Estimation of Sparsity via Simple Measurements

Abhishek Agarwal Larkin Flodin Arya Mazumdar

***Abstract*—We consider several related problems of estimating the 'sparsity' or number of nonzero elements d in a length n vector $\mathbf{x}$ by observing only $\mathbf{b} = M \odot \mathbf{x}$, where M is a predesigned test matrix independent of $\mathbf{x}$, and the operation $\odot$ varies between problems. We aim to provide a Δ-approximation of sparsity for some constant Δ with a minimal number of measurements (rows of M). This framework generalizes multiple problems, such as estimation of sparsity in group testing and compressed sensing. We use techniques from coding theory as well as probabilistic methods to show that $O(D \log D \log n)$ rows are sufficient when the operation $\odot$ is logical OR (i.e., group testing), and nearly this many are necessary, where D is a known upper bound on d. When instead the operation $\odot$ is multiplication over $\mathbb{R}$ or a finite field, we show that respectively $\Theta(D)$ and $\Theta(D \log \frac{n}{D})$ measurements are necessary and sufficient.**

I. Introduction

Suppose that we want to identify an n dimensional vector $\mathbf{x} \in \mathbb{F}^n$, however, we can only observe the output $\mathbf{b}$ where

$$M \odot \mathbf{x} = \mathbf{b} \tag{1}$$

for a designed matrix $M \in \mathbb{F}^{m \times n}$. Let M_{ij} denote the (i, j)th entry of M and let $\mathbf{x}_i$ denote the ith component of $\mathbf{x}$. We will frequently refer to a single row of M as a "test" or "measurement." If we define the operation $\odot$ in eq. (1) as standard matrix multiplication over the field, and $\mathbb{F} = \mathbb{R}$, the problem of identifying $\mathbf{x}$ is known as the compressed sensing problem (and solvable with $m \ll n$ tests when $\mathbf{x}$ is 'sparse', see e.g. [5]). If instead we define the operation $\odot$ as the logical OR, so that $\mathbf{b}_i := \bigvee_{j:M_{ij}=1} \mathbf{x}_j$, and $\mathbb{F} = \mathbb{F}_2$, the identification problem is known as the group testing problem. We can easily identify $\mathbf{x}$ when $m \geq n$, for example by taking M to be the identity matrix. Let d be the *sparsity* or number of nonzero entries of the vector $\mathbf{x}$. For the case when it is known a priori that $\mathbf{x}$ is sparse (that is to say $d \ll n$), it has been shown that $m = O(d^2 \log n)$ measurements are sufficient for identification in the group testing setting, and even fewer measurements, $m = 2d$, are necessary and sufficient in the compressed sensing setting. The group testing result is tight up to a $\log d$ factor (see, e.g., [10], [15]).

Without any information about $\mathbf{x}$, it would be desirable to first estimate d, and then use this estimate to choose a suitable strategy to identify the d-sparse vector $\mathbf{x}$. A considerable body of research (see surveys in [13], [10]) has focused on the identification problem when an upper bound on d is known in advance. In comparison less work has been done on the problem of simply estimating d, without trying to determine which specific entries are nonzero [9], [8], [6], [7], [17], [2]. For group testing in the adaptive setting (when each subsequent test can depend on the results of previous tests), a recent result of Falahatgar et al. [11] allows approximation of d in as few as $O(\log \log d)$ tests, though with a small probability of error. In this paper, we solely concentrate on the nonadaptive version.

In particular, we provide tight upper and lower bounds on the number of measurements required by any deterministic algorithm (a predesigned matrix) for estimating d within a constant multiplicative factor of Δ (Δ-approximation) in three different settings of this problem (i.e., definitions of the operation $\odot$ and field $\mathbb{F}$ in eq. (1)). Note that Δ-approximation implies an estimate $\hat{d}$ such that

$$\frac{1}{\Delta} \leq \frac{\hat{d}}{d} \leq \Delta. \tag{2}$$

Earlier results by Damaschke and Muhammad [9], [8] in the nonadaptive group testing setting show that when no upper bound on the number of defectives is known, $O(\log n)$ queries are needed to approximate d even when the scheme is allowed to fail for a small number of inputs, and that any deterministic strategy capable of exactly determining d requires enough information to reconstruct the vector $\mathbf{x}$. In the group testing model, we restrict our attention to the unstudied problem of bounds for nonadaptive approximation schemes which work for all inputs, i.e., those that always produce an estimate of d within the specified range of allowable estimates.

The majority of work in the area of compressed sensing is concerned with the more difficult identification problem of how to proceed when $\mathbf{x}$ is not exactly sparse, but instead is approximately sparse, meaning it is close in ℓ_2 norm to a sparse vector. To our knowledge all such works are concerned with recovery of the vector $\mathbf{x}$, rather than estimation of its sparsity. Another set of works from both the compressed sensing and signal processing literatures [14], [16], [18], [12] focus on the related problem of "sparsity pattern recovery," which involves identifying the positions of the nonzero (or largest) entries of the vector $\mathbf{x}$, but not their values.

We will require our tests to be non-adaptive, but in contrast to existing works, we focus only on estimating the sparsity of $\mathbf{x}$, under the assumptions of absolute sparsity (as opposed to approximate sparsity) and no additional noise, over both finite fields and $\mathbb{R}$.

First, in section II we identify a necessary and sufficient condition on the matrix M for Δ-approximation to be possible. This condition applies to all models which we consider. In section III, we look at the group testing model specifically; given an upper bound D on the number of defectives d, we demonstrate a lower bound $\Omega(\min\{n, D \log \frac{n}{D\Delta^2}\})$ on the number of measurements needed to Δ-approximate d without error. The lower bound uses a simple and elegant combinatorial approach to bound the size of a cover of the

College of Information and Computer Sciences, University of Massachusetts Amherst. {abhiag,lflodin,arya}@cs.umass.edu. This work was supported by NSF CCF Awards 1318093, 1642658, 1642550.

space of possible input vectors $\mathbf{x}$, by showing all such vectors with the same output $\mathbf{b} = M \odot \mathbf{x}$ must be elements of the same poset of the Boolean lattice. We also show that this lower bound is nearly tight by demonstrating the existence of a matrix M with $m = O(\min\{n, D \log n \frac{\log D}{\log \Delta}\})$ rows that Δ-approximates d. In section IV, we take the operation $\odot$ to be multiplication in either $\mathbb{F}_q$ or $\mathbb{R}$, and exhibit close connections between the number of tests needed to approximate d and existing quantities in coding theory; this allows us to prove lower and upper bounds on the number of tests that are tight up to constant factors. Our main results are summarized in table I, though some bounds are more fine-grained than the table implies.

TABLE I: Number of Measurements Needed

Model	$D = n$	$D = o(n)$	
		Lower Bound	Upper Bound
Group Testing	$\Theta(n)$	$D \log\left(\frac{n}{D\Delta^2}\right)$	$O(\frac{\log D}{\log \Delta} D \log n)$
Compressed Sensing over $\mathbb{F}_q, q < n$	$\Theta(n)$	$\frac{D}{2} \log_q \frac{n}{D}$	$2D \log_q \frac{n}{D}$
Compressed Sensing over $\mathbb{R}$ or $\mathbb{F}_q, q \geq n$	$\Theta(n)$	$D - 1^*$	$2D$

*Assuming $D \geq 2\lfloor \Delta^2 \rfloor - 4$.

II. Preliminaries and Condition for Approximability

Throughout, we write $\log x$ to mean $\log_2 x$, and $||\mathbf{x}||_{\ell_0}$ to mean the sparsity or number of nonzero entries of the vector $\mathbf{x}$. We will denote the set $\{1, 2, \ldots, n\}$ by $[n]$. An ℓ-subset of a set S is a simply a subset of S of size ℓ.

For the estimation problem defined in section I in all models, we have the following necessary and sufficient condition on the matrix M for it to be used to Δ-approximate d, the sparsity of the vector in question. For a set $S \subseteq [n]$, let $\mathbf{v}(S)$ denote the set of vectors with support S.

Theorem 1. *Let $\Delta > 1$. Let $M \in \mathbb{F}^{m \times n}$ be a matrix such that there exists a decoder producing a Δ-approximation $\hat{d}$ of $d = ||\mathbf{x}||_{\ell_0}$ from observing $\mathbf{b} = M \odot \mathbf{x}$, for any $\mathbf{x} \in \mathbb{F}^n$, and assuming D to be a known upper bound on d. Consider sets $S_1, S_2 \subseteq [n]$ such that $|S_1| \leq |S_2| \leq D$. Then M satifies,*

$$\frac{|S_1|}{|S_2|} > \Delta^2 \implies (M \odot \mathbf{v}(S_1)) \cap (M \odot \mathbf{v}(S_2)) = \emptyset. \quad (3)$$

where $M \odot \mathbf{v}(S) \coloneqq \{M \odot \mathbf{v} : \mathbf{v} \in \mathbf{v}(S)\}$. Conversely, for any matrix M satisfying eq. (3)*, there exists a decoder producing an estimate $\hat{d}$ that satisfies the approximation criteria in eq.* (2).

Proof. In the forward direction, given a matrix M we show that if eq. (3) is not satisfied for two vectors, then no decoder satisfying eq. (2) can exist. Therefore, assume that for $S_1, S_2 \subseteq [n]$, $\frac{|S_1|}{|S_2|} \geq \Delta^2$ and $M \odot \mathbf{v}(S_1) \cap M \odot \mathbf{v}(S_2) \neq \emptyset$. Then a deterministic decoder must output the same estimate $\hat{d}$ when the observation belongs to $M \odot \mathbf{v}(S_1) \cap M \odot \mathbf{v}(S_2)$. Now, any estimate $\hat{d}$ will violate eq. (2) for at least one of a pair of vectors in $\mathbf{v}(S_1)$ and $\mathbf{v}(S_2)$.

For the converse assume eq. (3) holds for the matrix M. When observing the result vector $\mathbf{b}$, we define

$$\hat{d} = \sqrt{\Big(\max_{\substack{\mathbf{y} \in \mathbb{F}^n: \\ ||\mathbf{y}||_{\ell_0} \leq D \\ \mathbf{b} = M \odot \mathbf{y}}} ||\mathbf{y}||_{\ell_0}\Big)\Big(\min_{\substack{\mathbf{y} \in \mathbb{F}^n: \\ \mathbf{b} = M \odot \mathbf{y}}} ||\mathbf{y}||_{\ell_0}\Big)},$$

i.e., we estimate d by the geometric mean of the minimum and maximum weight vectors $\mathbf{y}$ with weight $\leq D$ and $M \odot \mathbf{y} = \mathbf{b}$.

From eq. (3) we are guaranteed that the estimate above is within a Δ factor of both the lowest and highest weight vectors $\mathbf{y}$ with $\mathbf{b} = M \odot \mathbf{y}$, so will satisfy eq. (2). □

Note that the above proof of decoder existence does not imply the existence of an efficient decoder; the decoder described in the proof may take time exponential in n to determine the minimum and maximum weight vectors $\mathbf{y}$ with $\mathbf{b} = M \odot \mathbf{y}$. We will see that despite this, efficient decoding is possible for some specific matrices with this property.

III. Defective Approximation for Group Testing

Throughout this section we take $\mathbb{F} = \mathbb{F}_2$ and assume that the operation $\odot$ denotes logical OR as defined in section I. For a subset $S \subseteq [n]$, we write $\mathbf{v}(S)$ to mean the unique binary vector with support S. We establish upper and lower bounds on the number of rows of a matrix capable of Δ-approximating d without error. As is standard in nonadaptive group testing, we will typically assume that an initial upper bound D on the number of defectives d is known. We then show that the number of rows m of the matrix M in eq. (1) must satisfy $m = \Omega(D \log \frac{n}{D})$. On the other hand, we show that there exists a matrix M with $m = O(D \log D \log n)$ rows satisfying eq. (3). When no bound is known on the number of defectives, we can take $D = n$, and in this scenario our bound shows that a linear (in n) number of measurements is required to find an estimate $\hat{d}$ satisfying eq. (2). This implies only constant factor improvement is possible, as n measurements suffice to determine d exactly.

Theorem 2. *Suppose $M \in \mathbb{F}_2^{m \times n}$ is a matrix capable of Δ-approximating d when $d \leq D \leq \frac{n}{2}$. Then $m \geq D \log \frac{n}{D} - 2D \log \Delta - D \log e$.*

Proof. Assume two sets $S_1, S_2 \subseteq [n]$ satisfy $M \odot \mathbf{v}(S_1) = M \odot \mathbf{v}(S_2)$. Then using the definition of the operation $\odot$ as logical OR, we have $M \odot \mathbf{v}(S_1) = M \odot \mathbf{v}(S_2) = M \odot \mathbf{v}(S_1 \cup S_2)$.

Let $\mathbb{P}_D([n])$ denote the set $\{A \subseteq [n] : |A| \leq D\}$ of possible defective subsets with size at most D. Let $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_t$ denote a partition of $\mathbb{P}_D([n])$ such that $A, B \in \mathcal{P}_i$ if and only if $M \odot \mathbf{v}(A) = M \odot \mathbf{v}(B)$. Then we must have

$$A, B \in \mathcal{P}_i \implies A \cup B \in \mathcal{P}_i. \quad (4)$$

Since the matrix M is capable of Δ-approximation, by theorem 1 any two sets A, B with $|A| \geq |B|$ in the same part of the partition $\mathcal{P}_i$ must also satisfy

$$|A| \leq \Delta^2 \cdot |B|. \quad (5)$$

Consider the lattice L on $\mathbb{P}_D([n])$ ordered by inclusion. Then the $\mathcal{P}_i$ must form posets in L with unique maximal elements, as they are closed under union. For any matrix M, if t is the number of distinct vectors $\mathbf{b}$ arising as $\mathbf{b} = M \odot \mathbf{x}$ for any potential vector of defectives $\mathbf{x}$, then each row of M corresponds to a test with either a positive or negative result, so there can be at most 2^m distinct vectors $\mathbf{b}$. Thus $t \leq 2^m$, so $m \geq \log t$. Since any partition is also a cover, we lower bound t by finding a lower bound on the minimum size of a cover of $\mathbb{P}_D([n])$ with posets satisfying eqs. (4) and (5).

From eq. (5), we see that a poset $\mathcal{P}$ containing an ℓ-subset of $[n]$ cannot have an element of size greater than $\ell\Delta^2$. Then by [3, Sec 4.7] we know that the minimum size t_ℓ of a cover of all ℓ-subsets of $[n]$ with posets satisfying eqs. (4) and (5) satisfies $t_\ell \geq \binom{n}{\ell}/\binom{\ell\Delta^2}{\ell}$.

The minimum size of any cover of $\mathbb{P}_D([n])$ is at least the size of the minimum cover of the ℓ-subsets of $[n]$, for any particular value of $\ell \leq D$. Thus, $t \geq \max_{\ell \leq D} \binom{n}{\ell}/\binom{\ell\Delta^2}{\ell} \geq \binom{n}{D}/\binom{D\Delta^2}{D}$. Then as $m \geq \log t$, we have

$$
\begin{aligned}
m &\geq \log\left(\binom{n}{D}\Big/\binom{D\Delta^2}{D}\right) \geq \log\left(\left(\frac{n}{D}\right)^D\Big/\frac{(D\Delta^2)^D}{D!}\right) \\
&= D\log n - 2D\log D - 2D\log\Delta + \log D! \\
&\geq D\log n - D\log D - 2D\log\Delta - D\log_2(e).
\end{aligned}
$$

using Stirling's approximation in the final step. □

Now consider the case when no nontrivial upper bound on d is known. When $D = n$ we have the following result, proved straightforwardly by bounding $\max_{\ell \leq n} \binom{n}{\ell}/\binom{\ell\Delta^2}{\ell}$.

Corollary 3. *Consider a matrix $M \in \mathbb{F}_2^{m\times n}$ that can be used for Δ-approximation in the group testing model with only the trivial upper bound $D = n$ on the true number of defectives d, and Δ a constant. Then for large n,*

$$
m \geq n\left(h\left(\frac{1}{\Delta^4}\right) - \frac{1}{\Delta^2}h\left(\frac{1}{\Delta^2}\right)\right) = \Omega(n),
$$

where $h(x)$ denotes the binary entropy function.

We now show the existence of a matrix M capable of Δ-approximation with $m = O(\frac{D\log D}{\log\Delta}\log n)$ rows for any D and $\Delta > 1$. We modify a construction of Damaschke and Muhammad [8] for nonadaptive sparsity approximation with error, and show that with additional repetition of certain tests, we can achieve no error.

Theorem 4. *Given an initial upper bound D on the number of defectives, there exists a matrix $M \in \mathbb{F}_2^{m\times n}$ that can Δ-approximate d in the group testing model such that $m = O(\frac{D\log D}{\log\Delta}\log n)$. Furthermore, this matrix has a decoding scheme that requires only $O(m) = O(\frac{D\log D}{\log\Delta}\log n)$ time.*

Proof. Let $\text{Bern}(p)$ denote the Bernoulli variable on $\mathbb{F}_2$ such that $\mathbb{P}(\text{Bern}(p) = 1) = p$. Let $\hat{d}$ denote the estimate of d, which is specific to the matrix M. We show that a matrix constructed in a random way combined with a specific estimator $\hat{d}$ works simultaneously for all vectors of defectives with nonzero probability. To this end, we define the bad events E_1 and E_2 to be the events that our estimate is too far off, the complements of the estimation criteria in eq. (2):

$$
E_1 : \frac{\hat{d}}{d} > \Delta, \tag{6a}
$$

$$
E_2 : \frac{\hat{d}}{d} < \frac{1}{\Delta}. \tag{6b}
$$

Construction: We modify the random construction in [8] to construct $M \in \mathbb{F}_2^{m\times n}$. For this matrix M, we show that the probability of either bad event occurring over all defective vectors with $d \leq D$ is strictly less than 1 when $m = O(\frac{D\log D}{\log\Delta}\log n)$. Thus there exists a matrix M with m rows for which none of the bad events in eq. (6) occur for any defective vector of weight at most D.

Consider a fixed parameter $b > 1$, and let $\delta := \log_b\Delta$. Take s to be a fixed parameter such that $s < \delta - 1$, and let $l \in \{\lfloor\log_b\ln 2\rfloor, \lfloor\log_b\ln 2\rfloor + 1, \ldots, \lceil\log_b D\rceil\}$ denote indices for subsets of tests. For each index l, we construct t random identically and independently distributed (iid) tests such that each element is selected in the test for index l with probability $1 - (1 - \frac{1}{D})^{b^l}$. Then the total number of such indices is $N := \lceil\log_b D\rceil - \lfloor\log_b\ln 2\rfloor + 1$, so the matrix M consists of tN rows with the elements in row indices $j \in \{(l-1)t+1, (l-1)t+2, \ldots, lt\}$ selected randomly and independently as $\text{Bern}(1 - (1 - \frac{1}{D})^{b^l})$. Thus the probability of row j for $j \in \{(l-1)t+1, (l-1)t+2, \ldots, lt\}$ having a negative result (containing no defectives) is $q_l(d) := (1 - \frac{1}{D})^{db^l}$. Now, there exists an index $\ell(d) \in \{\lfloor\log_b\ln 2\rfloor, \lfloor\log_b\ln 2\rfloor + 1, \ldots, \lceil\log_b D\rceil\}$ (or simply ℓ) such that $\frac{1}{2} \leq q_{\ell-s}(d) < \frac{1}{2^{1/b}}$, because $q_{l+1} = q_l^b$.

Let $L \in \{\lfloor\log_b\ln 2\rfloor, \lfloor\log_b\ln 2\rfloor + 1, \ldots, \lceil\log_b D\rceil\}$ denote the random variable corresponding to the maximum index for which the majority of tests were negative. Our decoding algorithm will be to take the estimate $\hat{d}$ of d such that $q_{L-s}(\hat{d}) = \frac{1}{2}$, i.e., $\hat{d} = \frac{-1}{b^{L-s}\log(1-\frac{1}{D})}$. Then the probability of bad event E_1 for a defective set of size d is

$$
\begin{aligned}
\mathbb{P}(E_1) &= \mathbb{P}\left(\frac{\hat{d}}{d} > \Delta\right) \leq \mathbb{P}\left(\hat{d} \geq d\Delta\right) \\
&= \mathbb{P}\left(\frac{1}{2} \leq \left(1 - \frac{1}{D}\right)^{d\, b^{L+\delta-s}}\right) = \mathbb{P}\left(\frac{1}{2} \leq q_{L+\delta-s}(d)\right) \\
&= \mathbb{P}(L \leq \ell(d) - \delta) = \prod_{\ell(d)-\delta < l \leq \lceil\log_b D\rceil} F\left(\left\lceil\frac{t-1}{2}\right\rceil; t, q_l(d)\right),
\end{aligned}
$$

where $F(k; n, p) = \mathbb{P}(X \leq k)$ denotes the cumulative distribution function for $X = \sum_{i=1}^n X_i$ for $X_i \sim \text{Bern}(p)$ iid. Similarly, for bad event E_2 we have

$$
\begin{aligned}
\mathbb{P}(E_2) &= \mathbb{P}\left(\frac{\hat{d}}{d} < \frac{1}{\Delta}\right) \leq \mathbb{P}\left(\hat{d} \leq d/\Delta\right) \\
&= \mathbb{P}\left(\frac{1}{2} \geq \left(1 - \frac{1}{D}\right)^{d\, b^{L-\delta-s}}\right) = \mathbb{P}\left(\frac{1}{2} \geq q_{L-\delta-s}(d)\right) \\
&= \mathbb{P}(L \geq \ell(d) + \delta) = \sum_{l=\ell(d)+\delta}^{\lceil\log_b D\rceil} F\left(\left\lceil\frac{t-1}{2}\right\rceil; t, 1 - q_l(d)\right).
\end{aligned}
$$

Thus the probability of the event $\tilde{E}_2$, defined to be the union of the events E_2 for all defective sets of size $d \leq D$, can be upper bounded by union bound as follows:

$$
\mathbb{P}\left(\tilde{E}_2\right) \leq \sum_{1\leq i\leq D}\binom{n}{i}\mathbb{P}(E_2)
$$

$$\leq \sum_i \binom{n}{i} \sum_{\ell(d)+\delta \leq l \leq \lceil \log_b D \rceil} F\left(\left\lceil \frac{t-1}{2} \right\rceil; t, 1-q_l(i)\right)$$
$$\leq \sum_i \binom{n}{i} \cdot N \cdot F\left(\left\lceil \frac{t-1}{2} \right\rceil; t, 1-q_{\lceil \log_b D \rceil}(i)\right)$$
$$\leq \sum_i \binom{n}{i} \cdot N \cdot \exp\left(-t\, D\left(\frac{1}{2}||1-q_{\lceil \log_b D \rceil}(i)\right)\right),$$

where $D(a||p) := a\log(\frac{a}{p}) + (1-a)\log(\frac{1-a}{1-p})$ denotes the KL divergence between a and p, and the last inequality follows from the bound on $F(k;n,p)$ in [4]. Then applying the inequalities $\binom{n}{i} \leq n^i$ and $q_{\lceil \log_b D \rceil} \leq e^{-i}$, we have that $\mathbb{P}\left(\tilde{E_2}\right)$ is at most

$$N \sum_{1 \leq i \leq D} \exp\left(i \ln n + \frac{t}{2}\left(1 + \log(1 - q_{\lceil \log_b D \rceil}) - \frac{i}{\ln 2}\right)\right).$$

Now, it can be seen that for $t = O(\log n)$, $\mathbb{P}\left(\tilde{E_2}\right)$ goes to 0.

Similarly, we bound the probability of $\tilde{E_1}$, the union of the events E_1 for all defective sets of size $d \leq D$, recalling that $\delta > s+1$:

$$\mathbb{P}\left(\tilde{E_1}\right) \leq \sum_{1 \leq i \leq D} \binom{n}{i} \mathbb{P}(E_1)$$
$$\leq \sum_i \binom{n}{i} \prod_{\ell(i)-\delta < l \leq \lceil \log_b D \rceil} F\left(\left\lceil \frac{t-1}{2} \right\rceil; t, q_l(i)\right)$$
$$\leq \sum_i \binom{n}{i} \prod_{\ell(i)-\delta < l \leq \ell(i)-s} F\left(\left\lceil \frac{t-1}{2} \right\rceil; t, q_l(i)\right)$$
$$\leq \sum_i \binom{n}{i} \exp\left(-\frac{\delta - s - 1}{2q_{\ell(i)-s-1}} \cdot t \left(q_{\ell(d)-s-1} - \frac{1}{2}\right)^2\right),$$

where the last inequality follows from the Chernoff bound for Binomial distribution. Thus $\mathbb{P}\left(\tilde{E_1}\right)$ is at most

$$D \cdot \exp\left(D \ln n - \frac{\delta - s - 1}{2q_{\ell(d)-s-1}} \cdot t \left(q_{\ell(d)-s-1} - \frac{1}{2}\right)^2\right)$$
$$\leq D \cdot \exp\left(D \ln n - \frac{\delta - s - 1}{2^{1-b^{-1}}} \cdot t \left(\frac{1}{2^{b^{-1}}} - \frac{1}{2}\right)^2\right).$$

It can be seen that for $t = O(\frac{D}{\delta} \log n)$, $\mathbb{P}\left(\tilde{E_1}\right)$ goes to 0.

Therefore when the number of tests $m = O(\frac{D \log D}{\log \Delta} \log n)$, there exists a matrix which estimates d within a multiplicative factor of Δ for all defective vectors of weight $\leq D$.

The decoding requires only computing a function of the largest index l for which the corresponding block of t tests had a majority of test results negative, and this index can easily be determined in a single pass over the result vector, requiring $O(m)$ time. □

IV. Defective Approximation by Linear Operations

In this section we assume that the operation $\odot$ denotes standard matrix multiplication over a field $\mathbb{F}$, and as such will typically write $M\mathbf{x}$ instead of $M \odot \mathbf{x}$. We take $\mathbb{F}$ to be either finite or $\mathbb{R}$, the latter of which is the setting of the well-studied compressed sensing problem. Our aim now is to bound the size of a matrix M that satisfies the criteria for Δ-approximation in this model. Consider a vector space $V \subseteq \mathbb{F}^n$. Call such a vector space (Δ, D)-distinguishing if it has the property that for parameters $\Delta > 1$, $D \leq n$, any two vectors $\mathbf{x}$ and $\mathbf{y}$ in the same coset of V (so there exists $\mathbf{v}_1, \mathbf{v}_2 \in V$ such that $\mathbf{x} = \mathbf{z} + \mathbf{v}_1, \mathbf{y} = \mathbf{z} + \mathbf{v}_2$ for some $\mathbf{z} \in \mathbb{F}^n$) both having weight at most D differ in weight by a factor of at most Δ^2. In other words, if $||\mathbf{x}||_{\ell_0} \leq ||\mathbf{y}||_{\ell_0}$, then $||\mathbf{y}||_{\ell_0} \leq \Delta^2 ||\mathbf{x}||_{\ell_0}$. We can use this property to lower bound the rank (and thus the number of rows) of any Δ-approximating matrix.

Theorem 5. *For an $m \times n$ matrix M which can Δ-approximate the sparsity, d, of any vector with sparsity at most D in the linear operations model, it is necessary that*

$$\operatorname{rank}(M) \geq n - \max_{(\Delta, D)\text{-distinguishing } V} \dim(V). \tag{7}$$

Proof. We show that the nullspace of any such matrix M is a (Δ, D)-distinguishing vector space. By our necessary criteria for Δ-estimation, we have for any vectors $\mathbf{x}, \mathbf{y}$ with $||\mathbf{x}||_{\ell_0} \leq ||\mathbf{y}||_{\ell_0} \leq D$, that when $M\mathbf{x} = M\mathbf{y}$, then $||\mathbf{y}||_{\ell_0} \leq \Delta^2 ||\mathbf{x}||_{\ell_0}$. Now, since $M(\mathbf{x} - \mathbf{y}) = 0$ if and only if $\mathbf{x}$ and $\mathbf{y}$ are in the same coset of $\ker M$, we have $\operatorname{rank}(M) \geq \min\{\operatorname{rank}(N) : \ker N \text{ is } (\Delta, D)\text{-distinguishing}\}$. □

We use the condition in theorem 5 to bound the number of measurements needed, as $\operatorname{rank}(M)$ is a lower bound on the number of rows of M. Let $A_{\mathbb{F}}(n, d_{\min})$ denote the maximum dimension of a subspace in $\mathbb{F}^n$ that does not contain any nonzero vector in $\{\mathbf{x} \in \mathbb{F}^n : ||\mathbf{x}||_{\ell_0} < d_{\min}\}$. Note that for finite fields $\mathbb{F}_q$, the quantity $A_{\mathbb{F}_q}(n, d_{\min})$ denotes the maximum dimension of a linear q-ary code with minimum distance $d_{\min}$. A generalization of this quantity is studied in [1], which defines $A_{\mathbb{F}}(n, a, b)$ as the maximum dimension of a subspace in $\mathbb{F}^n$ that does not contain any vector in the annulus $\{\mathbf{x} \in \mathbb{F}^n : a < ||\mathbf{x}||_{\ell_0} < b\}$. Let U be a (Δ, D)-distinguishing vector space, and let $\mathbf{x}$ be a nonzero vector of maximum weight in U subject to $||\mathbf{x}||_{\ell_0} < D$. If no such vector exists, then $\dim U \leq A_{\mathbb{F}}(n, D)$. Otherwise, there exists a vector $\mathbf{y} \in \mathbb{F}^n$ with $||\mathbf{y}||_{\ell_0} = 1$ and $||\mathbf{y} + \mathbf{x}||_{\ell_0} = ||\mathbf{x}||_{\ell_0} + 1 \leq D$. Since $\mathbf{y}$ and $\mathbf{x} + \mathbf{y}$ are in the same coset of U, we have

$$||\mathbf{x} + \mathbf{y}||_{\ell_0} \leq \Delta^2 \underbrace{||\mathbf{y}||_{\ell_0}}_{=1} \implies ||\mathbf{x}||_{\ell_0} \leq \lfloor \Delta^2 \rfloor - 1.$$

Thus, $\dim U \leq A_{\mathbb{F}}(n, \lfloor \Delta^2 \rfloor - 1, D)$. As $A_{\mathbb{F}}(n, D) \leq A_{\mathbb{F}}(n, a, D)$ for any $a \leq D$, the above inequality applies to any (Δ, D)-distinguishing space U. Therefore we have

$$\operatorname{rank}(M) \geq n - A_{\mathbb{F}}(n, \lfloor \Delta^2 \rfloor - 1, D)$$

for any matrix M that can Δ-approximate d for every defective vector of sparsity up to D in the linear operations model. Define $m^\star_{\mathbb{F}} := n - A_{\mathbb{F}}(n, a, b)$. We have from [1, Proposition 1] that whenever $2a - 2 \leq b$, $m^\star_{\mathbb{F}}(a, b, n) = m^\star_{\mathbb{F}}(1, b, n - a + 1)$. The work in [1] is oriented towards finite fields, but their proof is independent of the choice of field $\mathbb{F}$, so in particular applies also when $\mathbb{F} = \mathbb{R}$. Note that $m^\star_{\mathbb{F}}(1, b, n)$ is just the rank of the parity check matrix of the largest dimension linear code on n coordinates with minimum distance b. Thus, as long as $D \geq 2\lfloor \Delta^2 \rfloor - 4$,

$$\operatorname{rank}(M) \geq m^{\star}_{\mathbb{F}}(1, D, n - \lfloor \Delta^2 \rfloor + 2).$$

If $\mathbb{F}$ is a finite field of size greater than or equal to n, or $\mathbb{F} = \mathbb{R}$, we know that $m^{\star}_{\mathbb{F}}(1, D, n - \lfloor \Delta^2 \rfloor + 2) \geq D - 1$ by the Singleton bound. Since we can exactly identify the set of defectives for $d \leq D$ in $2D$ queries using basic techniques from compressed sensing, there is at most a factor 2 improvement possible.

For $|\mathbb{F}| = q < n$, a stronger lower bound is possible since in general the Singleton bound is not tight. From the sphere-packing bound on $A_{\mathbb{F}}(n - \lfloor \Delta^2 \rfloor + 2, 1, D)$, we have that $m^{\star}_{\mathbb{F}_q}(1, D, n - \lfloor \Delta^2 \rfloor + 2)$ is lower bounded by

$$\left\lfloor \frac{D-1}{2} \right\rfloor \log_q(n - \lfloor \Delta^2 \rfloor + 2) - O(D \log D) = \Omega\left(D \log \frac{n}{D}\right),$$

as long as $\lfloor \frac{D-1}{2} \rfloor \leq \frac{n}{2}$.

For an upper bound on the minimum possible rank of M when $|\mathbb{F}| = q < n$, recall that such a matrix must have the property that in every coset of the nullspace, any two vectors of weight less than or equal to D must differ in weight by at most a Δ^2 factor. As the difference of any two vectors in the same coset lies in the nullspace, this condition is satisfied if every nonzero vector in the nullspace has weight at least $2D + 1$, so we have the following result.

Theorem 6. *In the linear operations model, when $|\mathbb{F}| = q < n$, and $D \leq \frac{n(q-1)-q}{2}$, there exist matrices M that can Δ-approximate the true number of defectives, d, for every possible vector of defectives of weight at most D, with at most $nH_q(\frac{2D+1}{n})$ rows, where H_q is the q-ary entropy function, $H_q(x) := x \log_q(q-1) - x \log_q(x) - (1-x) \log_q(1-x)$.*

Proof. We have

$$\min_{\substack{M \\ M\ \Delta-\text{approximates } d}} \operatorname{rank}(M) \leq \min_{\substack{M \\ \forall \mathbf{x} \neq 0 \in \ker M,\ ||\mathbf{x}||_{\ell_0} > 2D}} \operatorname{rank}(M) = m^{\star}_{\mathbb{F}_q}(1, 2D+1, n) \leq nH_q\left(\frac{2D+1}{n}\right),$$

where the last inequality follows from the asymptotic Gilbert-Varshamov bound, using the fact that $\frac{2D+1}{n} \leq 1 - \frac{1}{q}$. □

Corollary 7. *Under the conditions of the previous theorem, there exist matrices M that can Δ-approximate d with $(2D+1)\log_q\left(\frac{n}{2D+1}\right) + O(D) = O(D \log(n/D))$ rows for sufficiently large n.*

The above result is asymptotically tight with the lower bound given previously, so only improvements in the constant factor are possible. For certain settings of parameters these improvements follow easily from known results; for instance, the existence of binary BCH codes with $n - k \leq D \log_2(n+1)$ for certain values of n implies that when $q = 2$ the bound in corollary 7 improves by about a factor of 2.

Remark: Adaptive Tests. In the case of linear measurements over $\mathbb{R}$, a simple trick yields a very efficient adaptive scheme as well. First, suppose the entries of the vector $\mathbf{x}$ are nonnegative. In this case, the result of a test is nonzero if and only if some nonzero entry of $\mathbf{x}$ is included in the test. Then for the purpose of determining sparsity, each test tells us as least as much information as the corresponding test in the group testing model, as we can simply threshold the real-valued test result to 1 if it is nonzero and 0 otherwise. Thus we can exploit existing results for adaptive group testing sparsity approximation [11] to obtain an estimate of the sparsity that is accurate with high probability in as few as $O(\log \log d)$ measurements. However, if not all entries of the vector $\mathbf{x}$ are nonnegative, then there is the additional complication that it is possible to observe a test result of 0 even when a nonzero entry of $\mathbf{x}$ is included in the test. This will happen exactly when the vector corresponding to some test is orthogonal to $\mathbf{x}$. It is easy to counteract this by slightly perturbing each test vector, adding a small real-valued random vector that is nonzero only on the support of the test vector to each test. This ensures that the new test vector lies in the space orthogonal to $\mathbf{x}$ with probability 0.

References

[1] E. Abbe, N. Alon, A. S. Bandeira, and C. Sandon. Linear boolean classification, coding and the critical problem. *IEEE Transactions on Information Theory*, 62(4):1667–1673, 2016.

[2] J. Acharya, C. L. Canonne, and G. Kamath. Adaptive estimation in weighted group testing. In *IEEE International Symposium on Information Theory, ISIT*, pages 2116–2120. IEEE, 2015.

[3] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley Publishing, 4th edition, 2016.

[4] R. Arratia and L. Gordon. Tutorial on large deviations for the binomial distribution. *Bulletin of Mathematical Biology*, 51(1):125–131, 1989.

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

[6] Y. Cheng. An efficient randomized group testing procedure to determine the number of defectives. *Oper. Res. Lett.*, 39(5):352–354, 2011.

[7] Y. Cheng and Y. Xu. An efficient FPRAS type group testing procedure to approximate the number of defectives. *J. Comb. Optim.*, 27(2):302–314, 2014.

[8] P. Damaschke and A. S. Muhammad. Competitive group testing and learning hidden vertex covers with minimum adaptivity. *Discrete Math., Alg. and Appl.*, 2(3):291–312, 2010.

[9] P. Damaschke and A. S. Muhammad. Bounds for nonadaptive group tests to estimate the amount of defectives. *Discrete Math., Alg. and Appl.*, 3(4):517–536, 2011.

[10] D. Du and F. Hwang. *Combinatorial Group Testing and Its Applications*. Applied Mathematics. World Scientific, 2000.

[11] M. Falahatgar, A. Jafarpour, A. Orlitsky, V. Pichapati, and A. T. Suresh. Estimating the number of defectives with group testing. In *Information Theory (ISIT), 2016 IEEE International Symposium on*, pages 1376–1380. IEEE, 2016.

[12] A. K. Fletcher, S. Rangan, and V. K. Goyal. Necessary and sufficient conditions for sparsity pattern recovery. *IEEE Trans. Information Theory*, 55(12):5758–5772, 2009.

[13] A. C. Gilbert and P. Indyk. Sparse recovery using sparse matrices. *Proceedings of the IEEE*, 98(6):937–947, 2010.

[14] M. E. Lopes. Compressed sensing without sparsity assumptions. *CoRR*, abs/1507.07094, 2015.

[15] A. Mazumdar. Nonadaptive group testing with random set of defectives. *IEEE Trans. Information Theory*, 62(12):7522–7531, 2016.

[16] G. Reeves and M. Gastpar. Fundamental tradeoffs for sparsity pattern recovery. *CoRR*, abs/1006.3128, 2010.

[17] D. Ron and G. Tsur. The power of an example: Hidden set size approximation using group queries and conditional sampling. *CoRR*, abs/1404.5568, 2014.

[18] Y. Wang, Z. Tian, and C. Feng. Sparsity order estimation and its application in compressive spectrum sensing for cognitive radios. *IEEE Trans. Wireless Communications*, 11(6):2116–2125, 2012.