On the Fundamental Limits of Adaptive Sensing

Ery Arias-Castro[*], Emmanuel J. Candès[†] and Mark A. Davenport [‡]

November 2011 (Revised May 2012)

Abstract

Suppose we can sequentially acquire arbitrary linear measurements of an n-dimensional vector $\mathbf{x}$ resulting in the linear model $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{z}$, where $\mathbf{z}$ represents measurement noise. If the signal is known to be sparse, one would expect the following folk theorem to be true: choosing an *adaptive* strategy which cleverly selects the next row of $\mathbf{A}$ based on what has been previously observed should do far better than a *nonadaptive* strategy which sets the rows of $\mathbf{A}$ ahead of time, thus not trying to learn anything about the signal in between observations. This paper shows that the folk theorem is false. We prove that the advantages offered by clever adaptive strategies and sophisticated estimation procedures—no matter how intractable—over classical compressed acquisition/recovery schemes are, in general, minimal.

Keywords: sparse signal estimation, adaptive sensing, compressed sensing, support recovery, information bounds, hypothesis tests.

1 Introduction

This paper is concerned with the fundamental question of how well one can estimate a sparse vector from noisy linear measurements in the general situation where one has the flexibility to design those measurements at will (in the language of statistics, one would say that there is nearly complete freedom in designing the experiment). This question is of importance in a variety of sparse signal estimation or sparse regression scenarios, but perhaps arises most naturally in the context of compressive sensing (CS) [4, 5, 11]. In a nutshell, CS asserts that it is possible to reliably acquire sparse signals from just a few linear measurements selected a priori. More specifically, suppose we wish to acquire a sparse signal $\mathbf{x} \in \mathbb{R}^n$. A possible CS acquisition protocol would proceed as follows. (i) Pick an $m \times n$ random projection matrix $\mathbf{A}$ (the first m rows of a random unitary matrix) in advance, and collect data of the form

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{z}, \tag{1.1}$$

where $\mathbf{z}$ is a vector of errors modeling the fact that any real world measurement is subject to at least a small amount of noise. (ii) Recover the signal by solving an ℓ_1 minimization problem such as the Dantzig selector [6] or the LASSO [28]. As is now well known, theoretical results guarantee that such convex programs yield accurate solutions. In particular, when $\mathbf{z} = 0$, the recovery is exact, and the error degrades gracefully as the noise level increases.

A remarkable feature of the CS acquisition protocol is that the sensing is completely nonadaptive; that is to say, no effort whatsoever is made to understand the signal. One simply selects a collection $\{\mathbf{a}_i\}$

[*]Department of Mathematics, University of California, San Diego {eariasca@ucsd.edu}
[†]Departments of Mathematics and Statistics, Stanford University {candes@stanford.edu}
[‡]Department of Statistics, Stanford University {markad@stanford.edu}

of sensing vectors a priori (the rows of the matrix $\mathbf{A}$), and measures correlations between the signal and these vectors. One then uses numerical optimization—e.g., linear programming [6]—to tease out the sparse signal $\mathbf{x}$ from the data vector $\mathbf{y}$. While this may make sense when there is no noise, this protocol might draw some severe skepticism in a noisy environment. To see why, note that in the scenario above, most of the power is actually spent measuring the signal at locations where there is no information content, i.e., where the signal vanishes. Specifically, let $\mathbf{a}$ be a row of the matrix $\mathbf{A}$ which, in the scheme discussed above, has uniform distribution on the unit sphere. The dot product is

$$\langle \mathbf{a}, \mathbf{x} \rangle = \sum_{j=1}^{n} a_j x_j,$$

and since most of the coordinates x_j are zero, one might think that most of the power is wasted. Another way to express all of this is that by design, the sensing vectors are approximately orthogonal to the signal, yielding measurements with low signal power or a poor signal-to-noise ratio (SNR).

The idea behind adaptive sensing is that one should localize the sensing vectors around locations where the signal is nonzero in order to increase the SNR, or equivalently, not waste sensing power. In other words, one should try to "learn" as much as possible about the signal while acquiring it in order to design more effective subsequent measurements. Roughly speaking, one would (*i*) detect those entries which are nonzero or significant, (*ii*) progressively localize the sensing vectors on those entries, and (*iii*) estimate the signal from such localized linear functionals. This is akin to the game of 20 questions in which the search is narrowed by formulating the next question in a way that depends upon the answers to the previous ones. Note that in some applications, such as in the acquisition of wideband radio frequency signals, aggressive adaptive sensing mechanisms may not be practical because they would require near instantaneous feedback. However, there do exist applications where adaptive sensing is practical and where the potential benefits of adaptivity are too tantalizing to ignore.

The formidable possibilities offered by adaptive sensing give rise to the following natural "folk theorem."

> **Folk Theorem**. *The estimation error one can get by using a clever adaptive sensing scheme is far better than what is achievable by a nonadaptive scheme.*

In other words, learning about the signal along the way and adapting the questions (the next sensing vectors) to what has been learned to date is bound to help. In stark contrast, the main result of this paper is this:

> **Surprise**. *The folk theorem is wrong in general. No matter how clever the adaptive sensing mechanism, no matter how intractable the estimation procedure, in general it is not possible to achieve a fundamentally better mean-squared error (MSE) of estimation than that offered by a naïve random projection followed by ℓ_1 minimization.*

The rest of this article is mostly devoted to making this claim precise. In doing so, we shall also show that adaptivity does not help in obtaining a fundamentally better estimate of the signal support, which is of independent interest.

1.1 Main result

To formalize matters, we assume that the error vector $\mathbf{z}$ in (1.1) has i.i.d. $\mathcal{N}(0, \sigma^2)$ entries. Then if $\mathbf{A}$ is a random projection with unit-norm rows as discussed above, [6] shows that the Dantzig selector estimate $\widehat{\mathbf{x}}^{\mathrm{DS}}$ (obtained by solving a simple linear program) achieves an MSE obeying

$$\frac{1}{n}\,\mathbb{E}\,\|\widehat{\mathbf{x}}^{\mathrm{DS}} - \mathbf{x}\|_2^2 \le C\,\frac{k}{m}\,\log(n)\,\sigma^2, \tag{1.2}$$

where C is some numerical constant. The bound holds *universally* over all k-sparse signals[1] provided that the number of measurements m is sufficiently large (on the order of at least $k \log(n/k)$). Moreover, one can show that this result is essentially optimal in the sense that *any* possible nonadaptive choice of $\mathbf{A}$ (with unit-norm rows) and *any* possible estimation procedure $\widehat{\mathbf{x}}$ will satisfy

$$\frac{1}{n}\,\mathbb{E}\,\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \geq C'\,\frac{k}{m}\,\log(n/k)\,\sigma^2, \tag{1.3}$$

where C' is a numerical constant [3]. The fundamental question is thus: *how much lower can the MSE be* when (i) we are allowed to sense the signal adaptively and (ii) we can use any estimation algorithm we like to recover $\mathbf{x}$.

The distinction between adaptive and nonadaptive sensing can be expressed in the following manner. Begin by rewriting the statistical model (1.1) as

$$y_i = \langle \mathbf{a}_i, \mathbf{x} \rangle + z_i, \quad i = 1, \ldots, m, \tag{1.4}$$

in which a power constraint imposes that each $\mathbf{a}_i$ is of norm at most 1, i.e., $\|\mathbf{a}_i\|_2 \leq 1$; then in a nonadaptive sensing scheme the vectors $\mathbf{a}_1, \ldots, \mathbf{a}_m$ are chosen in advance and do not depend on $\mathbf{x}$ or $\mathbf{z}$ whereas in an adaptive setting, the measurement vectors may be chosen depending on the history of the sensing process, i.e., $\mathbf{a}_i$ is a (possibly random) function of $(\mathbf{a}_1, y_1, \ldots, \mathbf{a}_{i-1}, y_{i-1})$.

If we follow the principle that "you cannot get something for nothing," one might argue that giving up the freedom to adaptively select the sensing vectors would result in a far worse MSE. Our main contribution is to show that this is not the case. We prove that there are classes of vectors $\mathbf{x}$ for which the MSE of adaptive sensing strategies is within a logarithmic factor of that of the nonadaptive schemes.

Theorem 1. *Suppose that $k < n/2$ and let m be arbitrary. Assume that $\mathbf{x}$ is sampled with i.i.d. coordinates such that $x_j = 0$ with probability $1 - k/n$ and $x_j = \mu$ with probability k/n. Then for $\mu = \frac{4}{3}\sqrt{\frac{n}{m}}$, any sensing strategy and any estimate $\widehat{\mathbf{x}}$ obey*

$$\frac{1}{n}\,\mathbb{E}\,\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \geq \frac{4}{27}\,\frac{k}{m}\,\sigma^2 > \frac{1}{7}\,\frac{k}{m}\,\sigma^2.$$

Note that when n and k are sufficiently large, then from the normal approximation to the binomial we have that for at least 95% of draws we will have an $\mathbf{x}$ with approximately $k \pm 2\sqrt{k}$ nonzeros. The following result establishes a lower bound for exactly k-sparse signals.

Theorem 2. *For any $k < n/2$ and any m, the following minimax lower bound holds:*

$$\inf_{\widehat{\mathbf{x}}} \sup_{\mathbf{x}\,:\,\|\mathbf{x}\|_0 = k} \frac{1}{n}\,\mathbb{E}\,\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \geq C_0\,\frac{k}{m}\,\sigma^2.$$

In the case of general n and k we can take $C_0 = 1/216$, although if n is a multiple of $2k$ then this can be improved to $C_0 = 1/108$. Further improvements should be possible as we have made no effort to optimize this constant. See Section 2 for details.

In short, Theorems 1 and 2 say that if one ignores a logarithmic factor, then *adaptive measurement schemes cannot (substantially) outperform nonadaptive strategies.* While seemingly counterintuitive, we find that precisely the same sparse vectors which determine the minimax rate in the nonadaptive setting are essentially so difficult to estimate that by the time we have identified the support, we will have already exhausted our measurement budget (i.e., we will have acquired all m measurements).

[1] A signal is said to be *k-sparse* if it has at most k nonzero components. We also occasionally use the notation $\|\mathbf{x}\|_0$ to denote the number of nonzero components of $\mathbf{x}$.

A brief note is in order regarding what we mean by a *substantial* improvement. After all, the lower bound in Theorem 2 does improve upon the nonadaptive bound in (1.3) by a factor of $\log(n/k)$. However, this is a relatively modest improvement compared to what one might hope to gain by exploiting adaptivity. Specifically, consider a simple adaptive procedure that uses $m/2$ measurements to identify the support of $\mathbf{x}$ and uses the remaining $m/2$ measurements to estimate the values of the nonzeros. If such a scheme identifies the correct support, then it is easy to show that this procedure will yield an estimate satisfying

$$\frac{1}{n}\,\mathbb{E}\,\|\widehat{\mathbf{x}}-\mathbf{x}\|_2^2 = \frac{2k}{n}\,\frac{k}{m}\,\sigma^2.$$

Thus, there seems to be room for reducing the error by a factor of k/n beyond the $\log(n/k)$ factor. Theorem 2, however, shows that this gain is not possible in general.

On the one hand, our main result states that one cannot universally improve on bounds achievable via nonadaptive sensing strategies. Indeed, we will see that there are natural classes of sparse signals for which, even after applying the most clever sensing scheme and the most subtle testing procedure, one would still not be sure about where the nonzeros lie. This remains true even after having used up the entirety of our measurement budget. On the other hand, our result does not say that adaptive sensing *never* helps. In fact, there are many instances in which it will. For example, when some or most of the nonzero entries in $\mathbf{x}$ are sufficiently large, they may be detected sufficiently early so that one can ultimately get a far better MSE than what would be obtained via a nonadaptive scheme, see Section 3 for simple experiments in this direction and Section 4 for further discussion.

1.2 Connections with testing problems

The arguments we develop to reach our conclusions are quite intuitive, simple, and yet they seem different from the classical Fano-type arguments for obtaining information-theoretic lower bounds (see Section 1.3 for a discussion of the latter methods). Our approach involves proving a lower bound for the Bayes risk under the prior from Theorem 1. To obtain such a lower bound, we make a detour through testing—multiple testing to be exact. Our argument proceeds through two main steps:

- *Support recovery in Hamming distance.* We consider the multiple testing problem of deciding which components of the signal are zero and which are not. We show that no matter which adaptive strategy and tests are used, the Hamming distance between the estimated and true supports is large. Put differently, the multiple testing problem is shown to be difficult. In passing, this establishes that adaptive schemes are not substantially better than nonadaptive schemes for support recovery.
- *Estimation with mean-squared loss.* Any estimator with a low MSE can be converted into an effective support estimator simply by selecting the largest coordinates or those above a certain threshold. Hence, a lower bound on the Hamming distance immediately gives a lower bound on the MSE.

The crux of our argument is thus to show that it is not possible to choose sensing vectors adaptively in such a way that the support of the signal may be estimated accurately.

1.3 Differential entropies and Fano-type arguments

Our approach is significantly different from classical methods for getting lower bounds in decision and information theory. Such methods typically rely on Fano's inequality [9], and are all intimately related to methods in statistical decision theory (see [29, 31]). Before continuing, we would like to point out that Fano-type arguments have been used successfully to obtain (often sharp) lower bounds for some adaptive methods. For example, the work [8] uses results from [29] to establish a bound on the minimax rate for

binary classification (see the references therein for additional literature on active learning). Other examples include the recent paper [26], which derives lower bounds for bandit problems, and [24] which develops an information theoretic approach suitable for stochastic optimization, a form of online learning, and gives bounds about the convergence rate at which iterative convex optimization schemes approach a solution.

Following the standard approaches in our setting leads to major obstacles that we would like to briefly describe. Our hope is that this will help the reader to better appreciate our easy itinerary. As usual, we start by choosing a prior for $\mathbf{x}$, which we take having zero mean. Coming from information theory, one would want to bound the mutual information between $\mathbf{x}$ (what we want to learn about) and $\mathbf{y}$ (the information we have), for any measurement scheme $\mathbf{a}_1, \dots, \mathbf{a}_m$. Assuming a deterministic measurement scheme, by the chain rule, we have

$$I(\mathbf{x}, \mathbf{y}) = h(\mathbf{y}) - h(\mathbf{y} \,|\, \mathbf{x}) = \sum_{i=1}^{m} h(y_i \,|y_{[i-1]}) - h(y_i \,|y_{[i-1]}, \mathbf{x}), \tag{1.5}$$

where $y_{[i]} := (y_1, \dots, y_i)$. Since the history up to time $i-1$ determines $\mathbf{a}_i$, the conditional distribution of y_i given $y_{[i-1]}$ and $\mathbf{x}$ is then normal with mean $\langle \mathbf{a}_i, \mathbf{x} \rangle$ and variance σ^2. Hence, $h(y_i \,|\, y_{[i-1]}, \mathbf{x}) = \frac{1}{2} \log(2\pi e \sigma^2)$. This is the easy term to handle — the challenging term is $h(y_i \,|\, y_{[i-1]})$ and it is not clear how one should go about finding a good upper bound. To see this, observe that

$$\mathrm{Var}(y_i \,|\, y_{[i-1]}) = \mathrm{Var}(\langle \mathbf{a}_i, \mathbf{x} \rangle \,|\, y_{[i-1]}) + \sigma^2.$$

A standard approach to bound $h(y_i \,|\, y_{[i-1]})$ is to write

$$h(y_i \,|y_{[i-1]}) \le \frac{1}{2} \, \mathbb{E} \log\big(2\pi e \mathrm{Var}(\langle \mathbf{a}_i, \mathbf{x} \rangle \,|\, y_{[i-1]}) + 2\pi e \sigma^2\big),$$

using the fact that the Gaussian distribution maximizes the entropy among distributions with a given variance. If we simplify the problem by applying Jensen's inequality, we obtain

$$I(\mathbf{x}, \mathbf{y}) \le \sum_{i=1}^{m} \frac{1}{2} \log\big(\mathbb{E}\langle \mathbf{a}_i, \mathbf{x} \rangle^2 / \sigma^2 + 1\big). \tag{1.6}$$

The RHS needs to be bounded uniformly over all choices of measurement schemes, which is a daunting task given that $\mathbf{a}_i$ is a function of $y_{[i-1]}$ which is in turn a function of $\mathbf{x}$. We note however that the RHS *can* be bounded in the nonadaptive setting, which is the approach taken in [3] to establish (1.3). See also [1, 25, 30] for other asymptotic results in this direction.

We have presented the problem in this form to help information theorists see the analogy with the problem of understanding the role of feedback in a Gaussian channel [9]. Specifically, we can view the inner products $\langle \mathbf{a}_i, \mathbf{x} \rangle$ as inputs to a Gaussian channel where we observe the output of the channel via feedback. It is well-known that feedback does not substantially increase the capacity of a Gaussian channel, so one might expect this argument to be relevant to our problem as well. Crucially, however, in the case of a Gaussian channel the user has full control over the channel input—whereas in the absence of a priori knowledge of $\mathbf{x}$, in our problem we are much more restricted in our control over the "channel input" $\langle \mathbf{a}_i, \mathbf{x} \rangle$.

1.4 Connections with other works

A number of papers have studied the advantages (or sometimes the lack thereof) offered by adaptive sensing in the setting where one has *noiseless* data, see for example [11, 16, 23] and references therein. Of course, it is well known that one can uniquely determine a k-sparse vector from $2k$ linear nonadaptive noise-free measurements and, therefore, there is not much to dwell on. The aforementioned works of course do not study such a trivial problem. Rather, the point of view is that the signal is not exactly sparse,

only approximately sparse, and the question is thus whether one can get a lower approximation error by employing an adaptive scheme. Whereas we study a statistical problem, this is a question in approximation theory. Consequently, the techniques and results of this line of research have no bearing on our problem.

There is much research suggesting intelligent adaptive sensing strategies in the presence of noise and we mention a few of these works. In a setting closely related to ours—that of detecting the locations of the nonzeros of a sparse signal from noisy point samples (so that $m > n$)—[14] shows that by adaptively allocating sensing resources one can significantly improve upon the best nonadaptive schemes [12]. Lower bounds for nonadaptive and adaptive methods in this context were recently established in [20], with the adaptive lower bounds established through the sequential probability ratio test (SPRT) [27]. Closer to home, [13, 15] consider CS schemes (with $m < n$) which perform sequential subset selection via the random projections typical of CS, but which focus in on promising areas of the signal. When the signal is (i) very sparse (ii) has sufficiently large entries and (iii) has constant dynamic range, the method in [13] is able to remove a logarithmic factor from the MSE achieved by the Dantzig selector with (nonadaptive) i.i.d. Gaussian measurements. In a different direction, [7, 18] suggest Bayesian approaches where the measurement vectors are sequentially chosen so as to maximize the conditional differential entropy of y_i given $y_{[i-1]}$. Finally, another approach in [17] suggests a bisection method based on repeated measurements for the detection of 1-sparse vectors, subsequently extended to k-sparse vectors via hashing. None of these works, however, establish a lower bound on the MSE of the recovered signal.

1.5 Content

We prove all of our results in Section 2, trying to give as much insight as possible as to why adaptive methods are not much more powerful than nonadaptive ones for detecting the support of a sparse signal. We will also attempt to describe the regime in which adaptivity might be helpful via simple numerical simulations in Section 3. These simulations show that adaptive algorithms are subject to a fundamental phase transition phenomenon. Finally, we comment on open problems and future research in Section 4.

2 Limits of Adaptive Sensing Strategies

This section establishes nonasymptotic lower bounds for the estimation of a sparse vector from adaptively selected noisy linear measurements. To begin with, we remind ourselves that we collect possibly adaptive measurements of the form (1.4) of an n-dimensional signal $\mathbf{x}$ where $\|\mathbf{a}_i\|_2 \le 1$; from now on, we assume for simplicity and without loss of generality that $\sigma = 1$.

In our analysis below, we denote the total-variation metric between any two probability distributions $\mathbb{P}$ and $\mathbb{Q}$ by $\|\mathbb{P} - \mathbb{Q}\|_{\mathrm{TV}}$, and their KL divergence by $K(\mathbb{P}, \mathbb{Q})$ [22]. Our arguments will make use of Pinsker's inequality, which relates these two quantities via

$$\|\mathbb{P} - \mathbb{Q}\|_{\mathrm{TV}} \le \sqrt{K(\mathbb{Q}, \mathbb{P})/2}. \tag{2.1}$$

We shall also use the convexity of the KL divergence, which states that for $\lambda_i \ge 0$ and $\sum_i \lambda_i = 1$, we have

$$K\Big(\sum_i \lambda_i \mathbb{P}_i, \sum_i \lambda_i \mathbb{Q}_i\Big) \le \sum_i \lambda_i K(\mathbb{P}_i, \mathbb{Q}_i) \tag{2.2}$$

in which $\{\mathbb{P}_i\}$ and $\{\mathbb{Q}_i\}$ are families of probability distributions.

Before proceeding, we argue that when we are given a prior $\pi(\mathbf{x})$, we can restrict ourselves to deterministic measurement schemes in the sense that $\mathbf{a}_1$ is a deterministic vector and, for $i \ge 2$, $\mathbf{a}_i$ is a deterministic function of $y_{[i-1]} = (y_1, \ldots, y_i)$. In the general case we have $\mathbf{a}_i = F_i(y_{[i-1]}, U_i)$, where F_i is a deterministic

function and U_i is random and independent of $y_{[i-1]}$ and z_i. With $\mathbf{U} = (U_1, \ldots, U_m)$, it follows from the law of iterated expectation

$$\mathbb{E}\,\|\widehat{\mathbf{x}} - \mathbf{x}\|^2 = \mathbb{E}\Big[\mathbb{E}[\|\widehat{\mathbf{x}} - \mathbf{x}\|^2 | \mathbf{U}]\Big]$$

(the expectation in the left-hand side is taken over $\mathbf{x}, \mathbf{y}$ and $\mathbf{U}$) that there exists a fixed realization $\mathbf{u} = (u_1, \ldots, u_m)$ obeying

$$\mathbb{E}[\|\widehat{\mathbf{x}} - \mathbf{x}\|^2 | \mathbf{U} = \mathbf{u}] \le \mathbb{E}\,\|\widehat{\mathbf{x}} - \mathbf{x}\|^2.$$

Hence, we can construct an estimator based on a deterministic measurement scheme which is as good as any based on a randomized measurement scheme. Note that in a deterministic scheme, letting $\mathbb{P}_{\mathbf{x}}$ be the distribution of $y_{[i-1]}$ when the target vector is $\mathbf{x}$ and using the fact that y_i is conditionally independent of $y_{[i-1]}$ given $\mathbf{a}_i$, we see that the likelihood factorizes as

$$\mathbb{P}_{\mathbf{x}}(y_{[m]}) = \prod_{i=1}^{m} \mathbb{P}_{\mathbf{x}}(y_i | \mathbf{a}_i), \tag{2.3}$$

which will be of use in our analysis below.

2.1 The Bernoulli prior

We begin by studying the model in Theorem 1 which makes our argument most transparent. The ideas to prove Theorem 2 are essentially the same and differ only in technical complications.

In this model, we suppose that $\mathbf{x} \in \mathbb{R}^n$ is sampled from a product prior: for each $j \in \{1, \ldots, n\}$,

$$x_j = \begin{cases} 0 & \text{w.p. } 1 - k/n, \\ \mu & \text{w.p. } k/n, \end{cases}$$

and the x_j's are independent. In this model, $\mathbf{x}$ has on average k nonzero entries, all with known positive amplitudes equal to μ. This model is easier to study than the related model in which one selects k coordinates uniformly at random and sets those to μ. The reason is that in this Bernoulli model, the independence between the coordinates of $\mathbf{x}$ brings welcomed simplifications, as we shall see.

Our goal here is to establish a lower bound on the MSE when $\mathbf{x}$ is drawn from this prior. We do this in two steps. First, we look at recovering the support of $\mathbf{x}$, which is done via a reduction to multiple testing. Second, we show that a lower bound on the error for support recovery implies a lower bound on the MSE, leading to Theorem 1.

2.1.1 Support recovery in Hamming distance

We would like to understand how well we can estimate the support $S = \{j : x_j \neq 0\}$ of $\mathbf{x}$ from the data (1.4), and shall measure performance by means of the expected Hamming distance. Here, the error of a procedure $\widehat{S}$ for estimating the support S is defined as

$$\mathbb{E}\,|\widehat{S} \Delta S| = \sum_{j=1}^{n} \mathbb{P}(\widehat{S}_j \neq S_j)$$

where Δ denotes the symmetric difference, $S_j = 1$ if $j \in S$ and equals zero otherwise, and similarly for $\widehat{S}_j$. As we can see, this reduces our problem to a sequence of n independent hypothesis tests. We will obtain a lower bound on the number of errors among these tests by exploiting the following lemma.

Lemma 1. *Consider the testing problem of deciding between $H_0 : \mathbf{x} \sim \mathbb{P}_0$ and $H_1 : \mathbf{x} \sim \mathbb{P}_1$, where H_0 and H_1 occur with prior probabilities π_0 and π_1 respectively. Under the 0-1 loss, The Bayes risk B obeys*

$$B \geq \min(\pi_0, \pi_1) \left(1 - \| \mathbb{P}_1 - \mathbb{P}_0 \|_{\mathrm{TV}}\right).$$

Proof. Assume without loss of generality that $\pi_1 \leq \pi_0$. The test with minimum risk is the Bayes test rejecting H_0 if and only if

$$\Lambda = \frac{\pi_1 \, \mathbb{P}_1(\mathbf{x})}{\pi_0 \, \mathbb{P}_0(\mathbf{x})} > 1;$$

that is, if the adjusted likelihood ratio exceeds one; see [19, Pbm. 3.10]. A simple calculation shows that the Bayes risk obeys

$$B = \pi_0 \, \mathbb{E}_0 \left(\min(1, \Lambda)\right),$$

where $\mathbb{E}_0$ denotes expectation under $\mathbb{P}_0$. Using the fact that $\mathbb{E}_0 \, \Lambda = \pi_1/\pi_0$ together with

$$\min(1, \Lambda) = \frac{1 + \Lambda}{2} + \frac{|\Lambda - 1|}{2},$$

we obtain

$$B = \frac{1}{2} - \frac{\pi_0}{2} \, \mathbb{E}_0 \, |\Lambda - 1|. \tag{2.4}$$

Finally,

$$\begin{aligned} \pi_0 \, \mathbb{E}_0 \, |\Lambda - 1| = \int |\pi_1 \mathrm{d} \, \mathbb{P}_1 - \pi_0 \mathrm{d} \, \mathbb{P}_0 \,| &\leq \pi_1 \int |\mathrm{d} \, \mathbb{P}_1 - \mathrm{d} \, \mathbb{P}_0 \,| + \pi_0 - \pi_1 \\ &= 2\pi_1 \| \, \mathbb{P}_1 - \mathbb{P}_0 \, \|_{\mathrm{TV}} + \pi_0 - \pi_1, \end{aligned}$$

which when combined with (2.4) establishes the lemma. □

Theorem 3. *Suppose that $\mathbf{x}$ is sampled according to the Bernoulli prior with $k \leq n/2$, then any estimate $\widehat{S}$ obeys*

$$\mathbb{E} \, |\widehat{S} \Delta S| \geq k \Big(1 - \frac{\mu}{2} \sqrt{\frac{m}{n}}\Big). \tag{2.5}$$

Hence, if the amplitude of the signal is below $\sqrt{n/m}$, we expect a large number of errors; indeed, if $\mu = \sqrt{n/m}$, then $\mathbb{E} \, |\widehat{S} \Delta S| \geq k/2$.

Proof. [2] Let $\pi_1 = k/n$ and $\pi_0 = 1 - \pi_1$. For any j, set $\mathbb{P}_{0,j} = \mathbb{P}(\cdot | x_j = 0)$ and $\mathbb{P}_{1,j} = \mathbb{P}(\cdot | x_j \neq 0)$. Let B_j denote the Bayes risk of the decision problem $H_{0,j} : x_j = 0$ versus $H_{1,j} : x_j = 1$. From Lemma 1 we have that

$$\mathbb{E} \, |\widehat{S} \Delta S| = \sum_{j=1}^{n} \mathbb{P}(\widehat{S}_j \neq S_j) \geq \sum_{j=1}^{n} B_j \geq \pi_1 \sum_{j=1}^{n} \Big(1 - \|\mathbb{P}_{1,j} - \mathbb{P}_{0,j}\|_{\mathrm{TV}}\Big).$$

Applying the Cauchy-Schwartz inequality, we obtain

$$\mathbb{E} \, |\widehat{S} \Delta S| \geq k \Big(1 - \frac{1}{\sqrt{n}} \sqrt{\sum_{j=1}^{n} \|\mathbb{P}_{1,j} - \mathbb{P}_{0,j}\|_{\mathrm{TV}}^2}\Big). \tag{2.6}$$

[2] The main ideas of our proof are similar to those in that of Assouad's Lemma, see [2, 29] for instance. Note, however, that our approach yields a sharper constant.

The theorem is a consequence of (2.6) combined with

$$\sum_{j=1}^{n} \|\mathbb{P}_{1,j} - \mathbb{P}_{0,j}\|_{\mathrm{TV}}^2 \le \frac{\mu^2}{4}\, m. \tag{2.7}$$

To establish (2.7), we apply Pinsker's inequality twice to obtain

$$\|\mathbb{P}_{1,j} - \mathbb{P}_{0,j}\|_{\mathrm{TV}}^2 \le \frac{\pi_0}{2} K(\mathbb{P}_{0,j}, \mathbb{P}_{1,j}) + \frac{\pi_1}{2} K(\mathbb{P}_{1,j}, \mathbb{P}_{0,j}) \tag{2.8}$$

so that it remains to find an upper bound on the KL divergence between $\mathbb{P}_{0,j}$ and $\mathbb{P}_{1,j}$. Write $\mathbb{P}_0 = \mathbb{P}_{0,j}$ for short and likewise for $\mathbb{P}_{1,j}$. Then

$$\mathbb{P}_0(y_{[m]}) = \sum_{\mathbf{x}'} \mathbb{P}(\mathbf{x}')\, \mathbb{P}(y_{[m]} | x_j = 0, \mathbf{x}') := \sum_{\mathbf{x}'} \mathbb{P}(\mathbf{x}') \mathbb{P}_{0,\mathbf{x}'},$$

where $\mathbf{x}' = (x_1, \dots, x_{j-1}, x_{j+1}, \dots, x_n)$ and $\mathbb{P}_{0,\mathbf{x}'}$ is the conditional probability distribution of $y_{[m]}$ given $\mathbf{x}'$ and $x_j = 0$; $\mathbb{P}_1(y_{[m]})$ is defined similarly. The convexity of the KL divergence (2.2) gives

$$K(\mathbb{P}_0, \mathbb{P}_1) \le \sum_{\mathbf{x}'} \mathbb{P}(\mathbf{x}') K(\mathbb{P}_{0,\mathbf{x}'}, \mathbb{P}_{1,\mathbf{x}'}). \tag{2.9}$$

We now calculate this divergence. In order to do this, observe that we have $y_i = \langle \mathbf{a}_i, \mathbf{x} \rangle + z_i = c_i + z_i$ under $\mathbb{P}_{0,\mathbf{x}'}$ while $y_i = a_{i,j}\mu + c_i + z_i$ under $\mathbb{P}_{1,\mathbf{x}'}$. This yields

$$\begin{aligned}
K(\mathbb{P}_{0,\mathbf{x}'}, \mathbb{P}_{1,\mathbf{x}'}) &= \mathbb{E}_{0,\mathbf{x}'} \log \frac{\mathbb{P}_{0,\mathbf{x}'}}{\mathbb{P}_{1,\mathbf{x}'}} \\
&= \sum_{i=1}^{m} \mathbb{E}_{0,\mathbf{x}'} \left(\frac{1}{2}(y_i - \mu a_{i,j} - c_i)^2 - \frac{1}{2}(y_i - c_i)^2 \right) \\
&= \sum_{i=1}^{m} \mathbb{E}_{0,\mathbf{x}'} \left(-z_i \mu a_{i,j} + (\mu a_{i,j})^2/2 \right) \\
&= \frac{\mu^2}{2} \sum_{i=1}^{m} \mathbb{E}_{0,\mathbf{x}'}(a_{i,j}^2).
\end{aligned}$$

The first equality holds by definition, the second follows from (2.3), the third from $y_i = c_i + z_i$ under $\mathbb{P}_{0,\mathbf{x}'}$ and the last holds since z_i is independent of $a_{i,j}$ and has zero mean. Using (2.9), we obtain

$$K(\mathbb{P}_0, \mathbb{P}_1) \le \frac{\mu^2}{2} \sum_{i=1}^{m} \mathbb{E}[a_{i,j}^2 | x_j = 0].$$

Similarly,

$$K(\mathbb{P}_1, \mathbb{P}_0) \le \frac{\mu^2}{2} \sum_{i=1}^{m} \mathbb{E}[a_{i,j}^2 | x_j = \mu]$$

and, therefore, (2.8) shows that

$$\|\mathbb{P}_{1,j} - \mathbb{P}_{0,j}\|_{\mathrm{TV}}^2 \le \frac{\mu^2}{4} \Big(\sum_{i=1}^{m} \pi_0\, \mathbb{E}[a_{i,j}^2 | x_j = 0] + \pi_1\, \mathbb{E}[a_{i,j}^2 | x_j = \mu] \Big) = \frac{\mu^2}{4} \sum_{i=1}^{m} \mathbb{E}[a_{i,j}^2].$$

For any particular pair (i, j) with $i > 1$, we can say very little about $\mathbb{E}[a_{i,j}^2]$ since it can depend on all the previous measurements in a potentially very complicated manner. However, by summing this inequality over j we can obtain (2.7) by using the only constraint we have imposed on the $\mathbf{a}_i$, namely, $\|\mathbf{a}_i\|_2 = 1$, so that $\sum_{ij} a_{ij}^2 = m$. This establishes the theorem. □

2.1.2 Estimation in mean-squared error

It is now straightforward to obtain a lower bound on the MSE from Theorem 3.

Proof of Theorem 1. Let S be the support of $\mathbf{x}$ and set $\widehat{S} := \{j : |\widehat{x}_j| \geq \mu/2\}$. We have

$$\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 = \sum_{j \in S} (\widehat{x}_j - x_j)^2 + \sum_{j \notin S} \widehat{x}_j^2 \geq \frac{\mu^2}{4} |S \setminus \widehat{S}| + \frac{\mu^2}{4} |\widehat{S} \setminus S| = \frac{\mu^2}{4} |\widehat{S} \Delta S|$$

and, therefore,

$$\mathbb{E}\, \|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \geq \frac{\mu^2}{4}\, \mathbb{E}\, |\widehat{S} \Delta S| \geq \frac{\mu^2}{4} k \Big(1 - \frac{\mu}{2} \sqrt{\frac{m}{n}}\Big),$$

where the last inequality is from Theorem 3. We then plug in $\mu = \frac{4}{3}\sqrt{\frac{n}{m}}$ and simplify to conclude. □

2.2 The uniform-within-bins prior

We hope to have made clear how our argument above is considerably simpler than a Fano-type argument, at least for the Bernoulli prior. We now wish to prove a similar result for exactly k-sparse signals in order to establish Theorem 2. The natural prior in this context is the uniform prior over exactly k-sparse vectors with nonzero entries all equal. The main obstacle in following the previous approach is the lack of independence between the coordinates of the sampled signal $\mathbf{x}$.

To address this difficulty, we consider the following *uniform-within-bins* prior. Specifically, we let $N = \lfloor n/(2k) \rfloor$ and let $J_1, \ldots, J_{2k} \subset \{1, \ldots, n\}$ be disjoint subsets each of size N. For $q = 1, 2, \ldots, k$, the uniform-within-bins prior picks one of the two bins J_{2q-1} or J_{2q} with equal probability and places a single nonzero coefficient (equal to $\mu > 0$) at a random position in the selected set. We now prove a key result for this prior that will allow us to establish Theorem 2.

Theorem 4. *Suppose that* $\mathbf{x}$ *is sampled according to the uniform-within-bins prior with* $k \leq n/2$. *Set* $N = \lfloor n/(2k) \rfloor$, *then any estimate* $\widehat{S}$ *obeys,*

$$\mathbb{E}\, |\widehat{S} \Delta S| \geq k \Big(1 - \mu \sqrt{\frac{2m}{kN}}\Big). \tag{2.10}$$

Proof. Let $\widehat{S}$ be any estimate of S and let S_{J_s} denote the restriction of S to J_s. We begin by observing that we can write

$$|\widehat{S} \Delta S| \geq \sum_{q=1}^{k} |\widehat{S}_{J_{2q-1}} \Delta S_{J_{2q-1}}| + |\widehat{S}_{J_{2q}} \Delta S_{J_{2q}}|.$$

Put $T_s = 1$ if $S_j \neq 0$ for some $j \in J_s$ and $T_s = 0$ otherwise, and define $\widehat{T}_s$ in a similar way by using $\widehat{S}$. Note that $T_{2q-1} = 1$ and $T_{2q} = 0$ or $T_{2q-1} = 0$ and $T_{2q} = 1$ with equal probability. Thus we trivially have that if $\widehat{T}_{2q-1} = \widehat{T}_{2q}$ then

$$|\widehat{S}_{J_{2q-1}} \Delta S_{J_{2q-1}}| + |\widehat{S}_{J_{2q}} \Delta S_{J_{2q}}| \geq 1. \tag{2.11}$$

Hence, suppose that $\widehat{T}_{2q-1} \neq \widehat{T}_{2q}$. In this case we can view $\widehat{T}_{2q-1}$ as the output of a test to determine whether the nonzero is located in J_{2q-1} or J_{2q}. If the test fails, then we have that both $|\widehat{S}_{J_{2q-1}} \Delta S_{J_{2q-1}}| \geq 1$ and $|\widehat{S}_{J_{2q}} \Delta S_{J_{2q}}| \geq 1$. Combining this with (2.11) we have that,

$$\mathbb{E}\, \Big[|\widehat{S}_{J_{2q-1}} \Delta S_{J_{2q-1}}| + |\widehat{S}_{J_{2q}} \Delta S_{J_{2q}}| \Big] \geq \min(1, 2\mathbb{P}(\widehat{T}_{2q-1} \text{ fails})) = 2\mathbb{P}(\widehat{T}_{2q-1} \text{ fails}) \tag{2.12}$$

since we can restrict attention to tests with a 50% chance of failure or less. So once again, we have reduced matters to a series of hypothesis tests. If we let $\mathbb{P}_s$ denote the distribution of $y_{[m]}$ given that $T_s = 1$, then using Lemma 1 and (2.12) we obtain

$$\mathbb{E}\,|\widehat{S}\Delta S| \geq \sum_{q=1}^{k}\Big(1 - \|\mathbb{P}_{2q-1} - \mathbb{P}_{2q}\|_{\mathrm{TV}}\Big).$$

Applying the Cauchy-Schwartz inequality as before, we obtain

$$\mathbb{E}\,|\widehat{S}\Delta S| \geq k - \sqrt{k}\sqrt{\sum_{q=1}^{k}\|\mathbb{P}_{2q-1} - \mathbb{P}_{2q}\|_{\mathrm{TV}}^2}.$$

Lemma 2 develops an upper bound of the form

$$\|\mathbb{P}_{2q-1} - \mathbb{P}_{2q}\|_{\mathrm{TV}}^2 \leq \frac{2\mu^2}{N}\sum_{i=1}^{m}\sum_{j\in J_{2q-1}\cup J_{2q}}\mathbb{E}\,a_{i,j}^2. \tag{2.13}$$

In turn, this yields

$$\mathbb{E}\,|\widehat{S}\Delta S| \geq k - \mu\sqrt{\frac{2k}{N}}\sqrt{\sum_{q=1}^{k}\sum_{i=1}^{m}\sum_{j\in J_{2q-1}\cup J_{2q}}\mathbb{E}\,a_{i,j}^2} \geq k - \mu\sqrt{\frac{2km}{N}}.$$

Rearranging terms establishes the theorem. □

Using the same argument as in the proof of Theorem 1, we have that

$$\mathbb{E}\,\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \geq \frac{\mu^2}{4}\,\mathbb{E}\,|\widehat{S}\Delta S| \geq \frac{\mu^2}{4}k\Big(1 - \mu\sqrt{\frac{2m}{kN}}\Big).$$

Optimizing this bound with respect to μ results in selecting $\mu^2 = \frac{2}{9}\frac{kN}{m}$. Plugging this in and dividing by n we obtain

$$\frac{1}{n}\,\mathbb{E}\,\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \geq \frac{1}{108}\,\frac{k}{m}\,\frac{2k}{n}N.$$

If n is a multiple of $2k$, then $(2k/n)N = 1$. For general n, note that we can write $2k/n = 1/(N+r)$ with $r \in [0,1)$, and hence $(2k/n)N = N/(N+r) > N/(N+1) \geq 1/2$, thus we obtain

$$\frac{1}{n}\,\mathbb{E}\,\|\widehat{\mathbf{x}} - \mathbf{x}\|_2^2 \geq \frac{1}{216}\,\frac{k}{m}.$$

This establishes Theorem 2 since the minimax risk over a class of alternatives is always lower-bounded by the average risk over any subset of that same class.

It remains to justify (2.13) which is the purpose of the Lemma below.

Lemma 2. *Let J_1 and J_2 be two disjoint subsets of $\{1,\dots,n\}$ of size N each. Let π' denote an arbitrary prior on vectors supported on $\{1,\dots,n\}\setminus(J_1\cup J_2)$, and for $s = 1,2$, let π'_s denote the uniform prior on vectors supported on $J_1\cup J_2$ with only one nonzero coordinate in J_s equal to $\mu > 0$. Finally, set $\pi_s = \pi'_s\otimes\pi'$ and let $\mathbb{P}_s$ denote the distribution of $y_{[m]}$ under π_s. Then*

$$\|\mathbb{P}_2 - \mathbb{P}_1\|_{\mathrm{TV}}^2 \leq \frac{2\mu^2}{N}\sum_{i=1}^{m}\mathbb{E}\sum_{j\in J_1\cup J_2}a_{i,j}^2,$$

where $\mathbb{E}$ denotes the expectation both with respect to the noise in (1.1) and with respect to $\pi = \frac{1}{2}(\pi_1+\pi_2)$.

Proof. Let π_0' be the prior on the null vector supported on $J_1 \cup J_2$, and put $\pi_0 = \pi_0' \otimes \pi'$. Again, $\mathbb{P}_0$ denotes the distribution of $y_{[m]}$ under π_0. The triangle inequality gives

$$\|\mathbb{P}_2 - \mathbb{P}_1\|_{\mathrm{TV}} \le \|\mathbb{P}_0 - \mathbb{P}_1\|_{\mathrm{TV}} + \|\mathbb{P}_0 - \mathbb{P}_2\|_{\mathrm{TV}}.$$

Using Pinsker's inequality (2.1), we focus on bounding the KL divergence $K(\mathbb{P}_0, \mathbb{P}_s)$ for $s = 1, 2$. We study $s = 1$ as the other case is exactly the same. Let $\mathbf{x}'$ denote the restriction of $\mathbf{x}$ to $\{1, \dots, n\} \setminus (J_1 \cup J_2)$. By the law of total probability,

$$\mathbb{P}_1 = \sum_{j \in J_1, \mathbf{x}'} \frac{1}{N} \mathbb{P}(\mathbf{x}') \; \mathbb{P}_1(\cdot | \mathbf{x}', x_j \neq 0).$$

and similarly,

$$\mathbb{P}_0 = \sum_{\mathbf{x}'} \mathbb{P}(\mathbf{x}') \, \mathbb{P}_0(\cdot | \mathbf{x}').$$

The convexity of the KL divergence gives

$$K(\mathbb{P}_0, \mathbb{P}_1) \le \frac{1}{N} \, \mathbb{E} \sum_{j \in J_1} K(\mathbb{P}_0(\cdot | \mathbf{x}'), \mathbb{P}_1(\cdot | \mathbf{x}', x_j \neq 0)).$$

We now repeat the same calculations as in the proof of Theorem 3 to obtain

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

where expectation $\mathbb{E}_0$ is calculated under π_0. The problem is that we do not want an expectation under π_0 but under the mixture $\frac{1}{2}(\pi_1 + \pi_2)$. We thus need to change the measure and this is the delicate step.

Observe that to finish the proof of the theorem, we only need to establish that

$$\mathbb{E}_0 \sum_{j \in J_1} a_{i,j}^2 \le 2 \, \mathbb{E} \sum_{j \in J_1 \cup J_2} a_{i,j}^2. \tag{2.14}$$

We do this by constructing an auxiliary sensing strategy $\{\mathbf{w}_i\}_{1 \le i \le m}$, obeying $w_{i,j} = 0$ for all $j \in J_2$ and

$$\sum_{j \in J_1} w_{i,j}^2 = \sum_{j \in J_1 \cup J_2} a_{i,j}^2.$$

In words, whatever "sensing energy" was spent on J_2 is moved to J_1. It is obvious that such a sensing scheme can be constructed and our claim does not depend upon the details of a particular construction. The key identity is

$$\mathbb{E}_0 \sum_{j \in J_1} w_{i,j}^2 = \mathbb{E}_2 \sum_{j \in J_1} w_{i,j}^2.$$

This holds because under the sensing scheme $\{\mathbf{w}_i\}$, the observations generated by π_0 are the same as those generated by π_2. The reason why this is true is that under $\mathbb{P}_2$, $\mathbf{x}$ vanishes on J_1. It follows from this inequality, that

$$\mathbb{E}_0 \sum_{j \in J_1} a_{i,j}^2 \le \mathbb{E}_0 \sum_{j \in J_1} w_{i,j}^2 = \mathbb{E}_2 \sum_{j \in J_1} w_{i,j}^2 = \mathbb{E}_2 \sum_{j \in J_1 \cup J_2} a_{i,j}^2.$$

Inequality (2.14) follows from the simple fact that under π, $\mathbf{x}$ is supported on J_2 with probability $\frac{1}{2}$. □

We close by noting that Lemma 2 implies a lower bound on the risk of deciding whether a vector with a single nonzero entry is supported on the first half or second half of the index set. Proving such a result by directly looking at the likelihood ratio—arguably the standard route—seems quite delicate to say the least. The reason is that we are testing a mixture (supported on the first half) versus another one (supported on the first half). Lemma 2 completely avoids doing this.

Algorithm 1 Adaptive algorithm from [7]

input: $m \times n$ random matrix $\mathbf{B}$ with i.i.d. Rademacher (± 1 with equal probability) entries.
initialize: $\mathbf{p} = \frac{1}{n}(1, \ldots, 1)^T$.
for $i = 1$ to $i = m$ **do**
  Compute $\mathbf{a}_i = (b_{i,1}\sqrt{p_1}, \ldots, b_{i,n}\sqrt{p_n})^T$.
  Observe $y_i = \langle \mathbf{a}_i, \mathbf{x} \rangle + z_i$.
  Update posterior distribution $\mathbf{p}$ of $\mathbf{x}$ given $(\mathbf{a}_1, y_1), \ldots, (\mathbf{a}_i, y_i)$ using the rule in [7].
end for
output: Estimate for support($\mathbf{x}$) is the index where $\mathbf{p}$ attains its maximum value.

Algorithm 2 Recursive bisection algorithm of [17]

input: $m_1, \ldots, m_{s_{\max}}$.
initialize: $J_1^{(1)} = \{1, \ldots, \frac{n}{2}\}$, $J_2^{(1)} = \{\frac{n}{2} + 1, \ldots, n\}$.
for $s = 1$ to $s = s_{\max}$ **do**
  Construct the $m_s \times n$ matrix $\mathbf{A}^{(s)}$ with rows $|J_1^{(s)}|^{-\frac{1}{2}} \mathbf{1}_{J_1^{(s)}} - |J_2^{(s)}|^{-\frac{1}{2}} \mathbf{1}_{J_2^{(s)}}$.
  Observe $\mathbf{y}^{(s)} = \mathbf{A}^{(s)}\mathbf{x} + \mathbf{z}^{(s)}$.
  Compute $w^{(s)} = \sum_{i=1}^{m_s} y_i^{(s)}$.
  Subdivide: Update $J_1^{(s+1)}$ and $J_2^{(s+1)}$ by partitioning $J_1^{(s)}$ if $w^{(s)} \geq 0$ or $J_2^{(s)}$ if $w^{(s)} < 0$.
end for
output: Estimate for support($\mathbf{x}$) is $J_1^{(s_{\max})}$ if $w^{(s_{\max})} \geq 0$, $J_2^{(s_{\max})}$ if $w^{(s_{\max})} < 0$.

3 Numerical Experiments

In order to briefly illustrate the implications of the lower bounds in Section 2 and the potential limitations and benefits of adaptivity in general, we include a few simple numerical experiments. To simplify our discussion, we limit ourselves to existing adaptive procedures that aim at consistent support recovery: the adaptive procedure from [7] and the recursive bisection algorithm of [17].

We emphasize that in the case of a generic k-sparse signal, there are many possibilities for adaptively estimating the support of the signal. For example, the approach in [13] iteratively rules out indices and could, in principle, proceed until only k candidate indices remain. In contrast, the approaches in [7] and [17] are built upon algorithms for estimating the support of 1-sparse signals. An algorithm for a 1-sparse signal could then be run k times to estimate a k-sparse signal as in [7], or used in conjunction with a hashing scheme as in [17]. Since our goal is not to provide a thorough evaluation of the merits of all the different possibilities, but merely to illustrate the general limits of adaptivity, we simplify our discussion and focus exclusively on the simple case of one-sparse signals, i.e., where $k = 1$.

Specifically, in our experiments we will consider the uniform prior on the set of vectors with a single nonzero entry equal to $\mu > 0$ as in Section 2. Since we are focusing only on the case of $k = 1$, the algorithms in [7] and [17] are extremely simple and are shown in Algorithm 1 and Algorithm 2 respectively. Note that in Algorithm 1 the step of updating the posterior distribution $\mathbf{p}$ consists of an iterative update rule given in [7] and does not require any a priori knowledge of the signal $\mathbf{x}$ or μ. In Algorithm 2, we simplify the recursive bisection algorithm of [17] using the knowledge that $\mu > 0$, which allows us to eliminate the second stage of the algorithm aimed at detecting negative coefficients. Note that this algorithm proceeds through $s_{\max} = \log_2 n$ stages and we must allocate a certain number of measurements to each stage. In our experiments we set $m_s = \lceil \beta 2^{-s} \rceil$, where β is selected to ensure that $\sum_{s=1}^{\log_2 n} m_s \leq m$.

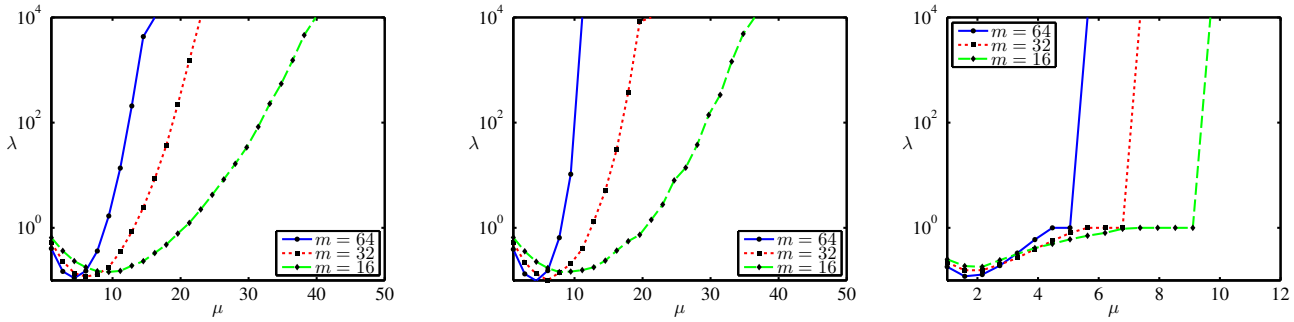

Figure 1: *Behavior of the posterior distribution as a function of μ for several values of m. (a) shows the results for nonadaptive measurements. (b) shows the results for Algorithm 1. (c) shows the results for Algorithm 2. We see that Algorithm 2 is able to detect somewhat weaker signals than Algorithm 1. However, for both cases we observe that once μ exceeds a certain threshold proportional to $\sqrt{n/m}$, the ratio λ of p_{j^*} to the second largest posterior probability grows exponentially fast, but that this does not differ substantially from the behavior observed in (a) when using nonadaptive measurements.*

3.1 Evolution of the posterior

We begin by showing the results of a simple simulation that illustrates the behavior of the posterior distribution of $\mathbf{x}$ as a function of μ for both adaptive schemes. Specifically, we assume that m is fixed and collect m measurements using each approach. Given the measurements $\mathbf{y}$, we then compute the posterior distribution $\mathbf{p}$ using the true prior used to generate the signal, which can be computed using the fact that

$$p_j \propto \exp\left(-\frac{1}{2\sigma^2}\|\mathbf{y} - \mu\mathbf{A}\mathbf{e}_j\|_2^2\right), \tag{3.1}$$

where σ^2 is the noise variance and $\mathbf{e}_j$ denotes the jth element of the standard basis. What we expect is that once μ exceeds a certain threshold (which depends on m), the posterior will become highly concentrated on the true support of $\mathbf{x}$. To quantify this, we consider the case where j^* denotes the true location of the nonzero element of $\mathbf{x}$ and define

$$\lambda = \frac{p_{j^*}}{\max_{j \neq j^*} p_j}.$$

Note that when $\lambda \leq 1$, we cannot reliably detect the nonzero, but when $\lambda \gg 1$ we can.

In Figure 1 we show the results for a few representative values of m (a) when using nonadaptive measurements, i.e., a (normalized) i.i.d. Rademacher random matrix $\mathbf{A}$, compared to the results of (b) Algorithm 1, and (c) Algorithm 2. For each value of m and for each value of μ, we acquire m measurements using each approach and compute the posterior $\mathbf{p}$ according to (3.1). We then compute the value of λ. We repeat this for 10,000 iterations and plot the median value of λ for each value of μ for all three approaches. In our experiments we set $n = 512$ and $\sigma^2 = 1$. We truncate the vertical axis at 10^4 to ensure that all curves are comparable. We observe that in each case, once μ exceeds a certain threshold proportional to $\sqrt{n/m}$, the ratio λ of p_{j^*} to the second largest posterior probability grows exponentially fast. As expected, this occurs for both the nonadaptive and adaptive strategies, with no substantial difference in terms of how large μ must be before support recovery is assured (although Algorithm 2 seems to improve upon the nonadaptive strategy by a small constant).

3.2 MSE performance

We have just observed that for a given number of measurements m, there is a critical value of μ below which we cannot reliably detect the support. In this section we examine the impact of this phenomenon

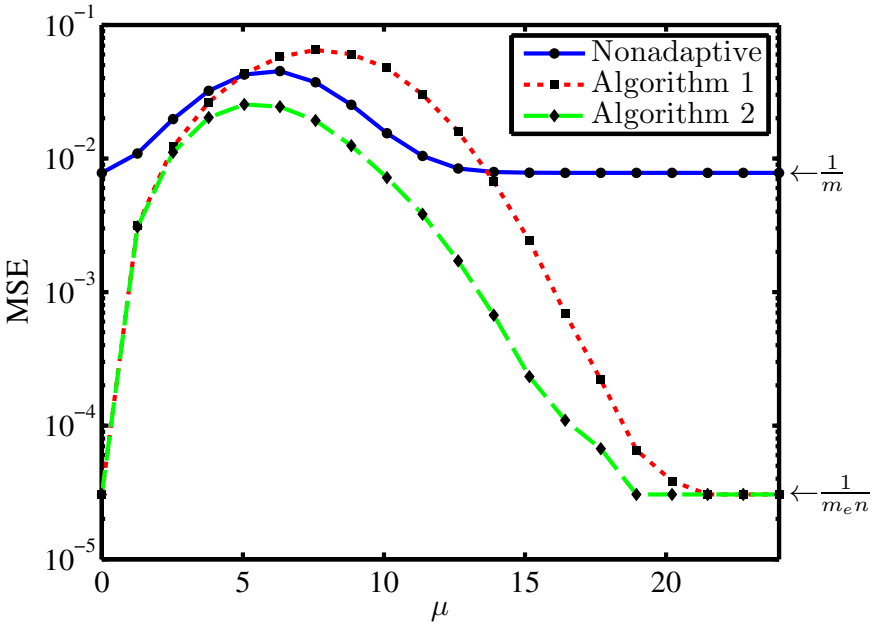

Figure 2: *The performance of Algorithm 1 and Algorithm 2 in the context of a two-stage procedure that first uses $m_d = \frac{m}{2}$ adaptive measurements to detect the location of the nonzero and then uses $m_e = \frac{m}{2}$ measurements to directly estimate the value of the identified coefficient. We show the resulting MSE as a function of the amplitude μ of the nonzero entry, and compare this to a nonadaptive procedure which uses a (normalized) i.i.d. Rademacher matrix followed by OMP. In the worst case, the MSE of the adaptive algorithms is comparable to the MSE obtained by the nonadaptive algorithm and exceeds the lower bound in Theorem 2 by only a small constant factor. When μ begins to exceed this critical threshold, the MSE of the adaptive algorithms rapidly decays below that of the nonadaptive algorithm and approaches $\frac{1}{m_e n}$, which is the MSE one would obtain given m_e measurements and a priori knowledge of the support.*

on the resulting MSE of a two-stage procedure that first uses $m_d = pm$ adaptive measurements to detect the location of the nonzero with either Algorithm 1 or Algorithm 2 and then reserves $m_e = (1-p)m$ measurements to directly estimate the value of the identified coefficient. It is not hard to show that if we correctly identify the location of the nonzero, then this will result in an MSE of $(m_e n)^{-1} = ((1-p)mn)^{-1}$. As a point of comparison, if an oracle provided us with the location of the nonzero a priori, we could devote all m measurements to estimating its value, with the best possible MSE being $\frac{1}{mn}$. Thus, if we can correctly detect the nonzero, this procedure will perform within a constant factor of the oracle.

We illustrate the performance of Algorithm 1 and Algorithm 2 in terms of the resulting MSE as a function of the amplitude μ of the nonzero in Figure 2. In this experiment we set $n = 512$ and $m = 128$ with $p = \frac{1}{2}$ so that $m_d = 64$ and $m_e = 64$. We then compute the average MSE over 100,000 iterations for each value of μ and for both algorithms. We compare this to a nonadaptive procedure which uses a (normalized) i.i.d. Rademacher matrix followed by orthogonal matching pursuit (OMP). Note that in the worst case the MSE of the adaptive algorithms is comparable to the MSE obtained by the nonadaptive algorithm and exceeds the lower bound in Theorem 2 by a only a small constant factor. However, when μ begins to exceed a critical threshold, the MSE rapidly decays and approaches the optimal value of $\frac{1}{m_e n}$. Note that when μ is large we can take $m_e \to m$ and hence can actually get arbitrarily close to $\frac{1}{mn}$ in the asymptotic regime.

4 Discussion

The contribution of this paper is to show that if one has the freedom to choose any adaptive sensing strategy and any estimation procedure no matter how complicated or computationally intractable, we would not be able to universally improve over a simple nonadaptive strategy that simply projects the signal onto a lower dimensional space and perform recovery via ℓ_1 minimization. This "negative" result should not conceal the fact that adaptivity may help tremendously if the SNR is sufficiently large, as illustrated in Section 3. Hence, we regard the design and analysis of effective adaptive schemes as a subject of important future research. At the methodological level, it seems important to develop adaptive strategies and algorithms for support estimation that are as accurate and as robust as possible. Further, a transition towards practical applications would need to involve engineering hardware that can effectively implement this sort of feedback, an issue which poses all kinds of very concrete challenges. Finally, at the theoretical level, it would be of interest to analyze the phase transition phenomenon we expect to occur in simple Bayesian signal models. For instance, a central question would be how many measurements are required to transition from a nearly flat posterior to one mostly concentrated on the true support.

We close by noting that after the submission of this paper, a variant of of Algorithm 2 was shown to recover the correct support of a 1-sparse vector with high probability provided that the amplitude μ of the nonzero entry satisfies $\mu \geq C\sqrt{n/m}$ [10, 21]. This implies that for the case of $k = 1$, the lower bound in Theorem 2 is tight up to constant factors. Thus, adaptive methods do have the potential to remove the $\log(n/k)$ factor that is required in the nonadaptive setting.

Acknowledgements

The authors would like to thank the reviewers as well as Rui Castro, Jarvis Haupt, and Alexander Tsybakov for their insightful feedback. The authors would also like to thank Xiaodong Li for suggesting an improvement in the proof of Theorem 3. E. A-C. is partially by ONR grant N00014-09-1-0258. E. C. is partially supported by NSF via grant CCF-0963835 and the 2006 Waterman Award, by AFOSR under grant FA9550-09-1-0643 and by ONR under grant N00014-09-1-0258. M. D. is supported by NSF grant DMS-1004718.

References

[1] S. Aeron, V. Saligrama, and M. Zhao. Information theoretic bounds for compressed sensing. *IEEE Trans. Inform. Theory*, 56(10):5111–5130, 2010.

[2] P. Assouad. Deux remarques sur l'estimation. *C. R. Acad. Sci. Paris Sér. I Math.*, 296(23):1021–1024, 1983.

[3] E. Candès and M. Davenport. How well can we estimate a sparse vector? *Arxiv preprint arXiv:1104.5246*, 2011.

[4] E. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, 2006.

[5] E. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5425, 2006.

[6] E. Candès and T. Tao. The Dantzig Selector: Statistical estimation when p is much larger than n. *Ann. Stat.*, 35(6):2313–2351, 2007.

[7] 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.

[8] R. Castro and R. Nowak. Minimax bounds for active learning. *IEEE Trans. Inform. Theory*, 54(5):2339–2353, 2008.

[9] T. Cover and J. Thomas. *Elements of information theory*. Wiley-Interscience, Hoboken, NJ, 2006.

[10] M. A. Davenport and E. Arias-Castro. Compressive binary search. In *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, Cambridge, MA, Jul. 2012.

[11] D. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.

[12] D. Donoho and J. Jin. Higher criticism for detecting sparse heterogeneous mixtures. *Ann. Stat.*, 32(3):962–994, 2004.

[13] 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.

[14] J. Haupt, R. Castro, and R. Nowak. Distilled sensing: Selective sampling for sparse signal recovery. In *Proc. Int. Conf. Art. Intell. Stat. (AISTATS)*, Clearwater Beach, FL, Apr. 2009.

[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] P. Indyk, E. Price, and D. Woodruff. On the power of adaptivity in sparse recovery. In *Proc. IEEE Symp. Found. Comp. Science (FOCS)*, Palm Springs, CA, Oct. 2011.

[17] 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.

[18] S. Ji, Y. Xue, and L. Carin. Bayesian compressive sensing. *IEEE Trans. Signal Processing*, 56(6):2346–2356, 2008.

[19] E. Lehmann and J. Romano. *Testing statistical hypotheses*. Springer Texts in Statistics. Springer, New York, 2005.

[20] M. Malloy and R. Nowak. On the limits of sequential testing in high dimensions. In *Proc. Asilomar Conf. Signals, Systems, and Computers*, Nov., 2011. Pacific Grove, CA.

[21] M. Malloy and R. Nowak. Near-optimal compressive binary search. *Arxiv preprint arXiv:1203.1804*, 2012.

[22] P. Massart. *Concentration inequalities and model selection*, volume 1896 of *Lecture Notes in Mathematics*. Springer, Berlin, 2007.

[23] E. Novak. On the power of adaptation. *J. Complexity*, 12(3):199–237, 1996.

[24] M. Raginsky and A. Rakhlin. Information complexity of black-box convex optimization: A new look via feedback information theory. In *Proc. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Oct. 2009.

[25] G. Raskutti, M. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q-balls. *Arxiv preprint arXiv:0910.2042*, 2009.

[26] P. Rigollet and A. Zeevi. Nonparametric bandits with covariates. *Arxiv preprint arXiv:1003.1630*, 2010.

[27] D. Siegmund. *Sequential analysis*. Springer Series in Statistics. Springer-Verlag, New York, 1985. Tests and confidence intervals.

[28] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288, 1996.

[29] A. Tsybakov. *Introduction to nonparametric estimation*. Springer Series in Statistics. Springer, New York, 2009.

[30] N. Verzelen. Minimax risks for sparse regressions: Ultra-high-dimensional phenomenons. *Electron. J. Statist.*, 6:38–90, 2012.

[31] B. Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, New York, 1997.