

On the Limits of Sequential Testing in High Dimensions

Matt Malloy

Electrical and Computer Engineering
University of Wisconsin-Madison
Email: mmalloy@wisc.edu

Robert Nowak

Electrical and Computer Engineering
University of Wisconsin-Madison
Email: nowak@ece.wisc.edu

Abstract—This paper presents results pertaining to sequential methods for support recovery of sparse signals in noise. Specifically, we show that *any* sequential measurement procedure fails provided the average number of measurements per dimension grows slower than $D(f_0||f_1)^{-1} \log s$ where s is the level of sparsity, and $D(f_0||f_1)$ the Kullback-Leibler divergence between the underlying distributions. For comparison, we show any *non-sequential* procedure fails provided the number of measurements grows at a rate less than $D(f_0||f_1)^{-1} \log n$, where n is the total dimension of the problem. Lastly, we show that a simple procedure termed *sequential thresholding* guarantees exact support recovery provided the average number of measurements per dimension grows faster than $D(f_0||f_1)^{-1}(\log s + \log \log n)$, a mere additive factor more than the lower bound.

I. INTRODUCTION

High-dimensional signal support recovery is a fundamental problem arising in many aspects of science and engineering. The goal of the basic problem is to determine, based on noisy observations, a sparse set of elements that somehow differ from the others.

In this paper we study the following problem. Consider a support set $\mathcal{S} \subset \{1, \dots, n\}$ and a random variable $y_{i,j}$ such that

$$y_{i,j} \sim \begin{cases} f_0(\cdot) & i \notin \mathcal{S} \\ f_1(\cdot) & i \in \mathcal{S} \end{cases} \quad (1)$$

where $f_0(\cdot)$ and $f_1(\cdot)$ are probability measures on \mathcal{Y} , and j indexes multiple independent measurements of any component $i \in \{1, \dots, n\}$. The dimension of the problem, n , is large – perhaps thousands or millions or more – but the support set \mathcal{S} is sparse in the sense that the number of elements following f_1 is much less than the dimension, i.e., $|\mathcal{S}| = s \ll n$. The goal of the sparse recovery problem is to identify the set \mathcal{S} .

In a *non-sequential* setting $m \geq 1$ independent observations of each component are made ($y_{i,1}, \dots, y_{i,m}$ are observed for each i) and the fundamental limits of reliable recovery are readily characterized in terms of Kullback-Leibler divergence and dimension.

Sequential approaches to the high dimensional support recovery problem have been given much attention recently (see [1], [2], [3], [4], etc). In the *sequential* setting, the decision to observe $y_{i,j}$ is based on prior observations, i.e. $y_{i,1}, \dots, y_{i,j-1}$. Herein lies the advantage of sequential methods: if prior measurements indicate a particular component belongs (or

doesn't belong) to \mathcal{S} with sufficient certainty, measurement of that component can cease, and resources can be diverted to a more uncertain element.

The results presented in this paper are in terms of asymptotic rate at which the average number of measurements per dimension, denoted m , must increase with n to ensure exact recovery of \mathcal{S} for any fixed distributions f_0 and f_1 . The main contributions are 1) to present a necessary condition for success of any sequential procedure in the sparse setting, 2) show success of a simple sequential procedure first presented in [2] is guaranteed provided the average number of measurements per dimension is within a small additive factor of the necessary condition for any sequential procedure, and compare this procedure to the known optimal sequential probability ratio test, and 3) lastly, compare these results to the performance limits of any non-sequential procedure. Table I summarizes these results.

TABLE I
AVERAGE NUMBER OF MEASUREMENTS PER DIMENSION FOR EXACT RECOVERY

<i>Non-sequential</i>	$m \geq \frac{\log n}{D(f_1 f_0)}$	necessary
<i>Sequential</i>	$m \geq \frac{\log s}{D(f_0 f_1)}$	necessary
<i>Sequential Thresholding</i>	$m > \frac{\log s}{D(f_0 f_1)} + \frac{\log \log n}{D(f_0 f_1)}$	sufficient if s sub-linear in n

Our results are striking primarily for two reasons. First, *sequential* procedures succeed when the number of measurements per dimension increases at a rate logarithmic in the level of *sparsity*, i.e. $\log s$. In contrast, *non-sequential* procedures require the average number of measurements per dimension to increase at a rate logarithmic in the *dimension*, i.e. $\log n$. For signals where sparsity is sublinear in dimension, the gains of sequential methods are polynomial; in scenarios where the sparsity grows logarithmically, the gains are exponential.

Secondly, and perhaps equally as surprising, a simple procedure dubbed sequential thresholding achieves nearly optimal performance provided minor constraints on the level of sparsity are met (specifically, that s is sublinear in n). In terms of the average number of measurements per dimension, the procedure comes within an additive factor, doubly logarithmic in dimension, of the lower bound of any sequential procedure.

II. PROBLEM FORMULATION

Let \mathcal{S} be a sparse subset of $\{1, \dots, n\}$ with cardinality $s = |\mathcal{S}|$. For any index $i \in \{1, \dots, n\}$, the random variable $y_{i,j}$ is independent, identically distributed according to (1). That is, for all j , $y_{i,j}$ follows distribution $f_1(\cdot)$ if i belongs to \mathcal{S} , and follows $f_0(\cdot)$ otherwise. We refer to f_0 as the null distribution, and f_1 the alternative.

In this paper, we limit our analysis to exact recovery of \mathcal{S} using coordinate wise methods. Defining $\hat{\mathcal{S}}$ as an estimate of \mathcal{S} , the family wise error rate is given as:

$$\mathbb{P}(\mathcal{E}) = \mathbb{P}(\hat{\mathcal{S}} \neq \mathcal{S}) = \mathbb{P}\left(\bigcup_{i \notin \mathcal{S}} \mathcal{E}_i \cup \bigcup_{i \in \mathcal{S}} \mathcal{E}_i\right)$$

where \mathcal{E}_i , $i \notin \mathcal{S}$ is a false positive error event and \mathcal{E}_i , $i \in \mathcal{S}$ a false negative error event. To simplify notation, we define the false positive and false negative probabilities in the usual manner: $\alpha = \mathbb{P}(\mathcal{E}_i | i \notin \mathcal{S})$, and $\beta = \mathbb{P}(\mathcal{E}_i | i \in \mathcal{S})$.

The test to decide if component i belongs to \mathcal{S} is based on the normalized log-likelihood ratio. For y_j distributed i.i.d. f_0 or f_1 ,

$$t^{(m)} := \frac{1}{m} \sum_{j=1}^m \log \frac{f_1(y_j)}{f_0(y_j)}$$

which is a function of $(y_1, \dots, y_m) \in \mathcal{Y}^m$. The superscript m explicitly indicates the number of measurements used to form the likelihood ratio and is suppressed when unambiguous. The log-likelihood ratio is compared against a scalar threshold γ to hypothesize if a component follows f_0 or f_1 :

$$t \underset{f_0}{\overset{f_1}{\geq}} \gamma.$$

Additionally, the Kullback-Liebler divergence of distribution f_0 from f_1 is defined as:

$$D(f_1 || f_0) = \mathbb{E}_1 \left[\log \frac{f_1(y)}{f_0(y)} \right]$$

where $\mathbb{E}_1[\cdot]$ is expectation with respect to distribution f_1 .

A. Measurement procedures

To be precise in characterizing a measurement procedure, we continue with three definitions.

Definition II.1. Measurement Procedure. A procedure, denoted π , used to determine if $y_{i,j}$ is measured. If $\pi_{i,j} = 1$, then $y_{i,j}$ is measured. Conversely, if $\pi_{i,j} = 0$, then $y_{i,j}$ is not measured.

Definition II.2. Non-sequential measurement procedure. Any measurement procedure π such that $\pi_{i,j}$ is not a function of $y_{i,j'}$ for any j' .

Definition II.3. Sequential measurement procedure. A measurement procedure π in which $\pi_{i,j}$ is allowed to depend on prior measurements, specifically, $\pi_{i,j} : \{y_{i,1}, \dots, y_{i,j-1}\} \mapsto \{0, 1\}$.

B. Measurement Budget

In order to make fair comparison between measurement schemes, we limit the total number of observations of $y_{i,j}$ in expectation. For any procedure π , we require

$$\mathbb{E} \left[\sum_{i,j} \pi_{i,j} \right] \leq nm \quad (2)$$

for some integer m . This implies, on average, we use m or fewer observations per dimension.

III. SEQUENTIAL THRESHOLDING

Sequential thresholding, first presented in [2], relies on a simple bisection idea. The procedure consists of a series of K measurement passes, where each pass eliminates from consideration a proportion of the components measured on the prior pass. After the last pass the procedure terminates and the remaining components are taken as the estimate of \mathcal{S} . Sequential thresholding is described in the following algorithm.

Sequential Thresholding

input: $K \approx \log n$ steps, threshold γ

initialize: $\pi_{i,1} = 1$ for all i

for $k = 1, \dots, K$ **do**

for $\{i : \pi_{i,k} = 1\}$ **do**

measure: t_i

threshold: $\pi_{i,k+1} = \begin{cases} 1 & t_i > \gamma \\ 0 & \text{else} \end{cases}$

end for

end for

output: $\hat{\mathcal{S}} = \{i : \pi_{i,K+1} = 1\}$

A. Example of Sequential Thresholding

Sequential thresholding is perhaps best illustrated by example. Consider a simple case with measurement budget $m = 2$, $f_0 \sim \mathcal{N}(0, 1)$ and $f_1 \sim \mathcal{N}(\theta, 1)$ for some $\theta > 0$.

On the first pass, the procedure measures $y_{i,1}$ for all i , using n measurements (half of the total budget as $mn = 2n$). On subsequent passes, the procedure observes $y_{i,k}$ if $\pi_{i,k} = 1$. To set $\pi_{i,k+1}$ the procedure thresholds observations that fall below, for example, $\gamma = 0$, eliminating a proportion (approximately half in this case) of components following the null distribution:

$$\pi_{i,k+1} = \begin{cases} 1 & y_{i,k} > 0 \\ 0 & y_{i,k} \leq 0. \end{cases}$$

In words, if a measurement of component i falls below the threshold on any pass, that component is *not* measured for the remainder of the procedure, and not included in the estimate of \mathcal{S} . After $K \approx \log n$ passes, the procedure terminates, and estimates \mathcal{S} as the set of indices that have not been eliminated from consideration: $\hat{\mathcal{S}} = \{i : \pi_{i,K+1} = 1\}$.

B. Details of Sequential Thresholding

Sequential thresholding requires two inputs: 1) K , the number of passes, and 2) γ , a threshold. We define ρ as the probability a component following the null is eliminated on any given pass, which is related to the threshold as

$$\mathbb{P}(t_i^{(\rho m)} \leq \gamma | i \notin \mathcal{S}) = \rho.$$

Additionally, we restrict our consideration to $\rho \in [1/2, 1)$ – that is, the probability a null component is eliminated on a given pass is one half or greater.

On each pass, ρm (which we assume to be an integer) measurements of a subset of components are made, and the log-likelihood ratio $t_i^{(\rho m)}$ is formed for each component. As measurements are made in blocks of size ρm , we use *boldface* $\pi_{i,k}$ to indicate a block of measurements are taken of component i on the k th measurement pass. $\pi_{i,k}$ can be interpreted as a vector:

$$\pi_{i,k} = (\pi_{i,(k-1)\rho m+1}, \dots, \pi_{i,\rho m}).$$

With γ and $K \approx \log n$ as inputs, sequential thresholding operates as follows. First, the procedure initializes, setting $\pi_{i,1} = \mathbf{1}$. For passes $k = 1, \dots, K$ the procedure measures $t_i^{(\rho m)}$ if $\pi_{i,k} = \mathbf{1}$. To set $\pi_{i,k+1}$, the procedure tests the corresponding log-likelihood ratio against the threshold γ :

$$\pi_{i,k+1} = \begin{cases} \mathbf{1} & \text{if } t_i^{(\rho m)} > \gamma \\ \mathbf{0} & \text{else.} \end{cases}$$

That is, if $t_i^{(\rho m)}$ is below γ , no further measurements of component i are taken. Otherwise, component i is measured on the subsequent pass. By definition of γ , approximately ρ times the number of remaining components following f_0 will be eliminated on each pass; if $s \ll n$, each thresholding step eliminates approximately ρ times the total number of components remaining.

After pass K , the procedure terminates and estimates \mathcal{S} as the indices still under consideration: $\hat{\mathcal{S}} = \{i : \pi_{i,K+1} = \mathbf{1}\}$.

C. Measurement Budget

Sequential thresholding satisfies the measurement budget in (2) provided s grows sublinearly with n . For brevity, we argue the procedure comes arbitrarily close to satisfying the measurement budget for large n :

$$\begin{aligned} \mathbb{E} \left[\sum_{i,j} \pi_{i,j} \right] &\leq \sum_{k=0}^{K-1} ((1-\rho)^k (n-s)\rho m + s\rho m) \\ &\leq m(n-s) + msK\rho. \end{aligned}$$

Letting $K = \log n$, the procedure comes arbitrarily close to satisfying the constraint as n grows large. To be rigorous in showing the procedure satisfies (2), m can be replaced by $m-1$, and the analysis throughout holds.

D. Ability of Sequential Thresholding

We present the first of the three main theorems of the paper to quantify the performance of sequential thresholding.

Theorem III.1. Ability of sequential thresholding. *Provided*

$$m > \frac{\log s}{D(f_0||f_1)} + \frac{\log \log n}{D(f_0||f_1)} \quad (3)$$

sequential thresholding recovers \mathcal{S} with high probability. More precisely, if

$$\lim_{n \rightarrow \infty} \frac{m}{\log(s \log n)} > \frac{1}{D(f_0||f_1)}$$

then $\mathbb{P}(\mathcal{E}) \rightarrow 0$.

Proof: From a union bound on the family wise error rate, we have

$$\mathbb{P}(\mathcal{E}) \leq (n-s)\alpha + s\beta. \quad (4)$$

Employing sequential thresholding, from the definition of γ , $\alpha = (1-\rho)^K$ and

$$\begin{aligned} \beta &= \mathbb{P} \left(\bigcup_{k=1}^K t_i^{(\rho m)} < \gamma | i \in \mathcal{S} \right) \\ &\leq K \mathbb{P} \left(t_i^{(\rho m)} < \gamma | i \in \mathcal{S} \right) \end{aligned}$$

where the inequality follows from a union bound.

We can further bound the false negative error event using the Chernoff-Stein Lemma [5], p. 384. Consider a simple binary hypothesis test with a fixed probability of false positive at $\alpha_0 = 1-\rho$. By the Chernoff-Stein Lemma, the false negative probability is then given as

$$\mathbb{P} \left(t_i^{(\rho m)} < \gamma | i \in \mathcal{S} \right) \doteq e^{-\rho m D(f_0||f_1)}$$

where $a \doteq e^{-mD}$ is equivalent to

$$\lim_{m \rightarrow \infty} \frac{1}{m} \log a = -D.$$

This implies, for any $\epsilon_1 > 0$, for sufficiently large m ,

$$\mathbb{P} \left(t_i^{(\rho m)} < \gamma | i \in \mathcal{S} \right) \leq e^{-\rho m (D(f_0||f_1) - \epsilon_1)}.$$

Letting $K = (1 + \epsilon_2) \log n$, for sufficiently large n and m , (4) becomes

$$\mathbb{P}(\mathcal{E}) \leq \frac{(n-s)}{n^{(1+\epsilon_2)}} + s(1+\epsilon_2) \log(n) e^{-\rho m (D(f_0||f_1) - \epsilon_1)}.$$

Hence, $\mathbb{P}(\mathcal{E})$ goes to zero provided

$$m \geq \frac{\log((1+\epsilon_2)s \log n)}{\rho(D(f_0||f_1) - \epsilon_1)}$$

which, as ϵ_1 and ϵ_2 can be made arbitrarily small, and ρ can be made arbitrarily close to 1, directly gives the theorem:

$$m > \frac{\log(s \log n)}{D(f_0||f_1)}.$$

■

IV. LOWER BOUND ON SEQUENTIAL PROCEDURES

In this section we derive a lower bound on the rate at which m must grow with n for any sequential procedure, and relate Sequential Thesholding to the high dimensional extension of the well known sequential probability ratio test (SPRT).

A. Limitation of any sequential procedure

The lower bound for any sequential procedure is presented in the following theorem.

Theorem IV.1. *Consider any sequential measurement procedure. Provided*

$$m < \frac{\log s}{D(f_0||f_1)}$$

the family wise error rate tends to one. More precisely, if

$$\lim_{n \rightarrow \infty} \frac{m}{\log s} < \frac{1}{D(f_0||f_1)} \quad (5)$$

then $\mathbb{P}(\mathcal{E}) \rightarrow 1$.

Proof: First, we show conditions under which the family wise error rate goes to one:

$$\begin{aligned} \mathbb{P}(\mathcal{E}) &= \mathbb{P}\left(\bigcup_{i \notin \mathcal{S}} \mathcal{E}_i \cup \bigcup_{i \in \mathcal{S}} \mathcal{E}_i\right) \\ &= 1 - \mathbb{P}\left(\bigcap_{i \notin \mathcal{S}} \mathcal{E}_i^c \cap \bigcap_{i \in \mathcal{S}} \mathcal{E}_i^c\right) \\ &= 1 - (1 - \beta)^s (1 - \alpha)^{n-s} \\ &\geq 1 - e^{-\beta s} e^{-\alpha(n-s)} \end{aligned} \quad (6)$$

which goes to one provided either

$$\alpha > \frac{1}{n-s} \quad \beta > \frac{1}{s}. \quad (7)$$

Second, for a simple binary hypothesis test, we can bound the expected number of measurements of any sequential procedure with false positive and false negative probabilities α and β . To simplify notation, define:

$$\mathbb{E}_0[N] = \mathbb{E}\left[\sum_j \pi_{i,j} | i \notin \mathcal{S}\right] \quad \mathbb{E}_1[N] = \mathbb{E}\left[\sum_j \pi_{i,j} | i \in \mathcal{S}\right]$$

that is, $\mathbb{E}_0[N]$ and $\mathbb{E}_1[N]$ are the expected number of measurements under f_0 and f_1 respectively. From [6] p.21, we have

$$E_0[N] \geq \frac{1}{D(f_0||f_1)} \left(\alpha \log \frac{\alpha}{1-\beta} + (1-\alpha) \log \frac{1-\alpha}{\beta} \right)$$

which is derived from a simple argument using Jensen's inequality. The total expected number of measurements, constrained by the measurement budget, is

$$(n-s)E_0[N] + sE_1[N] = E\left[\sum_{i,j} \pi_{i,j}\right] \leq mn \quad (8)$$

Dropping the $sE_1[N]$ term from (8), we need to find conditions under which the inequality

$$\frac{n-s}{D(f_0||f_1)} \left(\alpha \log \frac{\alpha}{1-\beta} + (1-\alpha) \log \frac{1-\alpha}{\beta} \right) \leq mn$$

implies $\mathbb{P}(\mathcal{E}) \rightarrow 1$. Dividing by $n \log s$, the inequality becomes

$$\frac{n-s}{D(f_0||f_1)n \log s} \left(\alpha \log \frac{\alpha}{1-\beta} + (1-\alpha) \log \frac{1-\alpha}{\beta} \right) \leq \frac{m}{\log s}.$$

Imposing the condition in (5) and cancelling $D(f_0||f_1)$ from both sides, the above inequality requires

$$\lim_{n \rightarrow \infty} \frac{n-s}{n \log s} \left(\alpha \log \frac{\alpha}{1-\beta} + (1-\alpha) \log \frac{1-\alpha}{\beta} \right) < 1. \quad (9)$$

It is sufficient to show that (9) implies either $\alpha > \frac{1}{n-s}$ or $\beta > \frac{1}{s}$ in the high dimensional limit.

With this in mind, let $\beta = \frac{1-\epsilon_1}{s}$, and $\alpha = \frac{1-\epsilon_2}{n-s}$ for some $\epsilon_1, \epsilon_2 \in [0, 1)$. Taking the limit as $n \rightarrow \infty$ in (9) and reducing terms we have:

$$\lim_{n \rightarrow \infty} (\cdot) = 1 \quad (10)$$

which contradicts (9), and negates our assumption that both $\beta = \frac{1-\epsilon_1}{s}$ and $\alpha = \frac{1-\epsilon_2}{n-s}$ for $\epsilon_1, \epsilon_2 \in [0, 1)$. Hence, by (7), the family wise error rate must go to one, completing the proof. ■

B. The SPRT

The sequential probability ratio test can be extended from simple binary hypothesis tests to the high dimensional case by simply considering n parallel SPRTs. Each individual SPRT operates by continuing to measure a component if the corresponding likelihood ratio is within an upper and lower boundary, and terminating measurement otherwise. For scalars A and B

$$\pi_{i,j+1} = \begin{cases} 1 & \text{if } \frac{A}{j} < t_i^{(j)}(y_{i,1}, \dots, y_{i,j}) < \frac{B}{j} \\ 0 & \text{else} \end{cases}$$

where $t_i^{(j)}$ is the normalized log-likelihood ratio comprised of all prior measurements (unlike sequential thresholding, in which the likelihood ratio is only formed using measurements from a single pass). If $t_i^{(j)} < A/j$, the SPRT labels index i as not belonging to \mathcal{S} , and if $t_i^{(j)} > B/j$, index i is assigned to \mathcal{S} . For a thorough discussion of the SPRT, see [6].

Sequential probability ratio tests are optimal for binary hypothesis tests in terms of minimum expected number of measurements for any error probabilities α and β (shown originally in [7]), and this optimality can be translated to the high dimensional case. Consider a single component i , and the corresponding binary hypothesis test. To be thorough, we restate the optimal property of the SPRT in the following lemma.

Lemma IV.2. *Optimality of the SPRT for simple binary tests [8] (p.63). Consider an SPRT with expected number of measurements $\mathbb{E}_0[N]$ and $\mathbb{E}_1[N]$, and corresponding error*

probabilities α and β . Any other sequential test with expected number of measurements $\mathbb{E}_0[N]'$ and $\mathbb{E}_1[N]'$ and error probabilities $\alpha' \leq \alpha$ and $\beta' \leq \beta$ will also have $\mathbb{E}_0[N] \geq \mathbb{E}_0[N]'$ and $\mathbb{E}_1[N] \geq \mathbb{E}_1[N]'$.

In short, no procedure with the smaller error probabilities can have fewer measurements in expectation than the SPRT. To translate the optimality of the SPRT to the high dimensional case, we introduce the following lemma.

Lemma IV.3. Optimality of the SPRT. Consider n component-wise sequential probability ratio tests used to estimate \mathcal{S} each with error probabilities α and β , and with a total of $\mathbb{E}[\sum_{i,j} \pi_{i,j}]$ measurements in expectation. Any other component wise test with $\alpha' \leq \alpha$ and $\beta' \leq \beta$ will also have expected number of measurements $\mathbb{E}[\sum_{i,j} \pi_{i,j}]' \geq \mathbb{E}[\sum_{i,j} \pi_{i,j}]$.

Proof: We can write the total expected number of measurements as:

$$\mathbb{E} \left[\sum_{i,j} \pi_{i,j} \right] = (n-s)\mathbb{E}_0[N] + s\mathbb{E}_1[N]$$

which is monotonically increasing in both $\mathbb{E}_0[N]$ and $\mathbb{E}_1[N]$. Together with IV.2, this implies the lemma. ■

C. Comparison of the SPRT to Sequential Thresholding

Although a fully rigorous proof is quite involved, using standard approximations for the sequential probability ratio test (again, see [6]) it is relatively straightforward to show the SPRT does achieve the lower bound presented above.

Sequential thresholding is similar in spirit to the SPRT. In many scenarios, however, implementing the SPRT can be substantially more complicated, if not infeasible, when compared to sequential thresholding. To set the stopping boundaries, an SPRT requires knowledge of the underlying distributions as well as the level of sparsity s . Even when these are available, only approximations relating error probabilities to the stopping boundaries can be derived in closed-form.

On the contrary, sequential thresholding does not require knowledge of s . Since its sample requirements are within a factor a small factor of the lower bound, sequential thresholding is automatically *adaptive* to unknown levels of sparsity. Moreover, in practice, sequential thresholding needs only approximate knowledge of the distributions to operate (such that a substantial number of components that follow f_0 can be eliminated on each pass).

V. LIMITATION OF NON-SEQUENTIAL METHODS

Our analysis would not be complete without comparison of sequential thresholding and the sequential lower bound to the performance limits of non-sequential methods. To do so, we analyze performance of any non-sequential method using *Chernoff Information*.

Theorem V.1. Limitation of non-sequential testing. Consider any non-sequential thresholding procedure. Provided

$$m < \frac{\log n}{D(f_1||f_0)} \quad (11)$$

the family wise error rate goes to 1. To be precise, (11) is equivalent to

$$\lim_{n \rightarrow \infty} \frac{m}{\log n} < \frac{1}{D(f_1||f_0)},$$

which implies $\lim_{n \rightarrow \infty} \mathbb{P}(\mathcal{E}) = 1$.

Proof: From [5], p. 386, (Chernoff Information) and by (7) any non-sequential test fails provided

$$\alpha \doteq e^{-mD(f_\lambda||f_0)} > \frac{1}{n-s}$$

or

$$\beta \doteq e^{-mD(f_\lambda||f_1)} > \frac{1}{s}$$

where

$$f_\lambda = \frac{f_0^\lambda f_1^{1-\lambda}}{\int_{\Omega} f_0^\lambda f_1^{1-\lambda} dy}$$

for $\lambda \in [0, 1]$. Hence, any sequential procedure fails provided

$$m < \min_{\lambda \in [0,1]} \max \left(\frac{\log(n-s)}{D(f_\lambda||f_0)}, \frac{\log s}{D(f_\lambda||f_1)} \right)$$

which is implied if

$$m < \frac{\log n}{D(f_1||f_0)}$$

completing the proof. ■

VI. CONCLUSION

This paper showed sequential methods for support recovery of high dimensional sparse signals in noise can succeed using far fewer measurements than non-sequential methods. Specifically, non-sequential methods require the number of measurements to grow logarithmically with the dimension, while sequential methods succeed if the number of measurements grows logarithmically with the level of sparsity. Additionally, a simple procedure termed sequential thresholding comes within a small additive factor of the lower bound in terms of number of measurements per dimension.

REFERENCES

- [1] J. Haupt, R. Castro, and R. Nowak, "Distilled sensing: Selective sampling for sparse signal recovery," <http://arxiv.org/abs/1001.5311>.
- [2] M. Malloy and R. Nowak, "Sequential analysis in high-dimensional multiple testing and sparse recovery," in *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, 31 2011-aug. 5 2011, pp. 2661–2665.
- [3] V. Srivastava, K. Plarre, and F. Bullo, "Randomized sensor selection in sequential hypothesis testing," *Signal Processing, IEEE Transactions on*, vol. 59, no. 5, pp. 2342–2354, may 2011.
- [4] A. R. Barron and A. Joseph, "Analysis of fast sparse superposition codes," *Proc. IEEE International Symposium on Information Theory*, 2011.
- [5] T. M. Cover and J. A. Thomas, *Elements of information theory*, 2nd ed. New York, NY, USA: Wiley-Interscience, 2005.
- [6] D. Siegmund, *Sequential Analysis*. New York, NY, USA: Springer-Verlag, 2010.
- [7] A. Wald and J. Wolfowitz, "Optimum character of the sequential probability ratio test," *The Annals of Mathematical Statistics*, vol. 19, no. 3, pp. 326–339, 1948.
- [8] H. Chernoff, *Sequential Analysis and Optimal Design*. Bristol, England: Society for Industrial and Applied Mathematics, 1979.