

Compressive Binary Search

Mark A. Davenport
Department of Statistics
Stanford University

Ery Arias-Castro
Department of Mathematics
University of California, San Diego

Abstract—In this paper we consider the problem of locating a nonzero entry in a high-dimensional vector from possibly adaptive linear measurements. We consider a recursive bisection method which we dub the *compressive binary search* and show that it improves on what any nonadaptive method can achieve. We also establish a non-asymptotic lower bound that applies to all methods, regardless of their computational complexity. Combined, these results show that the compressive binary search is within a double logarithmic factor of the optimal performance.

I. INTRODUCTION

How should one approach the problem of *finding a needle in a haystack*? Specifically, suppose that a high-dimensional vector $\mathbf{x} \in \mathbb{R}^n$ is known to have a single nonzero entry—how can we efficiently find the location of the nonzero? We will assume that we can learn about \mathbf{x} by taking m noisy linear measurements of the form

$$y_i = \langle \mathbf{a}_i, \mathbf{x} \rangle + z_i, \quad i = 1, \dots, m, \quad (1)$$

where the measurement vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$ have Euclidean norm at most 1 and z_1, \dots, z_m are i.i.d. according to $\mathcal{N}(0, 1)$. Our question reduces to the problem of choosing the vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$ and constructing an algorithm to estimate the location of the nonzero from the measurements y_1, \dots, y_m .

This is a special case of support recovery in compressive sensing (CS) [1, 2], since (1) is equivalent to the linear model

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{z}, \quad (2)$$

where $\mathbf{y} = (y_1, \dots, y_m)$, \mathbf{A} is the $m \times n$ matrix with row vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$ and $\mathbf{z} = (z_1, \dots, z_m)$. (Note that the *rows* of \mathbf{A} are normalized, as opposed to the columns, which is another common convention in the CS literature.) There are a variety of results on support recovery in the context of (2) where the measurement matrix \mathbf{A} is fixed in advance (i.e., is nonadaptive) and satisfies certain desirable properties [3–9]. As an example, one can show that if \mathbf{A} is generated by drawing i.i.d. $\pm 1/\sqrt{n}$ (symmetric) entries and the signal \mathbf{x} is 1-sparse with nonzero entry equal to $\mu > 0$, then the Lasso and Orthogonal Matching Pursuit (OMP) recover the support of \mathbf{x} with high probability provided that

$$\mu \geq C\sqrt{(n/m)\log n}, \quad (3)$$

with C sufficiently large. Moreover, any method based on such measurements requires μ to satisfy this lower bound for some constant $C > 0$ [10]. This is essentially the whole story when the measurements are nonadaptive.

In contrast, suppose now that the system implementing (1) can provide feedback in such a way as to allow for the measurements to be taken adaptively, meaning that \mathbf{a}_i may be chosen as a function of the observations up to time $i - 1$, that is, (y_1, \dots, y_{i-1}) . (This implicitly assumes that \mathbf{a}_i is a deterministic function of this vector, but there is no loss of generality in this assumption. See [11] for details.) This instance of active or online learning has received comparatively far less attention to date. However, in recent work [11] we have established lower bounds showing that *any* support recovery method under *any* adaptive sampling scheme (satisfying the conditions above) will be unable to recover the correct support unless the nonzero entry satisfies

$$\mu \geq C\sqrt{n/m}, \quad (4)$$

for some constant $C > 0$.

Our contribution in this paper is twofold. In Section II, we propose a *compressive binary search* algorithm which recursively tests whether the nonzero entry is on the left or right half of the current interval. We show that the method reliably recovers the support of a 1-sparse vector when the nonzero entry satisfies

$$\mu \geq C\sqrt{(n/m)\log \log_2 n}, \quad (5)$$

with a constant $C > 2$. We then verify this analysis via numerical simulations. Note that by using an adaptive measurement scheme we are able to improve upon the requirement in (3) by reducing the $\log n$ to $\log \log_2 n$, but our scheme does not eliminate the logarithmic factor entirely as in (4). A corollary of this result is that in contrast to the results of [11], which argued that in general adaptive strategies do not improve over nonadaptive strategies in terms of our ability to accurately recover \mathbf{x} , we see that when μ satisfies (5), adaptive strategies can significantly outperform nonadaptive ones by first identifying the location of the nonzero and then reserving a set of measurements to more accurately estimate the value of the nonzero.

In contrast to this upper bound, in Section III, we provide a simple proof that $\mu \geq C\sqrt{n/m}$ is necessary for any method to work. This novel proof is in some sense tailored to this binary method as it too is based on testing whether the nonzero component is in the left or right half of \mathbf{x} . In Section IV, we discuss related work in more detail and directions for future work.

II. COMPRESSIVE BINARY SEARCH

A. The algorithm

The algorithm is designed assuming that the target vector \mathbf{x} has exactly one nonzero entry equal to $\mu > 0$; both the location and magnitude are unknown. The methodology described here can be easily adapted to the case where the sign of the nonzero entry is unknown. For simplicity, we assume that n is dyadic, and let $s_0 = \log_2 n$, where \log_2 denotes the logarithm in base 2.

With a budget of $m \geq 2\log_2 n$ measurements of the form (1), the binary search method proceeds as follows. We divide our m measurements into a total of s_0 stages, allocating m_s measurements to stage s , where

$$m_s := \tilde{m}_s + 1, \quad \tilde{m}_s := \lfloor (m - s_0)2^{-s} \rfloor, \quad (6)$$

where $\lfloor a \rfloor$ denotes the largest integer not greater than a . Note that we do not exceed our total measurement budget since

$$\sum_{s=1}^{s_0} m_s = s_0 + \sum_{s=1}^{s_0} \tilde{m}_s \leq s_0 + (m - s_0) \sum_{s=1}^{s_0} 2^{-s} \leq m.$$

We also have $m_s \geq 1$ for all s , which is necessary for our algorithm to be able to run to completion. Starting with $J_0^{(1)} := \{1, \dots, n\}$, at stage $s = 1, \dots, s_0$, we have a dyadic interval $J_0^{(s)}$ and consider its left and right halves denoted $J_1^{(s)}$ and $J_2^{(s)}$. For example, $J_1^{(1)} := \{1, \dots, \frac{n}{2}\}$ and $J_2^{(1)} := \{\frac{n}{2} + 1, \dots, n\}$. Let $\mathbf{u}^{(s)}$ denote the vector with entries indexed by $J_1^{(s)}$ equal to $2^{-(s_0-s+1)/2}$ and with entries indexed by $J_2^{(s)}$ equal to $-2^{-(s_0-s+1)/2}$. Note that $\|\mathbf{u}^{(s)}\| = 1$, since $|J_1^{(s)}| = |J_2^{(s)}| = 2^{s_0-s}$. We measure m_s times with $\mathbf{u}^{(s)}$, meaning that we observe

$$y_i^{(s)} = \langle \mathbf{u}^{(s)}, \mathbf{x} \rangle + z_i^{(s)}, \quad i = 1, \dots, m_s.$$

Based on these measurements, we decide between going left or right, meaning we test whether the nonzero entry is in $J_1^{(s)}$ or $J_2^{(s)}$. We do so by simply computing

$$w^{(s)} = \sum_{i=1}^{m_s} y_i^{(s)}.$$

Specifically, we set $J_0^{(s+1)} = J_1^{(s)}$ if $w^{(s)} > 0$, and $J_0^{(s+1)} = J_2^{(s)}$ otherwise.

B. Performance analysis

The binary search improves on methods based on non-adaptive measurements by weakening the requirement (3) to (5).

Theorem 1. *In our setting, with a single nonzero entry equal to $\mu > 0$ and a measurement budget of $m \geq 2\log_2 n$, the probability that binary search fails to locate the nonzero entry (denoted \mathbb{P}_e) satisfies*

$$\mathbb{P}_e \leq \frac{\log_2 n}{2} \exp\left(-\frac{\mu^2 m}{8n}\right). \quad (7)$$

Proof: Since the binary search algorithm is equivariant with respect to the ordering of the entries, we can begin by assuming without loss of generality that $\mathbf{x} = (\mu, 0, \dots, 0)^T$, i.e., the nonzero is located in the first entry of \mathbf{x} . Thus, we can use a simple union bound to argue that

$$\mathbb{P}_e \leq \sum_{s=1}^{s_0} \mathbb{P}(w^{(s)} < 0), \quad (8)$$

where $J_1^{(s)} = \{1, \dots, 2^{-s}n\}$ and $J_2^{(s)} = \{2^{-s}n + 1, \dots, 2^{1-s}n\}$. Under our assumptions, we have that

$$w^{(s)} \sim \mathcal{N}\left(2^{(s-1)/2} \frac{m_s \mu}{\sqrt{n}}, m_s\right).$$

Thus we can bound

$$\begin{aligned} \mathbb{P}(w^{(s)} < 0) &= \bar{\Phi}\left(\mu \cdot \sqrt{\frac{m_s 2^s}{2n}}\right) \\ &\leq \frac{1}{2} \exp\left(-\frac{m_s \mu^2 2^s}{n}\right), \end{aligned} \quad (9)$$

since for all $t > 0$ we have

$$\bar{\Phi}(t) := \mathbb{P}(\mathcal{N}(0, 1) > t) \leq \frac{1}{2} \exp(-t^2/2).$$

We next note that by construction,

$$\tilde{m}_s + 1 \geq (m - s_0)2^{-s}.$$

Since $m \geq 2s_0$, we have that $m - s_0 \geq m/2$, and hence we obtain

$$m_s 2^s \geq (\tilde{m}_s + 1)2^s \geq m - s_0 \geq m/2.$$

Plugging $m_s 2^s \geq m/2$ into (9), we obtain

$$\mathbb{P}(w^{(s)} < 0) \leq \frac{1}{2} \exp\left(-\frac{\mu^2 m}{8n}\right).$$

Plugging this into (8) we arrive at

$$\mathbb{P}_e \leq \frac{s_0}{2} \exp\left(-\frac{\mu^2 m}{8n}\right),$$

as desired. \blacksquare

Note that we need (5) with $C > 2\sqrt{2}$ for the upper bound on \mathbb{P}_e in (7) to actually tend to zero as n increases. However, by taking additional measurements beyond the $2\log_2 n$ required by this theorem, we could loosen this requirement to be able to set C arbitrarily close to 2.




Fig. 1. Comparison between compressive binary search and OMP as a function of μ for $n = 4096$ and $m = 256$. Observe that compressive binary search can successfully identify the location of the nonzero for weaker values of μ than OMP, but still requires $\mu > \sqrt{n/m} = 4$.

C. Numerical experiments

To validate our theory, we perform some simple numerical experiments. Specifically, we compare the performance of the compressive binary search procedure to that of OMP (with \mathbf{A} constructed with random $\pm 1/\sqrt{n}$ entries). Note that in the 1-sparse case, OMP simply reduces to identifying the column of \mathbf{A} most highly correlated with the measurements \mathbf{y} . The performance of these two algorithms is shown in Figure 1, which shows the empirical probability of error as a function of μ computed by averaging over 100,000 trials. For these experiments, we set $n = 4096$ and $m = 256$. Note that for these values of n and m , we have that $\sqrt{n/m} = 4$ and $\sqrt{(n/m) \log \log_2 n} \approx 6.3$. Thus, ignoring the constant terms in (4) and (5), we see that the performance of the compressive binary search is largely consistent with our theory—namely, it cannot reliably identify the location of the nonzero when $\mu \leq 4$ but can for $\mu \gtrsim 6.3$. Moreover, recall that as noted in (3), the nonadaptive OMP algorithm requires that μ exceed $\sqrt{(n/m) \log n} \approx 11.5$ to succeed. Again ignoring constants, in our case this corresponds to requiring μ to be roughly 1.8 times larger than is required for the compressive binary search procedure, and this is precisely the behavior we observe in Figure 1.

III. LOWER BOUND: LEFT OR RIGHT?

We now establish an explicit, non-asymptotic lower bound for adaptive support recovery, valid for any recovery method based on adaptive measurements satisfying the conditions required here. Though such bounds were recently derived in [11], we provide a slightly simpler proof here for the case of 1-sparse signals that closely

aligns with the core idea of the compressive binary search.

Let $y_{[i]} = (y_1, \dots, y_i)$ denote the information available after taking i measurements. Let $\mathbb{P}_{\mathbf{x}}$ denote the distribution of these measurements when the target vector is \mathbf{x} . Without loss of generality, we assume that \mathbf{a}_i is a deterministic function of $y_{[i-1]}$. In that case, using the fact that y_i is independent of $y_{[i-1]}$ when \mathbf{a}_i is given, we have

$$\mathbb{P}_{\mathbf{x}}(y_{[m]}) = \prod_{i=1}^m \mathbb{P}_{\mathbf{x}}(y_i | \mathbf{a}_i). \quad (10)$$

For a subset $K \subset \{1, \dots, n\}$, let $K^c := \{1, \dots, n\} \setminus K$ and let \mathbf{x}_K be the part of \mathbf{x} indexed by K .

Let $\|\mathbb{P} - \mathbb{Q}\|_{\text{TV}}$ denote the total variation metric between distributions \mathbb{P} and \mathbb{Q} , and $K(\mathbb{P}, \mathbb{Q})$ their Kullback-Leibler divergence [12], related by Pinsker's inequality

$$\|\mathbb{P} - \mathbb{Q}\|_{\text{TV}}^2 \leq \frac{1}{2} K(\mathbb{P}, \mathbb{Q}). \quad (11)$$

Lemma 1. Suppose that n is even and let $J_1 = \{1, \dots, n/2\}$ and $J_2 = \{n/2 + 1, \dots, n\}$. For $r = 1, 2$, let π_r denote the uniform prior on the vectors $\mathbf{x} \in \mathbb{R}^n$ having a single nonzero entry equal to $\mu > 0$, located in J_r . Let \mathbb{P}_r denote the distribution of $y_{[m]}$ when $\mathbf{x} \sim \pi_r$. Then

$$\|\mathbb{P}_2 - \mathbb{P}_1\|_{\text{TV}}^2 \leq \frac{\mu^2 m}{n}.$$

Proof: Let \mathbb{P}_0 denote the distribution of $y_{[m]}$ when $\mathbf{x} = 0$, which is multivariate normal with zero mean and covariance \mathbf{I} . Using Pinsker's inequality (11), we have

$$\begin{aligned} \|\mathbb{P}_2 - \mathbb{P}_1\|_{\text{TV}}^2 &\leq 2\|\mathbb{P}_0 - \mathbb{P}_1\|_{\text{TV}}^2 + 2\|\mathbb{P}_0 - \mathbb{P}_2\|_{\text{TV}}^2 \\ &\leq K(\mathbb{P}_0, \mathbb{P}_1) + K(\mathbb{P}_0, \mathbb{P}_2). \end{aligned}$$

Let $\mathbb{P}_{(j)}$ denote the distribution of $y_{[m]}$ when the nonzero entry (equal to μ) is at $j \in \{1, \dots, n\}$. By the law of total probability,

$$\mathbb{P}_1 = \frac{2}{n} \sum_{j \in J_1} \mathbb{P}_{(j)},$$

and obviously

$$\mathbb{P}_0 = \frac{2}{n} \sum_{j \in J_1} \mathbb{P}_0,$$

which allows us to use the convexity of the KL divergence [13], to obtain

$$K(\mathbb{P}_0, \mathbb{P}_1) \leq \frac{2}{n} \sum_{j \in J_1} K(\mathbb{P}_0, \mathbb{P}_{(j)}).$$

Under $\mathbb{P}_{(j)}$, $y_i = \mu a_{i,j} + z_i$, while under \mathbb{P}_0 , $y_i = z_i$, so that

$$\begin{aligned} K(\mathbb{P}_0, \mathbb{P}_{(j)}) &= -\mathbb{E}_0 \log \frac{\mathbb{P}_{(j)}}{\mathbb{P}_0} \\ &= \sum_{i=1}^m \mathbb{E}_0 \left(\frac{1}{2} (y_i - \mu a_{i,j})^2 - \frac{1}{2} y_i^2 \right) \\ &= \sum_{i=1}^m \mathbb{E}_0 (-y_i \mu a_{i,j} + (\mu a_{i,j})^2 / 2) \\ &= \frac{\mu^2}{2} \sum_{i=1}^m \mathbb{E}_0 a_{i,j}^2. \end{aligned}$$

The first line is by definition; the second and third are consequences of (10) and the definition of the normal likelihood; the fourth line is because, under \mathbb{P}_0 , y_i is independent of $a_{i,j}$ and has zero mean. Hence,

$$K(\mathbb{P}_0, \mathbb{P}_1) \leq \frac{\mu^2}{n} \sum_{i=1}^m \mathbb{E}_0 \sum_{j \in J_1} a_{i,j}^2,$$

and similarly,

$$K(\mathbb{P}_0, \mathbb{P}_2) \leq \frac{\mu^2}{n} \sum_{i=1}^m \mathbb{E}_0 \sum_{j \in J_2} a_{i,j}^2,$$

so that

$$K(\mathbb{P}_0, \mathbb{P}_1) + K(\mathbb{P}_0, \mathbb{P}_2) \leq \frac{\mu^2}{n} \sum_{i=1}^m \mathbb{E}_0 \sum_{j=1}^n a_{i,j}^2 \leq \frac{\mu^2 m}{n},$$

since $\|\mathbf{a}_i\| \leq 1$ for all i . ■

Lemma 1 implies a lower bound on the risk of the problem of testing whether a vector $\mathbf{x} \in \mathbb{R}^n$ with a single nonzero entry equal to μ is supported on the first half or second half of the index set $\{1, \dots, n\}$. Proving this result by directly looking at the likelihood ratio, which would be the standard approach, seems quite delicate as we are testing a mixture (supported on the first half) versus a mixture (supported on the second half).

Theorem 2. *In the setting of Lemma 1, consider testing H_1 versus H_2 , where H_τ is the hypothesis by which \mathbf{x} is supported in J_τ . Then under the uniform prior, for any test T ,*

$$\mathbb{P}(T \text{ fails}) \geq \frac{1}{2} \left(1 - \mu \sqrt{m/n} \right).$$

Note that the lower bound is also valid in the minimax sense. In fact, the uniform prior is least favorable by invariance consideration [14, Sec. 8.4].

Proof: Under the uniform prior, we are effectively testing \mathbb{P}_1 versus \mathbb{P}_2 . The likelihood ratio test, which rejects when $L > 1$, with $L := \mathbb{P}_2 / \mathbb{P}_1$, has minimum risk, bounded by

$$\frac{1}{2} (1 - \|\mathbb{P}_2 - \mathbb{P}_1\|_{\text{TV}}).$$

(See Lemma 1 of [11].) We then apply Lemma 1 to bound the total variance distance on the RHS. ■

IV. DISCUSSION

Our main results can be cast as follows: Theorem 1 implies that, with probability at least $1/4$, the binary search method locates the nonzero entry (for $n \geq 4$) if

$$\mu \geq 4 \sqrt{\log \log_2 n} \sqrt{\frac{n}{m}},$$

while Theorem 2 shows that any method for locating the nonzero entry fails with probability at least $1/4$ when

$$\mu \leq \frac{1}{2} \sqrt{\frac{n}{m}}.$$

Clearly, the bounds do not match. Numerically, for $n \leq 10^6$, $\log \log_2 n \leq 3$, in which case the discrepancy is a multiplicative factor of $8\sqrt{3} < 14$.

We will return to the issue of whether this gap can be closed below, but first we wish to discuss an additional implication of Theorem 2. Specifically, one can show that Theorem 2 implies that for any estimator \hat{S} of the support of \mathbf{x} , $\mathbb{E}|\hat{S} \Delta S| \geq (1 - \mu \sqrt{m/n})$. Following the same argument as in Theorem 2 of [11], this implies that under the measurement model in (2) we have

$$\inf_{\hat{\mathbf{x}}} \sup_{\mathbf{x}: 1\text{-sparse}} \frac{1}{n} \mathbb{E} \|\hat{\mathbf{x}} - \mathbf{x}\|^2 \geq C \frac{1}{m},$$

where $C = 1/27$. In contrast, Theorem 1 implies that for sufficiently large μ , there exist estimators that do far better than this bound (by a factor of n).

While the problems of estimating or detecting the support of a 1-sparse vector might seem to have only limited applications, in fact one can extend any algorithm that identifies the support of a 1-sparse vector to one that works for vectors with $k \geq 2$ nonzero entries. This can be done by first exploiting a simple hashing scheme which will (with high probability) isolate each nonzero, and then applying the method for 1-sparse vectors to each hash separately. For an overview of this approach in a similar context, see [4].

We also note that [4] independently proposes a method very similar to the compressive binary search approach we describe. Though [4] considers a different setting with continuous signals (instead of vectors as we do), the method proposed is essentially the same, except that the measurement budget is partitioned differently. In particular, it is not obvious to us that the strategy in [4] will always succeed, since it does not account for rounding effects or enforce that a base number of measurements are reserved for each scale (stage) and so (to the best of our understanding) the method might exhaust its measurement budget before the algorithm terminates. Another key difference is that by considering the simpler setting of a vector in \mathbb{R}^n , we can significantly simplify the analysis. That being said, the conclusions of [4] are broadly similar to our own.

Finally, we also note that there are a few other adaptive algorithms that have been proposed in this setting. For example, [15] proposes an algorithm similar to the compressive binary search procedure but using a different procedure for allocating measurements to each stage. As another example, the Compressive Distilled Sensing (CDS) algorithm proposed in [16] considers a CS sampling algorithm which performs sequential subset selection via the random projections typical of CS. In a different direction, [17, 18] suggest Bayesian approaches where the measurement vectors are sequentially chosen so as to maximize the conditional differential entropy of y_i given (y_1, \dots, y_{i-1}) . While it remains a challenge to obtain performance bounds for the Bayesian methods suggested in [17, 18], CDS is analyzed in detail in [16] for the task of estimating a k -sparse vector \mathbf{x} . Following the proof with a view on support recovery, one can establish that CDS is reliable in our context when

$$\mu \geq C_n \sqrt{n/m},$$

with $C_n \rightarrow \infty$ arbitrarily slowly, coming extremely close to the lower bound in (4). However, the algorithm seems to require that $m \geq n^\alpha$ for a constant $\alpha > 0$ fixed, while binary search only requires $m \geq 2 \log_2 n$.

An important question would seem to be whether there exist methods which allow for both small m and μ approaching the bound in (4). After the submission of this paper, Malloy and Nowak proposed a slight modification of the compressive binary search approach (involving a different allocation of the measurements to each stage) which answers this question [19]. Specifically, [19] removes the $\log \log_2 n$ term at the cost of a slightly worse constant. Thus, the gap between the lower bound in (4) and the upper bound in [19] differs only by a constant factor. It would be interesting to know whether either of these bounds can be tightened.

ACKNOWLEDGEMENTS

Thanks to E. Candès for many insightful discussions. M. D. is supported by NSF grant DMS-1004718. E. A-C. is partially by ONR grant N00014-09-1-0258.

REFERENCES

- [1] R. Baraniuk, “Compressive sensing,” *IEEE Signal Processing Mag.*, vol. 24, no. 4, pp. 118–120, 124, 2007.
- [2] E. Candès and M. Wakin, “An introduction to compressive sampling,” *IEEE Signal Processing Mag.*, vol. 25, no. 2, pp. 21–30, 2008.
- [3] N. Verzelen, “Minimax risks for sparse regressions: Ultra-high-dimensional phenomena,” *Electron. J. Statist.*, vol. 6, pp. 38–90, 2012.
- [4] M. Iwen, “Group testing strategies for recovery of sparse signals in noise,” in *Proc. Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2009.
- [5] S. Aeron, V. Saligrama, and M. Zhao, “Information theoretic bounds for compressed sensing,” *IEEE Trans. Inform. Theory*, vol. 56, no. 10, pp. 5111–5130, 2010.
- [6] M. Wainwright, “Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting,” *IEEE Trans. Inform. Theory*, vol. 55, no. 12, pp. 5728–5741, 2009.
- [7] T. Cai and L. Wang, “Orthogonal matching pursuit for sparse signal recovery with noise,” *IEEE Trans. Inform. Theory*, vol. 57, no. 7, pp. 4680–4688, 2011.
- [8] P. Zhao and B. Yu, “On model selection consistency of lasso,” *J. Mach. Learn. Res.*, vol. 7, pp. 2541–2563, 2006.
- [9] E. Candès and Y. Plan, “Near-ideal model selection by ℓ_1 minimization,” *Ann. Stat.*, vol. 37, no. 5A, pp. 2145–2177, 2009.
- [10] E. Candès and M. Davenport, “How well can we estimate a sparse vector?” *Arxiv preprint arXiv:1104.5246*, 2011.
- [11] E. Arias-Castro, E. J. Candès, and M. A. Davenport, “On the fundamental limits of adaptive sensing,” *Arxiv preprint arXiv:1111.4646*, 2011.
- [12] P. Massart, *Concentration inequalities and model selection*, ser. Lecture Notes in Mathematics. Berlin: Springer, 2007, vol. 1896.
- [13] T. Cover and J. Thomas, *Elements of information theory*. Hoboken, NJ: Wiley-Interscience, 2006.
- [14] E. Lehmann and J. Romano, *Testing statistical hypotheses*, ser. Springer Texts in Statistics. New York: Springer, 2005.
- [15] J. Haupt, R. Nowak, and R. Castro, “Adaptive sensing for sparse signal recovery,” in *Proc. Digital Signal Processing Workshop*, Marco Island, FL, Jan. 2009.
- [16] J. Haupt, R. Baraniuk, R. Castro, and R. Nowak, “Compressive distilled sensing: Sparse recovery using adaptivity in compressive measurements,” in *Proc. Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2009.
- [17] R. Castro, J. Haupt, R. Nowak, and G. Raz, “Finding needles in noisy haystacks,” in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Las Vegas, NV, Apr. 2008.
- [18] S. Ji, Y. Xue, and L. Carin, “Bayesian compressive sensing,” *IEEE Trans. Signal Processing*, vol. 56, no. 6, pp. 2346–2356, 2008.
- [19] M. Malloy and R. Nowak, “Near-optimal compressive binary search,” *Arxiv preprint arXiv:1203.1804*, 2012.