

1-1-2015

One-Bit Compressive Sensing with Partial Support

Phillip North
Claremont McKenna College

Deanna Needell
Claremont McKenna College

Recommended Citation

P. North and D. Needell. "One-Bit Compressive Sensing with Partial Support." Proc. IEEE International Workshop on Computational Advances in Multi-sensor Adaptive Processing, 2015.

This Conference Proceeding is brought to you for free and open access by the CMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in CMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

One-Bit Compressive Sensing with Partial Support

Phillip North
Dept. of Mathematical Sciences
Claremont McKenna College
Claremont CA 91711
Email: PNorth15@students.claremontmckenna.edu

Deanna Needell
Dept. of Mathematical Sciences
Claremont McKenna College
Claremont CA 91711
Email: dneedell@cmc.edu

Abstract—The Compressive Sensing framework maintains relevance even when the available measurements are subject to extreme quantization, as is exemplified by the so-called one-bit compressed sensing framework which aims to recover a signal from measurements reduced to only their sign-bit. In applications, it is often the case that we have some knowledge of the structure of the signal beforehand, and thus would like to leverage it to attain more accurate and efficient recovery. This work explores avenues for incorporating such partial-support information into the one-bit setting. Experimental results demonstrate that newly proposed methods of this work yield improved signal recovery even for varying levels of accuracy in the prior information. This work is thus the first to provide recovery mechanisms that efficiently use prior signal information in the one-bit reconstruction setting.

I. INTRODUCTION

Compressed Sensing (CS) addresses the problem of accurately acquiring high dimensional signals from a set of relatively few linear measurements [1] [2] [3]. The problem can be formulated mathematically via the system $\tilde{y} = \Phi x$, where $\Phi \in \mathbb{C}^{m \times n}$ is the *measurement matrix*. In the compressed setting, $m \ll n$ but one utilizes the assumption that the signal x possesses some additional structure, such as *sparsity*; we say that x is *k-sparse* when

$$\|x\|_0 := |\text{supp}(x)| = k \ll n.$$

CS has seen a vast amount of progress (see e.g. [4], [5]), and it is now well-known that for suitable matrices Φ (for example i.i.d. Gaussian), any k -sparse vector x can be recovered from $\tilde{y} \in \mathbb{C}^m$ when $m \approx k \log(n/k)$.

Unfortunately, the majority of theoretical work in CS assumes that the measurements are acquired with infinite precision whereas in practice they must be quantized. The extreme quantization setting where only the sign bit is acquired is known as *one-bit compressed sensing* [6]. In this framework, the measurements now take the form $y_i = \text{sign}(\langle x, \phi_i \rangle)$ where ϕ_i denotes the i th row of the measurement matrix Φ . Typically one then loses the ability to recover the magnitude of x and thus assumes the signal has a fixed norm (e.g. unit-norm), although there are adaptive techniques to overcome this as well [7], [8].

A. Existing one-bit methods

Although one-bit CS is a relatively new technology, efficient recovery algorithms have been studied. There are mainly two

D.N. was partially supported by NSF CAREER DMS-1348721 and the Alfred P. Sloan Foundation.

types of methods, those based on linear programming [9], [10], [11] and on iterative approaches [12], [13]. In this work we focus on the iterative approach, and in particular the Binary Iterative Hard Thresholding Algorithm (BIHT) [13], an extension of the traditional Iterative Hard Thresholding Algorithm (IHT) [14]. Assuming that our desired signal is k -sparse, the objective of BIHT is to return a solution that is k -sparse and consistent with the given sign measurements. Viewed as solving an optimization problem, at each iteration BIHT computes and takes a step in the direction of the gradient to attain a new approximation. This approximation is then thresholded to retain only the k largest in magnitude entries. Finally, after a consistent approximation is attained or enough iterations have elapsed, the estimation is normalized and returned. Algorithm 1 presents a more detailed explanation of BIHT. Here and throughout we use the thresholding function $\text{prune}(z, k)$ which returns the vector z with all but the k largest in magnitude entries set to zero.

Algorithm 1 Binary Iterative Hard Thresholding (BIHT). Given: measurement matrix Φ , one-bit measurements y , assumed sparsity level k , gradient step-size τ

```

1: procedure BIHT( $\Phi, y, k$ )
2:    $\tilde{x} = 0$  ▷ Initialize trivial approximation of  $x$ 
3:   repeat
4:      $\Gamma = \tilde{x} + \frac{\tau}{2} \Phi'(y - \text{sign}(\Phi \tilde{x}_i))$  ▷ Gradient step
5:      $\tilde{x} = \text{prune}(\Gamma, k)$  ▷ Hard threshold
6:   until halting criterion satisfied
7:   return  $\frac{\tilde{x}}{\|\tilde{x}\|_2}$  ▷ Normalize
8: end procedure

```

II. METHODS USING PARTIAL SUPPORT

In many applications, it is not only known that the signal of interest is sparse, but additional information about the support of the signal may also be known. For example, it is well-known that the support of wavelet representations of natural images largely resides in the low frequencies. In distributed settings, the signals of interest may be highly correlated so that partial support information may be obtained from neighboring atoms. This framework can be modeled by assuming that there is some partial estimate \hat{T} of the signal support $T := \text{supp}(x)$ a priori. The work of Mansour et al. [15] [16] extend the conventional ℓ_1 -minimization method to a *weighted* ℓ_1 approach that effectively incorporates such a

support estimate. To incorporate the estimate \tilde{T} , consider the following program:

$$\min_x \sum_{i=1}^n w_i |x_i| \quad \text{s.t.} \quad \tilde{y} = \Phi x \quad \text{with} \quad w_i = \begin{cases} c \in [0, 1] & \text{if } w_i \in \tilde{T} \\ 1 & \text{if } w_i \notin \tilde{T}. \end{cases}$$

This program now imposes a penalty for placing non-zero entries in locations not specified in the support estimate \tilde{T} . The value c in the weight vector w can be determined by the confidence in each element of \tilde{T} . We build upon this notion of weighting the estimation vector by developing analogous methods for the one-bit setting, focusing on iterative methods.

A. Oracle estimation

As a first, most basic approach, let us assume that our support estimate \tilde{T} is completely accurate, i.e. $\tilde{T} = T$. At the pruning step, no matter our result, we could simply set all entries of Γ not in \tilde{T} to zero; instead of locating the k largest entries of Γ , we would naively only retain the entries of \tilde{T} . The entries of our new estimate would be determined as follows:

$$\tilde{x}_i = \begin{cases} \Gamma_i & \text{if } i \in \tilde{T} \\ 0 & \text{if } i \in \tilde{T}^C. \end{cases}$$

As a similar approach, we could instead try soft-thresholding the entries of Γ not in \tilde{T} , multiplying the entries of Γ not in \tilde{T} by some constant $0 < c < 1$. Figure 1 shows the result of this approach; here and throughout, unless otherwise noted the signal length is $n = 256$, sparsity level $k = 8$, Φ has standard normal entries, $\tau = 0.001$, the support of the signal x is distributed uniformly at random, the magnitudes of the non-zero entries are standard normal, and the algorithm is run until the estimate changes by less than 10^{-10} or after 1000 iterations. Figure 1 shows the mean squared error (MSE) averaged over 100 trials for various values of m , the number of measurements. At each iteration, entries of Γ that are not in the support estimate are scaled down by a factor c , so that when $c = 0$ the result is simple hard-thresholding. Observe unsurprisingly how powerful a perfect support estimate can be via this bold hard-thresholding strategy. Figure 1 also shows that soft-thresholding with the perfect support estimate performs identically to hard-thresholding. This result is not hard to understand: after j iterations the elements not in \tilde{T} have been scaled down by c^j , which approaches zero after many iterations. Of course, knowing the full support beforehand is not likely, but these examples show that this information can seriously expedite the recovery of x .

B. Soft thresholding

In practice our support estimate will seldom equal the true support, and thus at a given iteration we should consider both \tilde{T} and the locations of the k greatest in magnitude entries of Γ , denoted \hat{T} , when updating our estimate. Let us assume that we know the support estimate to be $(\rho \times 100)\%$ accurate. Now, in the pruning step of BIHT the entries of Γ may be thresholded

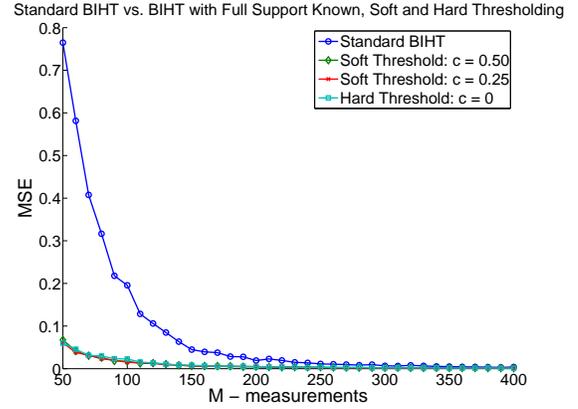


Fig. 1. Performance of BIHT when support estimate is exact and we employ soft and hard thresholding ($c = 0$).

according to four distinct possibilities:

$$\tilde{x}_i = \Gamma_i w_i \quad \text{where} \quad w_i = \begin{cases} 1 & \text{if } i \in \tilde{T} \cap \hat{T} \\ 1 & \text{if } i \in \tilde{T} \cap \hat{T}^C \\ 1 - \rho & \text{if } i \in \tilde{T}^C \cap \hat{T} \\ 0 & \text{if } i \in \tilde{T}^C \cap \hat{T}^C. \end{cases}$$

Additionally, we can use this 4-set framework when our support estimate includes erroneous elements. Suppose we have a support estimate that contains $(\rho \times k)$ correct elements but also includes $(1 - \rho) \times k$ incorrect elements. Figure 2 shows BIHT's performance when using this 4-set representation to incorporate prior support information. In the case of no false positives, when no elements of our support estimate are incorrect, we see that performance of BIHT is not bad: as ρ increases the MSE decays. Similarly, with the inclusion of false positives, performance is intuitive: improvements are seen when there are more correct estimates than incorrect estimates, i.e. $\rho \geq 0.5$; for lesser values of ρ the support estimate consists mostly of incorrect elements and performance is worse than standard BIHT. Also note that under this 4-set representation a k -sparse approximation is not necessarily returned. If $\tilde{T} \cap \hat{T} = \emptyset$, then in fact a $2k$ -sparse solution is returned, certainly this will result in a less accurate solution. Perhaps this is the reason for the slower rate of error decay in comparison with standard BIHT, as seen in Figure 2.

C. Supervised weighting

As a means for incorporating a partial support estimate and returning a k -sparse solution, we may use the weighting framework for traditional ℓ_1 -minimization as presented in [15]. We refer to this model as *supervised*, since the partial support estimate \tilde{T} must be obtained a priori, external to the algorithm itself. Again, suppose we believe our estimate \tilde{T} to be $(\rho \times 100)\%$ accurate. Let us create a weight vector w where

$$w_i = \begin{cases} 1 & \text{if } i \in \tilde{T} \\ 1 - \rho & \text{if } i \in \tilde{T}^C. \end{cases}$$

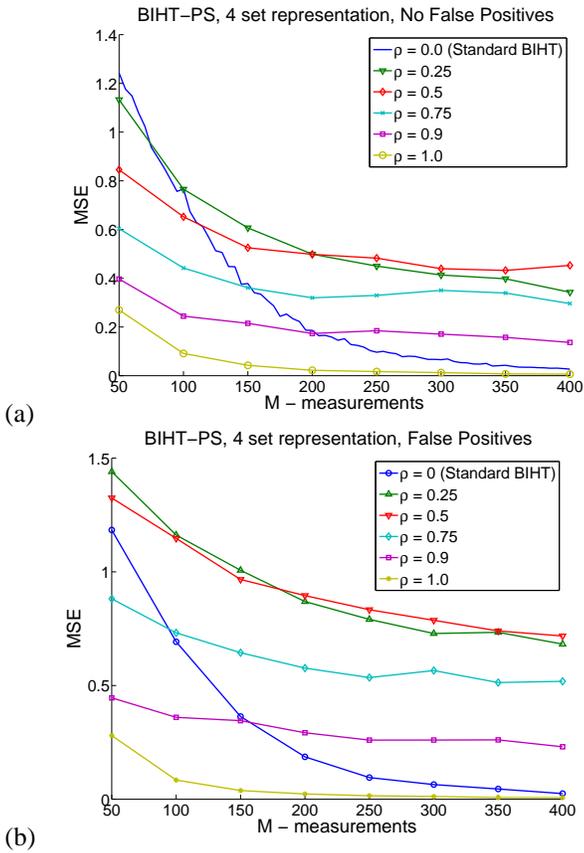


Fig. 2. Performance of BIHT using soft thresholding. In (a) we know the estimate to contain ρk correct elements (no false positives), in (b) there are an additional $(1 - \rho)k$ incorrect elements.

Then, at some iteration of BIHT, once we compute Γ let us multiply component-wise Γ and w to attain $\psi = \Gamma \odot w$. Now we locate the k largest elements of ψ , denoted $T_\psi := \text{prune}(\psi, k)$, and hard threshold the elements of $\Gamma \in T_\psi^C$, i.e. set them equal to zero. This approach is very similar to BIHT when there is no prior support estimate, except that here Γ is instead pruned to retain the k greatest in magnitude elements of $\Gamma \odot w$. A more thorough break down of this procedure is presented in Algorithm 2. We see that this supervised weighting approach outperforms the 4-set soft thresholding formulation both when false positives are and are not included. Figure 3 shows that for every value of m , when some prior information exists, *the weighting approach performs the same as or better than standard BIHT*.

This weighting framework for incorporating partial support information into BIHT (BIHT-PSW) performs well in the current context. However, we are assuming that the value for ρ is correct. This is a very bold assumption that, in practice, is not likely to hold. In the weighting step, the value of ρ determines how diminished the magnitude of an entry off the support estimate will be. If we are more confident in certain entries being non-zero, then other entries will be scaled by a constant closer to zero. Empirical results, as displayed in Figure 4, show that using the correct value of ρ is crucial to

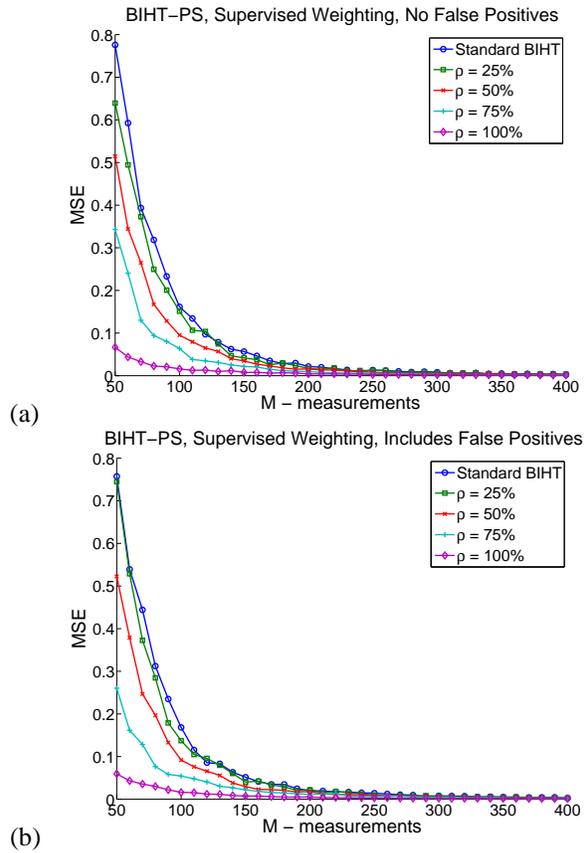


Fig. 3. Performance of BIHT when using supervised weighting. In (a) we know the estimate to contain ρk correct elements (no false positives), in (b) there are an additional $(1 - \rho)k$ incorrect elements.

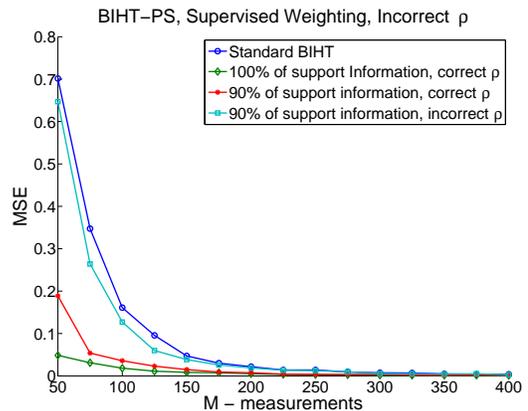


Fig. 4. Effect of inaccurate choice of ρ in supervised weighting. For incorrect ρ (cyan), $\rho = 0.1$ (rather than $\rho = 0.9$) was used in the weight vector.

the incorporation of a partial support estimate.

D. Unsupervised re-weighting

The improvements of this approach prompt one to ask whether such a result can be leveraged even when no support estimate is available. If we take m measurements of x and use BIHT to attain some estimate \tilde{x} , then we may use the support of \tilde{x} , denoted \tilde{T} , as an estimate for T . Then, we could

Algorithm 2 Binary Iterative Hard Thresholding with Partial Support Estimate Weighting (BIHT-PSW). Given: measurement matrix Φ , one-bit measurements y , sparsity level k , support estimate \tilde{T} , accuracy of support estimate ρ , step-size τ

```

1: procedure BIHT-PSW( $\Phi, y, k, \tilde{T}, \rho$ )
2:    $w_i = \begin{cases} (1 - \rho) & \text{if } i \in \tilde{T}^C \\ 1 & \text{if } i \in \tilde{T} \end{cases} \quad \triangleright \text{Construct weights}$ 
3:    $\tilde{x} = 0 \quad \triangleright \text{Initialize trivial approximation of } x$ 
4:   repeat
5:      $\Gamma = \tilde{x} + \frac{\tau}{2} \Phi'(y - \text{sign}(\Phi \tilde{x}^i)) \quad \triangleright \text{Gradient step}$ 
6:      $\Omega = \text{prune}(\Gamma \odot w, k) \quad \triangleright \text{Prune weighted update}$ 
7:      $\tilde{x}_j = \begin{cases} \Gamma_j & \text{if } j \in \Omega \\ 0 & \text{if } j \in \Omega^C \end{cases} \quad \triangleright \text{Update approximation}$ 
8:   until halting criterion is satisfied
9:   return  $\frac{\tilde{x}}{\|\tilde{x}\|_2} \quad \triangleright \text{Normalize}$ 
10: end procedure

```

use \tilde{T} to run BIHT-PSW, getting a more accurate approximation, and so on. This is reminiscent of the re-weighted ℓ_1 -minimization approach in classical compressed sensing [17], and is the general idea behind the BIHT Unsupervised Re-weighting (BIHT-URW) algorithm, presented as Algorithm 3. We refer to this model as *unsupervised*, since no outside information about the support is required. In this case, since ρ is certainly unknown, we utilize a parameter λ in place of ρ . The performance of BIHT-URW is displayed in Figure 5. Unfortunately, we observe no improvement from standard BIHT; we conjecture this is perhaps due to the arbitrary selection of λ (and thus ρ) within the method, which we have set to $\lambda = 0.5$ in Figure 5. As a benchmark for optimal performance of BIHT-URW we tested the method using a weight vector from the actual signal x (rather than its approximation), and then ran the algorithm as usual. This resulted in a significant improvement in performance, as shown in the cyan curve of Figure 5. Of course, this method is not applicable in practice, but demonstrates the potential of such an approach if better ways of estimating λ can be obtained, possibly adapting from iteration to iteration.

Algorithm 3 Binary Iterative Hard Thresholding with Unsupervised Re-weighting. Given: measurement matrix Φ , one-bit measurements y , sparsity level k , step-size τ , accuracy parameter λ , number of re-weighting iterations n

```

1: procedure BIHT-URW( $\Phi, y, k, n$ )
2:    $\tilde{T} = \text{supp}(\text{BIHT}(\Phi, y, k)) \triangleright \text{BIHT support estimate}$ 
3:   repeat
4:      $\tilde{x} = \text{BIHT-PSW}(\Phi, y, k, \tilde{T}, \lambda) \quad \triangleright \text{Estimate}$ 
5:      $\tilde{T} = \text{supp}(\tilde{x}) \quad \triangleright \text{Update support estimate}$ 
6:   until  $n$  iterations have completed
7:   return  $\frac{\tilde{x}}{\|\tilde{x}\|_2} \quad \triangleright \text{Normalize}$ 
8: end procedure

```

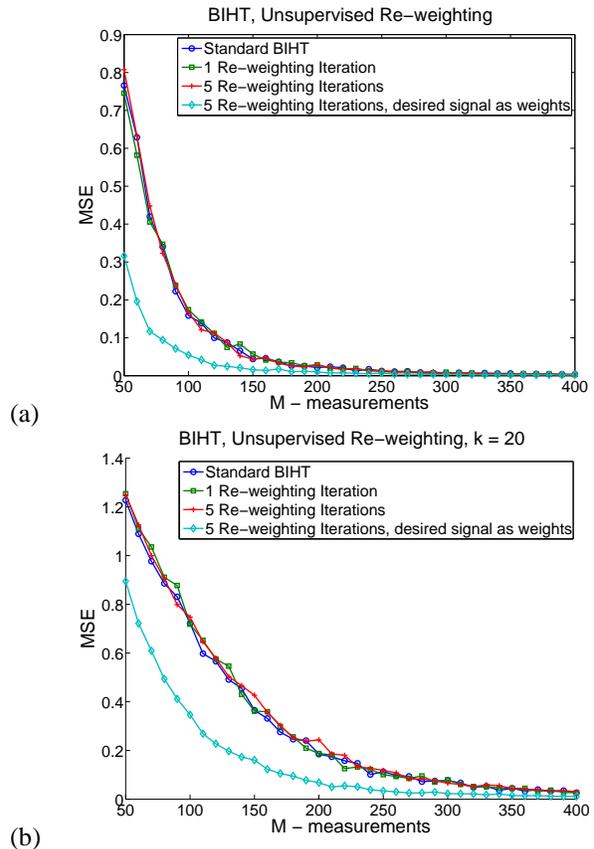


Fig. 5. Performance of BIHT-URW for various re-weighting iterations. In (a) the sparsity level is $k = 8$, and in (b) $k = 20$. In both figures, the cyan (\diamond) line is a result of creating a weight vector out of the desired signal x .

III. CONCLUSION

We presented several weighting variants of the BIHT algorithm for one-bit CS when partial support estimation is known. We demonstrate that our methods effectively utilize the support information, but that leveraging support estimates in an unsupervised fashion is not straightforward. We believe future work in this area could lead to re-weighted methods which improve upon existing approaches.

REFERENCES

- [1] E. J. Candès and T. Tao, “Near optimal signal recovery from random projections: Universal encoding strategies?” *IEEE T. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.
- [2] D. L. Donoho, “Compressed sensing,” Oct. 2004, unpublished manuscript.
- [3] E. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete fourier information,” *IEEE T. Inform. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [4] Y. C. Eldar and G. Kutyniok, *Compressed sensing: theory and applications*. Cambridge University Press, 2012.
- [5] S. Foucart and H. Rauhut, *A mathematical introduction to compressive sensing*. Birkhäuser, 2013.
- [6] P. T. Boufounos and R. G. Baraniuk, “1-bit compressive sensing,” in *42nd Annual Conference on Information Sciences and Systems (CISS)*. IEEE, 2008, pp. 16–21.
- [7] K. Knudson, R. Saab, and R. Ward, “One-bit compressive sensing with norm estimation,” *arXiv preprint arXiv:1404.6853*, 2014.

- [8] R. Baraniuk, S. Foucart, D. Needell, Y. Plan, and M. Wootters, "Exponential decay of reconstruction error from binary measurements of sparse signals," *arXiv preprint arXiv:1407.8246*, 2014.
- [9] Y. Plan and R. Vershynin, "One-bit compressed sensing by linear programming," *Communications on Pure and Applied Math*, 2011, to appear.
- [10] —, "Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach," *IEEE T. Inform.Theory*, 2012, to appear.
- [11] S. Gopi, P. Netrapalli, P. Jain, and A. Nori, "One-bit compressed sensing: Provable support and vector recovery," in *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, 2013, pp. 154–162.
- [12] L. Jacques, K. Degraux, and C. De Vleeschouwer, "Quantized iterative hard thresholding: Bridging 1-bit and high-resolution quantized compressed sensing," *arXiv preprint arXiv:1305.1786*, 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," *IEEE T. Inform. Theory*, vol. 59, no. 4, pp. 2082–2102, 2011.
- [14] T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *Appl. Comput. Harmon. A.*, vol. 27, no. 3, pp. 265–274, 2009.
- [15] H. Mansour and R. Saab, "Recovery analysis for weighted ℓ_1 -minimization using a null space property," *CoRR*, vol. abs/1412.1565, 2014. [Online]. Available: <http://arxiv.org/abs/1412.1565>
- [16] M. P. Friedlander, H. Mansour, R. Saab, and O. Yilmaz, "Recovering compressively sampled signals using partial support information," *CoRR*, vol. abs/1010.4612, 2010. [Online]. Available: <http://arxiv.org/abs/1010.4612>
- [17] E. J. Candès, M. Wakin, and S. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J.Fourier Anal.Appl.*, vol. 14, no. 5, pp. 877–905, Dec. 2008.