (GAMP) [15] and improves upon the state of the art. It allows the linear expansion to be undercomplete or overcomplete, and can incorporate a large class of priors.
- The usage of adaptive thresholds for one-bit quantizers, which extends the applicability of the framework. Properly chosen thresholds allow to recover signals of arbitrary dynamic range and to use a broad class of measurement matrices $\mathbf{A}$.
- The development of an efficient threshold-selection method. No transmission or storage of the thresholds is required because they are fully determined from the quantized measurements.

II. One-Bit Compressive Sensing

In compressive sensing (CS), the signal $\mathbf{x} \in \mathbb{R}^n$ is acquired with $m < n$ linear measurements

$$\mathbf{z} = \mathbf{A}\mathbf{x}, \tag{1}$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$ is the *measurement matrix*. The objective is to recover $\mathbf{x}$ from $\mathbf{z}$ and $\mathbf{A}$. Although the system of equations is underdetermined, it is possible to recover the signal if some favorable conditions on $\mathbf{x}$ and $\mathbf{A}$ are satisfied. The common assumption is that the signal is exactly or approximately sparse in some orthonormal basis $\mathbf{\Psi}$. This means that there is a vector $\mathbf{u} = \mathbf{\Psi}^{-1}\mathbf{x} \in \mathbb{R}^n$ with most of its elements equal or close to zero. Additionally, for certain guarantees on the recoverability of the signal to hold, the matrix $\mathbf{A}$ must satisfy the restricted isometry property (RIP) [16]. Some families of random matrices, like appropriately dimensioned matrices with i.i.d. Gaussian or Bernoulli elements, have been shown to satisfy the RIP with overwhelming probability.

The standard CS setting assumes measurements of infinite precision. In any realistic application, however, they have to be quantized. One-bit compressive sensing considers the extreme

1070-9908/$31.00 © 2012 IEEE

example of quantization, where the measurements are represented by their signs

$$\mathbf{y} = \operatorname{sign}(\mathbf{z}). \tag{2}$$

Unfortunately, by keeping only the sign of the measurements, the amplitude of the signal is lost. Therefore, the standard reconstruction algorithms seek sparse vectors $\mathbf{x}$ satisfying (2) under some constraints on the dynamic range of the signal (e.g., $\|\mathbf{x}\|_2 = 1$). In general, such reconstructions are non-convex, and practical implementations aim at finding an approximate solution. Finally, compared to the standard case, the one-bit CS framework further restricts the class of allowed measurement matrices. For example, it cannot be generalized to Bernoulli matrices $\mathbf{A}$ when the signal is sparse in the canonical basis [8].

III. Quantization With Adaptive Thresholds

In this section, we describe extensions of one-bit CS. We first allow the one-bit quantizer to use adaptive thresholds. This is useful for extending the framework to more general signals and measurement matrices. We then provide the Bayesian formulation of the recovery problem from one-bit measurements.

A. Signal and Measurement Model

We present in Fig. 1 a generalization of the one-bit compressive-sensing framework. The input signal $\mathbf{x} \in \mathbb{R}^n$ is random with a separable distribution

$$p_{\mathbf{x}}(\mathbf{x}) = \prod_{i=1}^{n} p_x(x_i). \tag{3}$$

The noiseless measurement vector $\mathbf{z}$ is obtained via the matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$. Each entry z_i of $\mathbf{z} = \mathbf{A}\mathbf{x}$ is then compared to some scalar τ_i and set to $+1$ if it is larger that $-\tau_i$ or to -1 if it is smaller. Formally, this can be written as

$$\mathbf{y} = \mathrm{Q}(\mathbf{A}\mathbf{x}; \boldsymbol{\tau}), \tag{4}$$

where $\boldsymbol{\tau} \in \mathbb{R}^m$ is the vector of thresholds and the quantizer $\mathrm{Q} : \mathbb{R}^m \to \{-1, 1\}^m$ is the Cartesian product of m scalar quantizer components

$$q(z;\tau) = \begin{cases} +1 & \text{when } z \geq -\tau \\ -1 & \text{when } z < -\tau. \end{cases} \tag{5}$$

We define the inverse image of the component quantizer q as

$$q^{-1}(y;\tau) = \begin{cases} [-\tau, +\infty) & \text{if } y = +1 \\ (-\infty, -\tau) & \text{if } y = -1. \end{cases} \tag{6}$$

The best performance is achieved when the binary measurements are obtained sequentially. Accordingly, the previously obtained y_i can be used as feedback to adapt the next threshold value. Thus, τ_{i+1} does not constitute additional storage. It is noteworthy that the proposed formulation is compatible with standard one-bit compressive sensing when the signal prior is sparse and the thresholds are all set to zero.

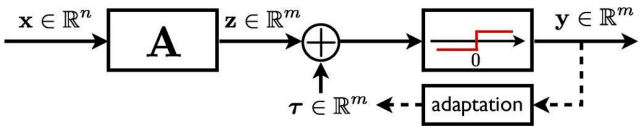

Fig. 1. Extended one-bit CS model considered in this work. The vector $\mathbf{x} \in \mathbb{R}^n$ with i.i.d. prior is estimated from scalar one-bit quantized measurements $\mathbf{y} \in \mathbb{R}^m$. The quantizer simply compares each input $z_i = [\mathbf{A}\mathbf{x}]_i$ to a threshold τ_i and sets the output to either $+1$ or -1. The best performance is achieved when the thresholds are selected adaptively.

B. Bayesian Formulation

We construct the conditional probability distribution for the signal $\mathbf{x}$ given the measurements $\mathbf{y}$ as

$$\begin{aligned} p_{\mathbf{x}|\mathbf{y}}(\mathbf{x}|\mathbf{y}) &\propto p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z}) p_{\mathbf{x}}(\mathbf{x}) \\ &\propto \prod_{i=1}^{m} \mathbb{1}_{\{z_i \in q^{-1}(y_i;\tau_i)\}}(z_i) \prod_{j=1}^{n} p_x(x_j), \end{aligned} \tag{7}$$

where $\mathbb{1}(\cdot)$ is the indicator function, and where $\propto$ denotes identity after normalization to unity. The distribution (7) provides a complete statistical characterization of the problem. In particular, the MMSE estimator of $\mathbf{x}$ is specified as

$$\hat{\mathbf{x}}_{\mathrm{MMSE}} = \mathbb{E}[\mathbf{x}|\mathbf{y}]. \tag{8}$$

Since (8) is intractable in direct form, we develop a simple computational approximation in the sequel.

IV. Estimation With Generalized Approximate Message Passing

High-dimensional integration complicates the evaluation of the posterior mean. We approximate it iteratively with a simple message-passing algorithm based on the Gaussian approximated belief propagation (BP) called *generalized approximate message passing* (GAMP) [15]. The algorithm is an extension of previous methods to nonlinear measurement channels [17], [18]. Recently, GAMP was successfully applied to reconstruct data from multibit quantized linear measurements [14]. For the complete analysis and optimality conditions of Gaussian approximated BP methods, we refer the reader to [15], [18], [19].

Given the measurements $\mathbf{y} \in \mathbb{R}^m$, the measurement matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, the vector of thresholds $\boldsymbol{\tau} \in \mathbb{R}^m$, and the prior $p_{\mathbf{x}}$, the GAMP-based MMSE estimation proceeds as follows:

1) *Initialization*: Set $t = 0$ and evaluate

$$\hat{\mathbf{x}}^t = \mathbb{E}[\mathbf{x}], \quad \mathbf{v}_x^t = \operatorname{var}[\mathbf{x}], \quad \hat{\mathbf{s}}^t = 0, \tag{9}$$

where the expected value and the variance are with respect to $p_{\mathbf{x}}$ in (3).

2) *Measurement update*: First, compute the linear step

$$\mathbf{v}_p^{t+1} = (\mathbf{A} \bullet \mathbf{A})\mathbf{v}_x^t, \tag{10a}$$

$$\hat{\mathbf{p}}^{t+1} = \mathbf{A}\mathbf{x}^t - \mathbf{v}_p^{t+1} \bullet \hat{\mathbf{s}}^t, \tag{10b}$$

where $\bullet$ denotes the Hadamard product (i.e., element-wise multiplication). Then, evaluate the nonlinear step

$$\hat{\mathbf{s}}^{t+1} = \mathrm{F}_1\left(\mathbf{y}, \boldsymbol{\tau}, \hat{\mathbf{p}}^{t+1}, \mathbf{v}_p^{t+1}\right), \tag{11a}$$
$$\mathbf{v}_s^{t+1} = \mathrm{F}_2\left(\mathbf{y}, \boldsymbol{\tau}, \hat{\mathbf{p}}^{t+1}, \mathbf{v}_p^{t+1}\right), \tag{11b}$$

where the scalar functions F_1 and F_2 are applied componentwise and given by

$$\mathrm{F}_1(y, \tau, \hat{p}, v_p) = \frac{1}{v_p}\left(\mathbb{E}\left[z | z \in q^{-1}(y;\tau)\right] - (\tau + \hat{p})\right),$$
$$\mathrm{F}_2(y, \tau, \hat{p}, v_p) = \frac{1}{v_p}\left(1 - \frac{\mathrm{var}\left[z | z \in q^{-1}(y;\tau)\right]}{v_p}\right).$$

The expected value and the variance are evaluated with respect to $z \sim \mathcal{N}(\tau + \hat{p}, v_p)$.

3) *Estimation update*: First, compute the linear step

$$\mathbf{v}_r^{t+1} = \left((\mathbf{A} \bullet \mathbf{A})^{\mathrm{T}} \mathbf{v}_s^{t+1}\right)^{-1}, \tag{12a}$$
$$\hat{\mathbf{r}}^{t+1} = \hat{\mathbf{x}}^t + \mathbf{v}_r^{t+1} \bullet (\mathbf{A}^{\mathrm{T}} \hat{\mathbf{s}}^{t+1}), \tag{12a}$$

where the inversion is componentwise. Then, evaluate the nonlinear step

$$\hat{\mathbf{x}}^{t+1} = \mathrm{G}_1\left(\hat{\mathbf{r}}^{t+1}, \mathbf{v}_r^{t+1}; p_x\right), \tag{13a}$$
$$\mathbf{v}_x^{t+1} = \mathrm{G}_2\left(\hat{\mathbf{r}}^{t+1}, \mathbf{v}_r^{t+1}; p_x\right), \tag{13b}$$

where the scalar functions G_1 and G_2 are applied componentwise and given by

$$\mathrm{G}_1(\hat{r}, v_r; p_x) = \mathbb{E}[x|\hat{r}],$$
$$\mathrm{G}_2(\hat{r}, v_r; p_x) = \mathrm{var}[x|\hat{r}].$$

The expected value and the variance are evaluated with respect to $p_{x|\hat{r}}(\cdot|\hat{r}) \propto \mathcal{G}(\cdot; \hat{r}, v_r) p_x(\cdot)$, where $\mathcal{G}(\cdot; \hat{r}, v_r)$ is the Gaussian pdf of mean $\hat{r}$ and variance v_r. This is essentially an AWGN denoising problem with noise $n \sim \mathcal{N}(0, v_r)$.

4) Set $t \leftarrow t + 1$ and proceed to step 2).

For each iteration $t = 0, 1, 2, \ldots$, the proposed update rules produce estimates $\hat{\mathbf{x}}^t$ of the true signal $\mathbf{x}$. Thus, the algorithm reduces the intractable high-dimensional integration to a sequence of matrix-vector products and scalar non-linearities.

V. Adaptive Thresholds

Depending on $\mathbf{A}$ and $\boldsymbol{\tau}$, each y_i defines a particular half-space containing $\mathbf{x}$ that is delimited by a hyperplane. The intersection of these m domains forms a convex *consistent set* Ω inside which any solution $\hat{\mathbf{x}}$ is associated with the same vector $\mathbf{y}$. According to this geometrical interpretation, the theoretical MMSE solution that is estimated using GAMP is the center of mass of the probability distribution in Ω.

An efficient way to reduce the quantization error is to select $\boldsymbol{\tau}$ adaptively. As a computationally tractable solution, we propose to adapt each next τ_i such that the corresponding hyperplane passes through the center of mass $\hat{\mathbf{x}}_i$ of the currently known consistent set Ω_i. Given the currently known measurement vector $\mathbf{y}_i$, we estimate $\hat{\mathbf{x}}_i$ using GAMP according to the joint posterior probability $p_{\mathbf{x}|\mathbf{y}_i}$. For large-scale problems, several thresholds can be updated simultaneously. Note that, due to the adaptive nature of the approach, the full knowledge of the measurements is required at reconstruction. Moreover, correct recovery of the thresholds relies on the fact that the same stopping criteria are used for GAMP during the measurement and reconstruction processes.

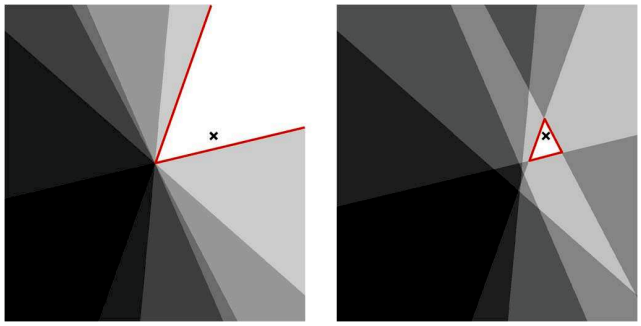

Fig. 2. Geometrical representation of one-bit CS with $n = 2$ and $m = 6$. The use of fixed versus adaptive thresholds is illustrated on the left and right, respectively. The consistent sets (white zones) are obtained from the same $\mathbf{A}$, assuming a bounded distribution on $\mathbf{x}$. The oracle is represented by a cross.

In Fig. 2, we illustrate how the use of adaptive thresholds yields consistent sets that are closed as well as substantially smaller than when applying zero thresholds on the same measurements. These observations also corroborate the performance of reconstruction that is addressed in Section VI.

VI. Experimental Results

A. Sparse Estimation

We consider the estimation of an n-dimensional sparse signal $\mathbf{x}$ from m one-bit measurements. We perform 1000 random trials and plot the average signal-to-noise ratio (SNR) of the reconstruction against m/n. For each trial, we generate a signal that has length $n = 100$ with 50 nonzero components drawn from the standard normal distribution, and form a measurement matrix $\mathbf{A}$ from i.i.d. zero-mean Gaussian random variables of variance $1/m$.

In Fig. 3, we compare the reconstruction performance of GAMP with Gauss-Bernoulli prior against the binary iterative hard thresholding (BIHT) algorithm introduced in [9].[1] BIHT has been shown to yield state-of-the-art performance for reconstructing data from one-bit measurements. We consider the standard scenario where all the thresholds are set to zero. For fair comparison, the signal $\mathbf{x}$ is normalized to lie on the unit ball. We also normalize the reconstructed signals for both algorithms. The results show that GAMP significantly outperforms BIHT over the whole range of m/n.

In Fig. 4, we compare the reconstruction performance of GAMP with and without adaptive thresholds. The thresholds are set according to the procedure described in Section V. As

[1]The source code for the BIHT algorithm can be downloaded from http://dsp.rice.edu/software/binary-iterative-hard-thresholding-biht-demo/

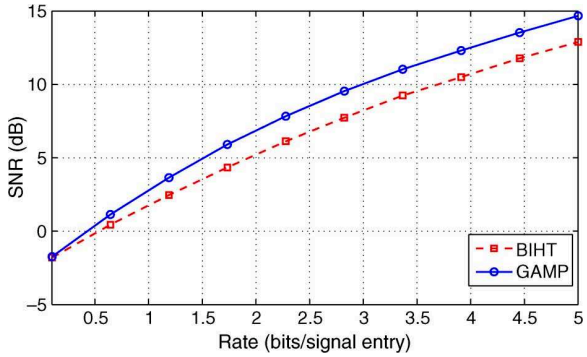

Fig. 3. Standard scenario. The average reconstruction SNR is plotted against the measurement ratio m/n for GAMP (solid) and BIHT (dashed) estimations of sparse signals. The plot demonstrates that GAMP yields considerable improvement (up to 2 dB).

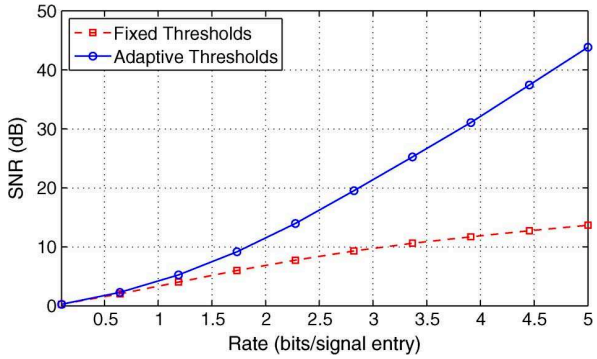

Fig. 4. Adaptive scenario. The average reconstruction SNR is plotted against the measurement ratio m/n for GAMP estimation with (solid) and without (dashed) adaptive thresholds. The plot illustrates that significant gains can be achieved by using one-bit quantizers with adaptive thresholds.

Fig. 5. Reconstruction of *Cameraman* from one-bit measurements: (a) original image, (b) reconstruction with zero thresholds ($\mathrm{SNR} = 13.20$ dB), (c) reconstruction with adaptive thresholds ($\mathrm{SNR} = 24.42$ dB).

expected, the adaptive choice of the thresholds considerably improves the quality of reconstruction.

B. Image Reconstruction

We now consider the problem of image recovery from one-bit measurements. We use the standard 8-bit grayscale test image Cameraman of size 128×128 pixels shown in Fig. 5. We form the measurement matrix $\mathbf{A}$ from i.i.d. $\pm 1/\sqrt{m}$ random variables that follow uniform Bernoulli distribution. The reconstruction is performed from $m = 3n$ one-bit measurements, which corresponds to the rate of 3 bits per image pixel. We compare the SNR performance of the GAMP reconstruction with zero thresholds and with adaptive thresholds. GAMP is applied with i.i.d. Gauss-Bernoulli prior on the Haar-wavelet coefficients of the signal. The Gauss-Bernoulli variables are nonzero with probability 0.3 with a variance that is matched to the average wavelet coefficient variance. Adaptive thresholds were determined simultaneously in groups of 1000. The results in Fig. 5 confirm that well-chosen thresholds improve the reconstruction significantly.

VII. Conclusion

We have presented a method that attains a high-quality signal recovery from one-bit linear measurements. The method relies on the selection of adaptive thresholds and on the GAMP reconstruction algorithm. The overall algorithm is computationally simple and general, allowing essentially arbitrary priors on the signal. The practical relevance of the method has been illustrated through several numerical evaluations.

References

[1] M. Lustig, D. L. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, December 2007.

[2] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 83–91, Mar. 2008.

[3] E. Bostan, U. S. Kamilov, and M. Unser, "Reconstruction of biomedical images and sparse stochastic modelling," in *Proc. Int. Symp. Biomedical Imaging*, Barcelona, Spain, May 2012.

[4] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[5] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[6] P. T. Boufounos and R. G. Baraniuk, "1-bit compressive sensing," in *Proc. Conf. Information Sciences and Systems*, Princeton, NJ, Mar. 2008, pp. 16–21.

[7] J. N. Laska and R. G. Baraniuk, Regime Change: Bit-Depth Versus Measurement-Rate in Compressive Sensing October 2011, arXiv:1110.3450v1 [cs.IT].

[8] Y. Plan and R. Vershynin, One-Bit Compressed Sensing by Linear Programming September 2011, arXiv:1109.4299v4 [cs.IT].

[9] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, Robust 1-Bit Compressive Sensing Via Binary Stable Embeddings of Sparse Vectors April 2011, arXiv:1104.3160v2 [cs.IT].

[10] Y. Plan and R. Vershynin, Dimension reduction by random hyperplane tessellations November 2011, arXiv:1111.4452 [math.PR].

[11] P. T. Boufounos, "Greedy sparse signal reconstruction from sign measurements," in *Proc. of Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2009, pp. 1305–1309.

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

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

[14] U. S. Kamilov, V. K. Goyal, and S. Rangan, Message-Passing Estimation From Quantized Samples Nov. 2011, arXiv:1105.4652 [cs.IT].

[15] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Information Theory*, St. Petersburg, Russia, July–August 2011, pp. 2168–2172.

[16] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.

[17] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," in *Proc. Nat. Acad. Sci.*, Nov. 2009, vol. 106, no. 45, pp. 18 914–18 919.

[18] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 764–785, Feb. 2011.

[19] S. Rangan, "Estimation with random linear mixing, belief propagation and compressed sensing," in *Proc. Conf. on Inform. Sci. & Sys.*, Princeton, NJ, Mar. 2010, pp. 1–6.