One-bit Compressive Sampling with Time-Varying Thresholds for Multiple Sinusoids

Christopher Gianelli, Luzhou Xu, Jian Li
Electrical and Computer Engineering
University of Florida
Gainesville, Florida 32611 USA
Email: (gianelli04@ufl.edu, xuoasis@gmail.com, li@dsp.ufl.edu)

Petre Stoica
Department of Information Technology
Uppsala University
Uppsala, Sweden
Email: (ps@it.uu.se)

***Abstract*—Wide-band spectral sensing is a challenging task that will be required in future cognitive radio and radar applications. Recent research has shown that sampling using only one-bit of amplitude precision can be realized at an extremely high rate [1] in an affordable manner. In this work, one-bit sampling using time-varying thresholds is considered for line spectral estimation. The time-varying thresholds allow for amplitude estimation. A novel one-bit RELAX algorithm is developed for multi-tone parameter estimation. This algorithm is shown to have excellent performance via a numerical example.**

I. Introduction

Quantization of signals of interest is a necessary first step in digital signal processing applications. Typically signal quantization is implemented as high-resolution amplitude quantization at a sampling frequency above the Nyquist rate [2] after suitable low-pass filtering to avoid aliasing. In this sampling regime the amplitude quantization is fine enough such that the sampling error can be modeled as additive noise, and has little impact on the performance of algorithms which are typically developed under the assumption of infinite precision sampling. However, in applications such as spectral sensing for cognitive radio and radar [3], [4] which require extremely high sampling rates, it may be impractical or impossible to achieve Nyquist sampling at even moderate amplitude precision.

The most extreme form of quantization is reduction of the signal to one-bit per sample, which may be accomplished via repeatedly comparing the signal to some reference level, and recording whether the signal is above or below the reference. One-bit sampling allows for an extremely high sampling rate at a low cost. Problems involving data quantized to one-bit have been studied from a classical statistical viewpoint in [5]–[9], from a sampling and reconstruction viewpoint in [10], [11], and from a compressive sensing viewpoint in [12]–[21]. It is important to note that many of the cited works use only comparisons to zero, which obliterates information about the amplitude of the signals of interest.

The authors of this work have recently published two papers [22], [23] that investigate the problem of sinusoidal signal parameter estimation using one-bit samples with *time-varying* thresholds. The use of one-bit sampling with time-varying thresholds allows for amplitude estimation, and has been demonstrated in hardware systems [1]. In [22] a sparse semi-parametric approach was introduced for sinusoidal parameter estimation in a one-bit sampling with time-varying thresholds system. It was shown that a logarithm penalty offered more parsimonious models and more accurate amplitude estimates than the common ℓ_1 penalty. In [23] a parametric approach to the spectral sensing is developed, and a maximum likelihood (ML) estimator for sinusoidal parameter estimation was proposed.

Here, the problem of multiple sinusoid parameter estimation using one-bit samples subjected to time-varying thresholds is investigated. It is shown that the true ML estimator is computationally inefficient for signals with numerous frequency components. Two computationally simpler algorithms are introduced for this case. One is analogous to the periodogram approach from infinite precision spectral estimation [24]. The second approach is analogous to the RELAX algorithm from the spectral estimation literature [25]. Finally, the previously developed semi-parametric sparse approaches [22] are applied to this problem, and slightly modified to improve their amplitude estimates.

II. Signal Model

Consider a signal, $s(t,\boldsymbol{\beta})$, where t denotes time and $\boldsymbol{\beta} \in \mathbb{R}^N$ is a vector of unknown parameters. Let $s(t,\boldsymbol{\beta})$ be compared to a known reference function $h(t)$ via a comparator, whose output $y(t)$ takes the form

$$y(t) = \text{sign}\left(s(t,\boldsymbol{\beta}) + e(t) - h(t)\right), \qquad (1)$$

where the term $e(t)$ denotes unknown additive noise. The sign function is characterized by

$$\text{sign}\left(x\right) = \begin{cases} 1 & x \geq 0 \\ -1 & x < 0 \end{cases} \qquad (2)$$

In this model, T samples of $y(t)$ are captured, not necessarily uniformly in time, and are stored in the vector $\mathbf{y} = \left[y(t_0), y(t_1), \ldots, y(t_{T-1})\right]^T$, where $(\cdot)^T$ denotes matrix transpose. The problem of interest is to estimate the parameter vector $\boldsymbol{\beta}$ using the one-bit samples $\mathbf{y}$ and the knowledge of the thresholds $\mathbf{h} = \left[h(t_0), h(t_1), \ldots, h(t_{T-1})\right]^T$.

Let $s(t,\boldsymbol{\beta})$ be a sum of sinusoids, or

$$s(t,\boldsymbol{\beta}) = \sum_{n=0}^{N-1} \left[a_n \sin(\omega_n t) + b_n \cos(\omega_n t)\right], \qquad (3)$$

where $\boldsymbol{\beta} = \left[a_0, b_0, \omega_0, a_1, b_1, \omega_1, \ldots, a_{N-1}, b_{N-1}, \omega_{N-1}\right]^T \in \mathbb{R}^{3N}$.

III. Likelihood Function

Suppose that the additive noise $e(t)$ is an i.i.d. Gaussian noise, or $e(t) \sim \mathcal{N}\left(0, \sigma^2\right)$. Define the sets I_+ and I_- such that $I_+ = \{k : y(t_k) = 1\}$, and $I_- = \{k : y(t_k) = -1\}$.

Then, the likelihood function can be expressed as

$$L(\boldsymbol{\beta}) = \prod_{k \in I_+} \Phi\left(\frac{s(t_k, \boldsymbol{\beta}) - h(t_k)}{\sigma}\right) \prod_{k \in I_-} \Phi\left(\frac{-s(t_k, \boldsymbol{\beta}) + h(t_k)}{\sigma}\right), \tag{4}$$

where $\Phi(\cdot)$ is the standard Normal cumulative distribution function (CDF). After some simple manipulations, the final expression for the likelihood function of the measurements is given by

$$L(\boldsymbol{\beta}) = \prod_{k=0}^{T-1} \Phi\left(y(t_k)\frac{s(t_k, \boldsymbol{\beta}) - h(t_k)}{\sigma}\right). \tag{5}$$

IV. Parameter Estimation

A. Maximum Likelihood Estimation

The maximum likelihood (ML) approach to parameter estimation has been applied to this problem [23], and been shown to offer good performance. ML estimators are attractive as they have many desirable properties, including consistency, asymptotic efficiency (attainment of the Cramér-Rao Bound (CRB)), and asymptotic normality (see e.g., [26]). In this approach, the parameters of interest are estimated by finding the values which minimize the negative log-likelihood function. Therefore, the problem under consideration is given in (6).

Let $\boldsymbol{\omega} = [\omega_0, \omega_1, \ldots, \omega_{N-1}]^T$ be a vector containing the angular frequencies of the input signal. Note that for given $\boldsymbol{\omega}$ the optimization problem in (6) is convex in $\{a_n\}_{n=0}^{N-1}$ and $\{b_n\}_{n=0}^{N-1}$ (see e.g., [27]). This holds true regardless if σ is known or unknown.

Given the convexity of the problem in the amplitude parameters for fixed $\boldsymbol{\omega}$, ML estimation of $\boldsymbol{\beta}$ can be carried out by performing an N dimensional search over the space of angular frequencies for each of the sinusoidal signal components. The search may be carried out as follows. Begin by establishing a discrete set of L points in $[0, \pi)$, e.g., $\phi_l = \frac{\pi l}{L}$, $l = 0, 1, \ldots, L-1$. Fix an $\boldsymbol{\omega}$ such that each ω_n, $n = 0, 1, \ldots, N-1$, is equal to some ϕ_l. For this fixed $\boldsymbol{\omega}$, solve (6) using, e.g., Newton's method (see e.g., [28]). Store the optimal $\{\hat{a}_n\}_{n=0}^{N-1}$ and $\left\{\hat{b}_n\right\}_{n=0}^{N-1}$, as well as the value of the log likelihood function at this optimal solution. Finally, repeat the previous steps after fixing a new $\boldsymbol{\omega}$. After all the $\boldsymbol{\omega}$ have been searched, the optimal $\boldsymbol{\omega}$ and $\{\hat{a}_n\}_{n=0}^{N-1}$ and $\left\{\hat{b}_n\right\}_{n=0}^{N-1}$ correspond to the minimum negative log-likelihood value obtained.

The previously described algorithm requires an N dimensional search, with L points to search along each dimension. This means that (6), which is a $2N$ dimensional convex problem, must be solved $\mathcal{O}\left(L^N\right)$ times. As the number of sinusoids in the model increase, the ML approach rapidly becomes computationally inefficient if a global optimal solution is desired. In this case, more efficient algorithms must be used.

B. One-Bit Periodogram

The well-known periodogram approach to spectral estimation involves computing the squared magnitude of the discrete Fourier transform (DFT) of a sampled data sequence. As described in [24], the periodogram may serve as an "approximate ML" approach, provided that the frequencies of the sinusoids are spaced sufficiently far apart. Furthermore, note that, in the case of a single sinusoid ($N = 1$), the location and complex amplitude corresponding to the maximum of the periodogram is the ML estimator when the noise is Gaussian.

In a similar spirit, the one-bit periodogram is defined as follows. First, solve (6) assuming only a single sinusoid exists for $\omega = \phi_l$, $l = 0, 1, \ldots, L-1$. This process yields the optimum a and b for each frequency ϕ_l, as well as the value of the optimized log-likelihood function. Then, find the N largest peaks in the computed log-likelihood function. Finally, the frequencies corresponding to the N largest peaks are used to estimate the amplitude and phase of the strongest signals via solving (6).

Clearly, this approach is far less computationally expensive than the ML technique, as the one-bit periodogram only requires solving L 2-dimensional convex problems. However, similar to the infinite precision periodogram, the algorithm may not offer good estimates when multiple signals are too closely spaced, or if there is a strong signal very near to a weaker one. In the following subsection, an approach which offers better performance than the one-bit periodogram with less computational complexity than the true ML algorithm will be presented.

C. One-Bit RELAX

As described in [24], ML parameter estimation for sinusoidal signals with infinite precision samples can be expressed as a non-linear least squares (NLS) optimization problem. The infinite precision case is thus similar to the one-bit sample case under consideration, in that each case requires high dimensional searches. The paper [25] developed a relaxation approach for this problem that can offer good parameter estimation performance and acceptable computational complexity. The algorithm, named RELAX, is extended to the one-bit precision sampling case in this work. Let $\boldsymbol{\beta}_n = [a_n, b_n, \omega_n]^T \in \mathbb{R}^3$ be a subvector of $\boldsymbol{\beta}$ containing the parameters of the n^{th}sinusoidal component. The operation of the one-bit RELAX algorithm is described below.

The algorithm begins by assuming that $N = 1$, and performs parameter estimation under this assumption. It has been shown above that the ML estimate of the amplitude, frequency, and phase of a single sinusoid requires only a one-dimensional search over ω. That is, (6) is solved for $N = 1$. This process will yield an estimate of $\boldsymbol{\beta}_1$ of the strongest sinusoidal signal component.

With an estimate of the dominant sinusoidal component, the algorithm moves to estimate the *next* strongest signal component. To this end, the algorithms assumes that $N = 2$, and now solves a slightly modified version of (6) given in (7). That is, the estimates of the dominant signal from the first stage of the algorithm are used to suppress the impact of this strong signal, improving the estimates of the parameters of the second strongest signal. The problem in (7) requires a search over one dimension, and solving a 2 dimensional convex problem at each search location.

$$\hat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta}} \sum_{k=0}^{T-1} -\log\left(\Phi\left(y(t_k)\frac{\sum_{n=0}^{N-1}\left[a_n \sin(\omega_n t_k) + b_n \cos(\omega_n t_k)\right] - h(t_k)}{\sigma}\right)\right). \quad (6)$$

$$\hat{\boldsymbol{\beta}}_2 = \arg\min_{\boldsymbol{\beta}_2} \sum_{k=0}^{T-1} -\log\left(\Phi\left(y(t_k)\frac{\hat{a}_1 \sin(\hat{\omega}_1 t_k) + \hat{b}_1 \cos(\hat{\omega}_1 t_k) + a_2 \sin(\omega_2 t_k) + b_2 \cos(\omega_2 t_k) - h(t_k)}{\sigma}\right)\right). \quad (7)$$

In a similar fashion, the estimate of the dominant sinusoid may be refined by considering the problem in (8). The algorithm continues iteratively refining the estimates of the two strongest signal components until practical convergence, i.e., until the change in cost or the estimates is small enough.

Once good estimates of the two strongest signals have been obtained, the algorithm may proceed to $N = 3$. Similar to the case of $N = 2$, the algorithm begins by estimating the third strongest signal component's parameters via assuming the estimates from the previous stages of the algorithms are known. Then, the algorithm iterates between the various signal components and refines the estimates until practical convergence is achieved. The model order is increased again, and the algorithm proceeds until the known model order is reached.

The benefit of this approach in comparison to the true ML method is in its computational efficiency. The proposed one-bit RELAX approach requires solving $\mathcal{O}\left(CLN^2 + CLN\right)$ 2-dimensional convex problems, where C is the number of iterations required to achieve practical convergence at each model order. Note that L, the size of the frequency grid used for the search, is proportional to the number of available samples for parameter estimation, and much larger than N. Thus, the one-bit RELAX algorithm's computation burden grows much less rapidly than that of the true ML estimator.

D. Sparse Estimation and ML

The final parameter estimation technique considered in this work is a combination between the sparse parameter estimation approaches from [22] with a final step of ML parameter estimation. The particulars of the sparse parameter estimation approaches are not provided here due to space constraints. The interested reader can consult [22] for the details on the ℓ_1-norm and logarithm approaches used in this work.

This approach begins by estimating the spectrum using either the ℓ_1-norm or log-penalty approach from [22]. After the spectrum has been estimated using either of these algorithms, the N strongest (in terms of amplitude) peaks are extracted, and the frequencies of these peaks are stored. Then, (6) is solved using the frequencies of the N strongest peaks, which yields the final estimates of the amplitude and phase of the sinusoidal signal components.

This method is attractive from a computation standpoint because the algorithm's complexity increases very little with model order, which is different than the true ML and RELAX approaches. Furthermore, the final step of ML estimation offers more accurate estimates of the amplitude of the signal components, as the sparse estimation techniques have the well-known effect of "shrinkage" [29], [30]. That is, the estimates offered by these sparse estimation approaches are biased downwards. The final step of ML estimation improves these amplitude estimates at a relatively small computational penalty (only one $2N$ dimensional convex problem must be solved).

TABLE I. SIGNAL PARAMETERS

Parameter	**1**	**2**	**3**	**4**	**5**	**6**
Amplitude	1	1	1	1	1	1
Freq. (Hz)	0.11	0.45	0.3	0.15	0.2	0.37
Phase (rad)	$\frac{2\pi}{3}$	$\frac{2\pi}{4}$	$\frac{14\pi}{8}$	$\frac{10\pi}{6}$	0	$\frac{4\pi}{3}$

V. NUMERICAL RESULTS

Results obtained by applying the previously described approaches are now presented and discussed. The case considered is a signal composed of six sinusoids with different frequencies and phases, each with an amplitude of 1. Table I displays the parameters of the six sinusoids. One hundred independent Monte Carlo trials were run, using each of the aforementioned algorithms to estimate the parameters of the signal. The number of one-bit samples and thresholds used was varied with $T = 128$, 256, and 512. Each trial had an independent noise realization, with the signal-to-noise ratio (SNR) taking values of 10 dB and 15 dB. The SNR was computed as

$$\text{SNR} = \sum_{n=1}^{N} \frac{\sqrt{a_n^2 + b_n^2}}{2\sigma^2}. \quad (9)$$

Note that the amplitude of the n^{th}sinusoidal component is defined as $A_n = \sqrt{a_n^2 + b_n^2}$.

Both the one-bit samples and the threshold levels are used for parameter estimation. The same set of thresholds are used for each of the trials run. The thresholds are drawn from a discrete set of 8 values evenly distributed in [-1, 1]. The threshold value for a particular sample is drawn randomly and uniformly from this set of 8 possible threshold levels. It should be noted that the model order (i.e., the number of sinusoidal signal components) is known. Model order selection and estimation will be considered in future works.

Figures 1(a) and (b) display the average frequency estimation mean square error (MSE) over the six sinusoidal signal components as a function of T for SNR values of 10 dB and 15 dB, respectively. In these figures, the blue curve labeled "1-b Per.-ML" corresponds to the one-bit periodogram with a final step of ML estimation, the green curve labeled "1-b RELAX" corresponds to the one-bit RELAX algorithm, the red curve labeled "ℓ_1-ML" corresponds to the ℓ_1-norm approach, and the cyan curve labeled "log-ML" corresponds to the log-penalty approach. Finally, the black curve labeled "Avg.-CRB" corresponds to the Cramér-Rao Bound for the set of signals and thresholds averaged over the six signals considered.

Inspecting these results, it can be seen that the one-bit RELAX algorithm provides the best frequency estimation performance, with the log-penalty and one-bit periodogram

$$\hat{\boldsymbol{\beta}}_1 = \arg\min_{\boldsymbol{\beta}_1} \sum_{k=0}^{T-1} -\log\left(\Phi\left(y(t_k)\frac{a_1\sin(\omega_1 t_k) + b_1\cos(\omega_1 t_k) + \hat{a}_2\sin(\hat{\omega}_2 t_k) + \hat{b}_2\cos(\hat{\omega}_2 t_k) - h(t_k)}{\sigma}\right)\right). \quad (8)$$

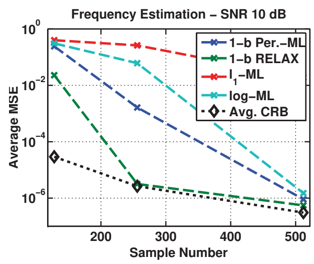

(a)

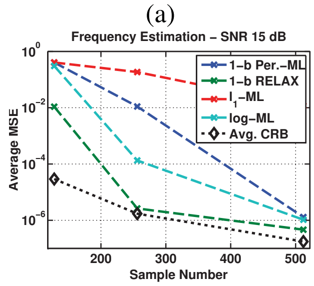

(b)

Fig. 1. Average (over signal components) frequency estimation MSE as a function of T for an SNR of (a) 10 dB and (b) 15 dB.

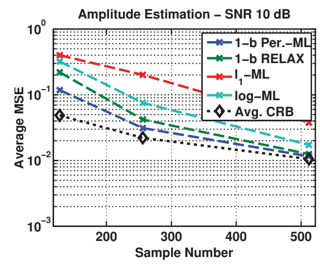

(a)

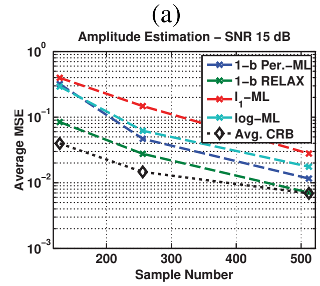

(b)

Fig. 2. Average (over signal components) amplitude estimation MSE as a function of T for an SNR of (a) 10 dB and (b) 15 dB.

approaches offering good performance when the sample support is large. Note that the increase in SNR by 5 dB has little effect on the frequency estimation performance in the example considered, as the one-bit RELAX and log-penalty approaches perform similarly in both cases. Note that the one-bit periodogram performance actually worsens slightly with improving SNR. After inspecting the raw outputs of the one-bit periodogram, it was seen that for the particular signal and threshold case considered, a large sidelobe was present. As the SNR increased, the average strength of this sidelobe increased, leading it to compete with the true signals for detection.

Figures 2(a) and (b) display the average amplitude estimation mean square error (MSE) over the six sinusoidal signal components as a function of T for SNR values of 10 dB and 15 dB, respectively. Inspecting these results, it can be seen that the one-bit periodogram appears to offer the lowest amplitude estimation error for the 10 dB SNR case. However, note that the one-bit periodogram frequency estimates are poor at low sample support, and the algorithm fails to detect some of the signals in the $T = 128$ and $T = 256$ cases considered. The missed detection problem is not taken into account in the computation of the average amplitude estimation MSE, making the one-bit periodogram amplitude estimation seem to perform better than it does. Inspecting the one-bit RELAX curves, it can be seen that this approach offers good amplitude estimation performance, along with good frequency estimation performance when $T \geq 256$ for the signal and threshold case considered.

Table II displays the computed average run time for each of the algorithms. Clearly, the one-bit periodogram is the most computationally efficient method, with the ℓ_1-ML algorithms requiring approximately quadruple the run time. The log penalty approach requires iteratively solving re-weighted ℓ_1

TABLE II. Average running time for algorithms as a function of T.

Algorithm	1-b Per.	1-b RELAX	ℓ_1-ML	log.-ML
$T = 128$	0.49	70.15	1.96	10.26
$T = 256$	1.34	174.22	13.20	60.01
$T = 512$	3.87	418.74	154.86	568.04

problems, hence the longer run times. Furthermore, while the one-bit RELAX algorithm generally requires more computation than the other methods considered, it provides the best estimation performance of all the approaches considered with manageable computational complexity.

VI. Conclusion

A one-bit sampling scheme which uses time-varying thresholds has been proposed for wide-band spectral sensing. This architecture allows for very high-rate sampling enabled by the use of only one-bit of amplitude precision, while preserving critical information about the amplitude of the signal. The problem of line spectra parameter estimation was considered, and several techniques from infinite precisions spectral estimation have been extended to this problem. It was shown that the one-bit RELAX algorithm can offer excellent estimation performance at a reasonable computational cost. The simple one-bit periodogram also provided reasonable performance for the well-spaced sinusoidal tones with very little computational expense.

Acknowledgement

This work was supported in part by the National Science Foundation (NSF) under Grant No. 1704240.

References

[1] Novelda, "X2 impulse radar transceiver," X2 Impulse Radar Transceiver Data Sheet, Rev. G, pp. 1–51, November 2015.

[2] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing, Fourth Edition*. Upper Saddle River, New Jersey 07458: Pearson Prentice Hall, 2007.

[3] H. Sun, A. Nallanathan, C. X. Wang, and Y. Chen, "Wideband spectrum sensing for cognitive radio networks: A survey," *IEEE Wireless Communications*, vol. 20, pp. 74–81, April 2013.

[4] J. Lunden, V. Koivunen, and H. V. Poor, "Spectrum exploration and exploitation for cognitive radio: Recent advances," *IEEE Signal Processing Magazine*, vol. 32, pp. 123–140, May 2015.

[5] A. Ribeiro and G. B. Giannakis, "Bandwidth-constrained distributed estimation for wireless sensor networks - Part II: Unknown probability density function," *IEEE Transactions on Signal Processing*, vol. 54, pp. 2784–2796, July 2006.

[6] A. Host-Madsen and P. Handel, "Effects of sampling and quantization on single-tone frequency estimation," *IEEE Transactions on Signal Processing*, vol. 48, pp. 650–662, March 2000.

[7] O. Bar-Shalom and A. J. Weiss, "DOA estimation using one-bit quantized measurements," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, pp. 868–884, July 2002.

[8] O. Dabeer and A. Karnik, "Signal parameter estimation using 1-bit dithered quantization," *IEEE Transactions on Information Theory*, vol. 52, pp. 5389–5405, December 2006.

[9] O. Dabeer and E. Masry, "Multivariate signal parameter estimation under dependent noise from 1-bit dithered quantized data," *IEEE Transactions on Information Theory*, vol. 54, pp. 1637–1654, April 2008.

[10] E. Masry, "The reconstruction of analog signals from the sign of their noisy samples," *IEEE Transactions on Information Theory*, vol. 27, pp. 735–745, November 1981.

[11] Z. Cvetkovic, I. Daubechies, and B. F. Logan, "Single-bit oversampled A/D conversion with exponential accuracy in the bit rate," *IEEE Transactions on Information Theory*, vol. 53, pp. 3979–3989, November 2007.

[12] P. T. Boufounos and R. G. Baraniuk, "1-bit compressive sensing," *42nd Annual Conference on Information Sciences and Systems,* Princeton, NJ, pp. 16–21, March 2008.

[13] A. Zymnis, S. Boyd, and E. Candes, "Compressed sensing with quantized measurements," *IEEE Signal Processing Letters*, vol. 17, pp. 149–152, February 2009.

[14] Y. Plan and R. Vershynin, "One-bit compressed sensing by linear programming," *Communications on Pure and Applied Mathematics*, vol. 66, pp. 1275–1297, August 2013.

[15] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, "Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors," *IEEE Transactions on Information Theory*, vol. 58, pp. 2082–2102, March 2013.

[16] Y. Plan and R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: a convex programming approach," *IEEE Transactions on Information Theory*, vol. 59, pp. 482–494, January 2013.

[17] M. Yan, Y. Yang, and S. Osher, "Robust 1-bit compressive sensing using adaptive outlier pursuit," *IEEE Transactions on Signal Processing*, vol. 60, pp. 3868–3875, April 2012.

[18] U. S. Kamilov, A. Bourquard, A. Amini, and M. Unser, "One-bit measurements with adaptive thresholds," *IEEE Signal Processing Letters*, vol. 19, pp. 607–610, October 2012.

[19] A. Ai, A. Lapanowski, Y. Plan, and R. Vershynin, "One-bit compressed sensing with non-Gaussian measurements," *Linear Algebra and its Applications*, vol. 441, pp. 222–239, January 2014.

[20] L. Zhang, J. Yi, and R. Jin, "Efficient algorithms for robust one-bit compressive sensing," *Proceedings of the 31st International Conference on Machine Learning,* Beijing, China, pp. 820–828, June 2014.

[21] K. Knudson, R. Saab, and R. Ward, "One-bit compressive sensing with norm estimation," *arXiv preprint arXiv:1404.6853*, May 2015.

[22] C. D. Gianelli, L. Xu, J. Li, and P. Stoica, "One-bit compressive sampling with time-varying thresholds for sparse parameter estimation," *9th IEEE Workshop on Sensor Array and Multi-channel Processing,* Rio De Janeiro, Brazil, July 10-13 2016.

[23] C. Gianelli, L. Xu, J. Li, and P. Stoica, "One-bit compressive sampling with time-varying thresholds: Maximum likelihood and the Cramér-Rao bound," *2016 Asilomar Conference on Signals, Systems and Computers,* Pacific Grove, CA, pp. 399–403, November 2016.

[24] P. Stoica and R. L. Moses, *Spectral Analysis of Signals*. Upper Saddle River, NJ: Prentice-Hall, 2005.

[25] J. Li and P. Stoica, "Efficient mixed-spectrum estimation with applications to target feature extraction," *IEEE Transactions on Signal Processing*, vol. 44, pp. 281–295, February 1996.

[26] G. Casella and R. L. Berger, *Statistical Inference* (2nd Edition). Pacific Grove, CA: Thomson Learning, Inc., 2002.

[27] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, United Kingdom: Cambridge University Press, 2004.

[28] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty, *Nonlinear Programming Theory and Algorithms, 3rd ed.* Hoboken, New Jersey: John Wiley & Sons, Inc., 2006.

[29] R. Tibshirani, "Regression shrinkage and selection via the LASSO," *Journal of the Royal Statistical Society*, vol. 58, no. 1, pp. 267–288, 1996.

[30] E. J. Candés, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *Journal of Fourier Analysis and Applications*, vol. 14, pp. 877–905, December 2008.