

Efficient and Near-Optimal Noisy Group Testing: An Information-Theoretic Framework

Jonathan Scarlett and Volkan Cevher

Abstract—The group testing problem consists of determining a small set of defective items from a larger set of items based on a number of tests, and is relevant in applications such as medical testing, communication protocols, pattern matching, and more. In this paper, we revisit an efficient algorithm for noisy group testing in which each item is decoded separately (Malyutov and Mateev, 1980), and develop novel performance guarantees via an information-theoretic framework for general noise models. For the special cases of no noise and symmetric noise, we find that the asymptotic number of tests required for vanishing error probability is within a factor $\log 2 \approx 0.7$ of the information-theoretic optimum at low sparsity levels, and that with a small number of allowed errors, this guarantee extends to all sublinear sparsity levels. In addition, we provide a converse bound showing that if one tries to move slightly beyond our low-sparsity achievability threshold using separate decoding of items and i.i.d. randomized testing, the average number of items decoded incorrectly approaches that of a trivial decoder.

I. INTRODUCTION

The group testing problem consists of determining a small subset S of “defective” items within a larger set of items $\{1, \dots, p\}$ based on a number of tests. This problem has a history in medical testing [1], and has regained significant attention with following applications in areas such as communication protocols [2], pattern matching [3], and database systems [4], and new connections with compressive sensing [5], [6]. In the noiseless setting, each test takes the form

$$Y = \bigvee_{i \in S} X_i, \quad (1)$$

where the test vector $X = (X_1, \dots, X_p) \in \{0, 1\}^p$ indicates which items are included in the test, and Y is the resulting observation. That is, the output indicates whether at least one defective item was included in the test. One wishes to minimize the total number of tests n while still ensuring the reliable recovery of S . We focus on the non-adaptive setting, in which all tests must be designed in advance. The corresponding test vectors $X^{(1)}, \dots, X^{(n)}$ are represented by the matrix $\mathbf{X} \in \{0, 1\}^{n \times p}$.

Following both classical works [7]–[10] and recent advances [11]–[13], the information-theoretic performance limits of group testing have become increasingly well-understood, and several practical near-optimal algorithms for the *noiseless* setting have been developed. In contrast, practical algorithms for *noisy* settings have generally remained less well-understood, with the best known theoretical guarantees usually being far from the information-theoretic limits (though sometimes matching in scaling laws) [14]–[16].

A notable exception to these limitations is the technique of Malyutov and Mateev [8] based on *separate decoding of*

items,¹ in which each given item $j \in \{1, \dots, p\}$ is decoded based only on the j -th column \mathbf{X}_j of \mathbf{X} , along with \mathbf{Y} (see Section I-B). This approach is computationally efficient, and was also proved to come with strong theoretical guarantees *in the case that* $k := |S| = O(1)$ [8], [17] (see Section I-C).

In this paper, we develop a theoretical framework for understanding separate decoding of items, and move beyond the work of [8] in several important directions: (i) We consider the general case of $k = o(p)$, thus handling much more general scenarios corresponding to “denser” settings, and leading to non-trivial challenges in the theoretical analysis; (ii) We consider not only exact recovery, but also partial recovery, often leading to much milder requirements on the number of tests; (iii) We provide a novel converse bound revealing that under separate decoding of items, our achievability bounds cannot be improved in several cases of interest.

Before discussing the previous work and our contributions in more detail, we formally state the setup.

A. Problem Setup

We let the defective set S be uniform on the $\binom{p}{k}$ subsets of $\{1, \dots, p\}$ of cardinality k . For convenience, we will sometimes equivalently refer to a vector $\beta \in \{0, 1\}^p$ whose j -th entry indicates whether or not item j is defective:

$$\beta_j = \mathbb{1}\{j \in S\}. \quad (2)$$

We consider i.i.d. Bernoulli testing, where each item is placed in a given test independently with probability $\frac{\nu}{k}$ for some constant $\nu > 0$. The vector of n observations is denoted by $\mathbf{Y} \in \{0, 1\}^n$, and the corresponding measurement matrix (each row of which contains a single measurement vector $X = (X_1, \dots, X_p)$) is denoted by $\mathbf{X} \in \{0, 1\}^{n \times p}$. Denoting the i -th entry of \mathbf{Y} by $Y^{(i)}$ and the i -th row of \mathbf{X} by $X^{(i)}$, the measurement model is given by

$$(Y^{(i)} | X^{(i)}) \sim P_{Y|N(S, X^{(i)})}, \quad (3)$$

where $N(S, X^{(i)}) = \sum_{j=1}^p \mathbb{1}\{j \in S \cap X_j^{(i)} = 1\}$ denotes the number of defective items in the test. That is, we consider arbitrary noise distributions $P_{Y|N}$ for which $Y^{(i)}$ depends on $X^{(i)}$ only through $N(S, X^{(i)})$, with conditional independence among the tests $i = 1, \dots, n$. For each item $j = 1, \dots, p$, the j -th column of \mathbf{X} is written as $\mathbf{X}_j \in \{0, 1\}^n$.

While most of our results will be written in terms of general noise models of the form (3), we also pay particular attention

¹This is referred to as *separate testing of inputs* in more recent works [17], [18], but there the word “testing” refers to a hypothesis test performed at the decoder, as opposed to testing the sense of designing \mathbf{X} .

to two specific models: The noiseless model in (1), and the *symmetric noise model* with parameter $\rho > 0$:

$$Y = \bigvee_{i \in S} X_i \oplus Z, \quad (4)$$

where $Z \sim \text{Bernoulli}(\rho)$, and \oplus denotes modulo-2 addition.

In the general case (i.e., not necessarily using separate decoding of items), given \mathbf{X} and \mathbf{Y} , a *decoder* forms an estimate \hat{S} of S , or equivalently, an estimate $\hat{\beta}$ of β . We consider two related performance measures. In the case of *exact* recovery, the error probability is given by

$$P_e := \mathbb{P}[\hat{S} \neq S], \quad (5)$$

and is taken over the realizations of S , \mathbf{X} , and \mathbf{Y} (the decoder is assumed to be deterministic). In addition, we consider a less stringent performance criterion in which we allow for up to $d_{\text{pos}} \in \{0, \dots, p-k-1\}$ false positives and $d_{\text{neg}} \in \{0, \dots, k-1\}$ false negatives, yielding an error probability of

$$P_e(d_{\text{pos}}, d_{\text{neg}}) := \mathbb{P}[|\hat{S} \setminus S| > d_{\text{pos}} \cup |S \setminus \hat{S}| > d_{\text{neg}}]. \quad (6)$$

In some cases (particularly for the converse) it will be convenient to consider yet another criterion in which we only seek to bound the *average number of incorrectly-decoded items* N_{err} :

$$\mathbb{E}[N_{\text{err}}] = \sum_{j=1}^p \mathbb{P}[\hat{\beta}_j \neq \beta_j]. \quad (7)$$

B. Separate Decoding of Items

We use the terminology *separate decoding of items* to mean any decoding scheme in which $\hat{\beta}_j$ is only a function of \mathbf{X}_j and \mathbf{Y} , i.e.,

$$\hat{\beta}_j = \phi_j(\mathbf{X}_j, \mathbf{Y}), \quad j = 1, \dots, p \quad (8)$$

for some functions $\{\phi_j\}_{j=1}^p$. All of our achievability results will choose ϕ_j not depending on j ; more specifically, following [8], each decoder is of the following form for some $\gamma > 0$:

$$\phi_j(\mathbf{X}_j, \mathbf{Y}) = \mathbb{1} \left\{ \sum_{i=1}^n \log \frac{P_{Y|X_j, \beta_j}(Y^{(i)} | X_j^{(i)}, 1)}{P_Y(Y^{(i)})} > \gamma \right\}, \quad (9)$$

where P_Y is the unconditional distribution of a given observation, and $P_{Y|X_j, \beta_j}(\cdot | \cdot, 1)$ is the conditional distribution given $\beta_j = 1$ and the value of X_j . This can be interpreted as the Neyman-Pearson test for binary hypothesis testing with hypotheses $H_0 : \beta_j = 0$ and $H_1 : \beta_j = 1$.

The computational complexity of (9) is $O(n)$ for each $j = 1, \dots, p$, for a total of $O(np)$. This matches the runtime of typical group testing algorithms [12], though it is slower than recent *sublinear-time* algorithms [15], [16]. Moreover, considerable speedups are possible via distributed implementations, as shown in [44], [45].

C. Related Work

Figure 1 summarizes the main results known for the noiseless and symmetric noise models (both information-theoretic and practical), along with our novel contributions. We proceed by outlining the relevant related work, and then describe our contributions in more detail and further discuss Figure 1.

Information-theoretic limits. The information-theoretic limits of group testing have long been well-understood for $k = O(1)$ in the Russian literature [7], [8], and have recently become increasingly well-understood for more general $k = o(p)$ [11], [13], [19], [20]. Here we highlight two results from [13], [20], [21] that are particularly relevant to our work:

- For the noiseless model (1) with Bernoulli testing and $k = \Theta(p^\theta)$ for some $\theta \in (0, 1)$, the minimal number of tests n^* ensuring high-probability exact recovery satisfies

$$n^* = \inf_{\nu > 0} \max \left\{ \frac{k \log \frac{p}{k}}{H_2(e^{-\nu})}, \frac{\frac{\theta}{1-\theta} k \log \frac{p}{k}}{e^{-\nu \nu}} \right\} (1 + o(1)), \quad (10)$$

where $H_2(\cdot)$ is the binary entropy function.² In particular, for $\theta \leq \frac{1}{3}$, we have $n^* = (k \log_2 \frac{p}{k})(1 + o(1))$, which is optimal even beyond Bernoulli testing.

- For the symmetric noise model (4) with $\rho > 0$, we have for *sufficiently small* θ that

$$n^* = \frac{k \log \frac{p}{k}}{\log 2 - H_2(\rho)} (1 + o(1)). \quad (11)$$

Moreover, if we move to partial recovery with $d_{\text{pos}} = d_{\text{neg}} = \alpha^* k$ for arbitrarily small $\alpha^* > 0$, then the right-hand side of (11) is achievable for all $\theta \in (0, 1)$ (including in the noiseless case $\rho = 0$).

Practical algorithms. In the noiseless setting, numerous practical group testing algorithms have been proposed with various theoretical guarantees. The best known bounds under Bernoulli testing were given for the *definite defectives* (DD) algorithm in [20], in particular matching (10) for $\theta \geq \frac{1}{2}$. Moreover, in [22], it was shown that the same bound is attained by the linear programming (LP) relaxation techniques of [18]. The DD algorithm was also shown to yield improved bounds under a random non-Bernoulli design in [23], but in this paper we focus on Bernoulli testing.

In noisy settings, less is known. For the symmetric noise model, an algorithm called *noisy combinatorial orthogonal matching pursuit* (NCOMP) [14] (also referred to as *noisy column matching* in [24]) was shown to achieve optimal $O(k \log p)$ scaling, but the constant factors in this result are quite suboptimal. NCOMP is in fact also an algorithm with separate decoding of items, albeit different than that of [8].

Some heuristic algorithms have been proposed for noisy settings without theoretical guarantees, including belief propagation [25] and a noisy LP relaxation [18]. A different LP relaxation was also given in [24] that only makes use of the negative tests; as a result, we found that it does not perform as well in practice. On the other hand, it was shown to yield optimal scaling laws in the symmetric noise model, with the constant factors again left loose to simplify the analysis.

²Here and subsequently, all logarithms have base e , and all information measures are in units of nats.

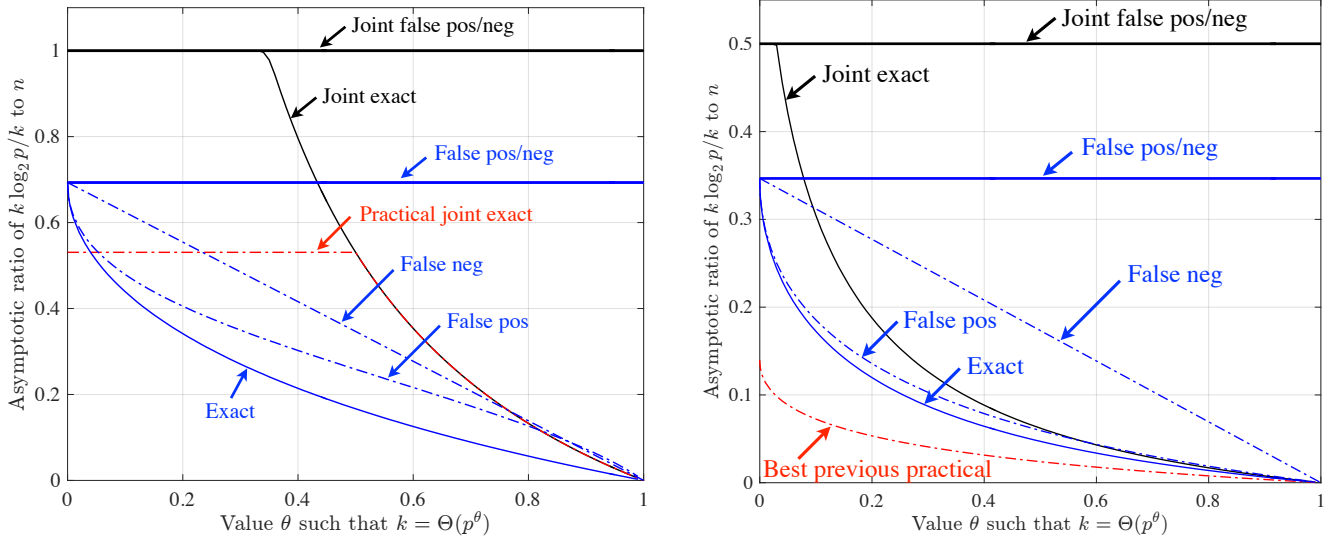


Figure 1: Asymptotic thresholds on the number of tests required for vanishing error probability in (Left) the noiseless setting, and (Right) the symmetric noise setting with $\rho = 0.11$. The number of defective items is $k = \Theta(p^\theta)$ for some $\theta > 0$. The vertical axis represents the constant $c(\theta)$ such that the number of tests is $(\frac{1}{c(\theta)} k \log_2 \frac{p}{k})(1 + o(1))$. The blue curves correspond to separate decoding of items, with exact recovery (Exact), $\Theta(k)$ false positives only (False pos), $\Theta(k)$ false negatives only (False neg), or both $\Theta(k)$ false positives and $\Theta(k)$ false negatives (False pos/neg). The “practical joint exact” curve is DD [12] or LP [18] in the noiseless case, and the “best previous practical” curve is NCOMP [14] in the noisy case.

Another related algorithm is the *column-based algorithm* of [26], which separately computes the number of agreements and disagreements with \mathbf{Y} for each column \mathbf{X}_j (though a final sorting step is also performed, so this is not a “separate decoding” algorithm according to our definition). The main focus in [26] is on the regime that $d_{\text{neg}} = 0$ and $d_{\text{pos}} = \Theta(p)$ (i.e., a very large number of false positives), which we do not consider in this paper.

Recently, algorithms based on *sparse graph codes* (with non-Bernoulli testing) have been proposed with various guarantees [15], [16]. For the noiseless non-adaptive case with exact recovery, however, the scaling laws are $O(k \log k \log p)$, thus failing to match the information-theoretic scaling $\Theta(k \log \frac{p}{k})$. Nevertheless, for partial recovery, a $O(k \log p)$ guarantee was proved in [16] (with loose constant factors). While these algorithms are suboptimal, they have the notable advantage of running in *sublinear time*.

We briefly mention that a rather different line of works has considered group testing with *adversarial* errors [27]–[29]. This is a fundamentally different setting to that of random errors, and the corresponding test designs and algorithms are less relevant to the present paper.

Separate decoding of items. The idea of separate decoding of items for sparse recovery problems (including group testing) was initiated by Maljutov and Mateev [8], who showed that when $k = O(1)$ and the decoder (9) is used with suitably-chosen γ , one can achieve exact recovery with vanishing error probability provided that

$$n \geq \frac{\log p}{I_1} (1 + o(1)), \quad (12)$$

where the *single-item mutual information* I_1 is defined as

follows, with implicit conditioning on item 1 being defective:

$$I_1 := I(X_1; Y). \quad (13)$$

As noted in [17], in the noiseless setting with $\nu = \ln 2$ we have $I_1 = \frac{(\log 2)^2}{k} (1 + o(1))$ as $k \rightarrow \infty$, and in this case, (12) matches (10) up to a factor of $\log 2 \approx 0.7$. In fact, the same turns out to be true for the symmetric noise model, matching (11) up to a factor of $\log 2$, or even better if ν is further optimized (see Appendix A). In more recent works [17], [30], similar results were shown when the rule (9) is replaced by a *universal* rule (i.e., not depending on the noise distribution) based on the empirical mutual information. However, we stick to (9) in this paper, as we found it to be more amenable to general scalings of the form $k = o(p)$.

Characterizations of the mutual information I_1 for several specific noise models were given in [31].

D. Contributions

In this paper, we provide an information-theoretic framework for studying separate decoding of items with general scalings of the form $k = o(p)$, as opposed to the case $k = O(1)$ considered in [8], [30]. As with joint decoding [13], [21], the regime $k = o(p)$ comes with significant challenges, with additional requirements of n arising from *concentration inequalities* and often dominating (12). In addition, we consider the novel aspect of partial recovery, as well as presenting a converse bound that is specific to separate decoding.

As mentioned above, our contributions for the noiseless and symmetric noise models are summarized in Figure 1, where we plot the asymptotic number of tests for achieving $P_e \rightarrow 0$ or $P_e(d_{\text{pos}}, d_{\text{neg}}) \rightarrow 0$ under Bernoulli testing with $\nu =$

$\log 2$. Note that in this figure, the number of allowed false positives and/or false negatives (if any) is always assumed to be $\Theta(k)$, though we allow for an arbitrarily small implied constant. Moreover, the horizontal line at the top of each plot also represents a converse for joint decoding with an arbitrarily small fraction of false positives and false negatives.

We make the following observations:

- In the noiseless case, our asymptotic bounds are within a factor $\log 2$ of the optimal threshold for joint decoding as $\theta \rightarrow 0$, are reasonable for all $\theta \in (0, 1)$ with improvements when false positives or false negatives are allowed, and are within a factor $\log 2$ of the optimal joint decoding threshold for all θ when both are allowed. Moreover, with exact recovery and $\theta < 0.0398$, we strictly improve on the best known bound for any efficient algorithm under Bernoulli testing.
- For the symmetric noise model, the general behavior is similar, but we significantly outperform the best known previous bound (NCOMP [14])³ for all $\theta \in (0, 1)$. Once again, when both false positives and false negatives are allowed, we are within a factor $\log 2$ of the optimal threshold for joint decoding.
- Although it is not shown in Figure 1, we provide a converse bound showing that if one tries to move beyond our achievability threshold for $\theta \rightarrow 0$ (or equivalently, the threshold obtained at all θ with both false positive and false negatives), then any separate decoding of items scheme with Bernoulli testing must have $\mathbb{E}[N_{\text{err}}] = k(1 - o(1))$, i.e., the average number of errors is close to the trivial value of k that would be obtained by declaring every item as non-defective. In contrast, below the same threshold, we show that $\mathbb{E}[N_{\text{err}}] = o(k)$ is achievable.

The exact recovery results are given in Section II, the partial recovery results are given in Section III, and the converse bounds are given in Section IV. In addition to these theoretical developments, we evaluate the performance of separate decoding via numerical experiments in Section V, showing it to perform well despite being outperformed by the LP relaxation technique of [18].

II. ACHIEVABILITY RESULTS WITH EXACT RECOVERY

In this section, we develop the theoretical results leading to the asymptotic bounds for the noiseless and noisy settings in Figure 1. To do this, we first establish non-asymptotic bounds on the error probability, then present the tools for performing an asymptotic analysis, and finally give the details of the applications to specific models.

³The final bound stated in [14, Thm. 6] appears to have a term omitted due an incorrect claim at the end of the proof stating that $\frac{\log k}{\log p}$ tends to zero (which is only true if $k = o(p^\theta)$ for all $\theta > 0$). Upon correcting this, the final bound increases from $n \approx 4.36(1 - 2\rho)^{-2} \log_2 p$ to $n \approx 4.36(1 + \sqrt{\theta})^2(1 - 2\rho)^{-2} \log_2 p$. This correction was also made in the follow-up paper [24], but that paper only considered the suboptimal choice $\nu = 1$, yielding a bound with worse constants. We also note that the bounds in [24] for an LP relaxation (using negative tests only) are strictly worse than those stated above for NCOMP.

A. Additional Notation

We define some further notation in addition to that in Section I-A. Our analysis will apply for any given choice of the defective set S , due to the symmetry of the observation model (3) and the i.i.d. test matrix \mathbf{X} . Hence, throughout this section we will focus on the specific set $S = \{1, \dots, k\}$. In particular, we assume that item 1 is defective, and we define $P_{Y|X_1}$ accordingly:

$$P_{Y|X_1}(y|x_1) = P_{Y|X_1, \beta_1}(y|x_1, 1). \quad (14)$$

Hence, the summation in (9) can be written as

$$i_1^n(\mathbf{X}_j, \mathbf{Y}) := \sum_{i=1}^n i(X_j^{(i)}, Y^{(i)}), \quad (15)$$

where

$$i_1(x_1, y) := \log \frac{P_{Y|X_1}(y|x_1)}{P_Y(y)}. \quad (16)$$

Following the terminology of the channel coding literature [32]–[34], we refer to this quantity as the *information density*. Denoting the distribution of a single entry of \mathbf{X} by $P_X \sim \text{Bernoulli}(\frac{\nu}{k})$, we find that the average of (16) with respect to $(X_1, Y) \sim P_X \times P_{Y|X_1}$ is the mutual information I_1 in (13). With the above distributions in place, we define $P_X^n(\mathbf{x}_1) = \prod_{i=1}^n P_X(x_1^{(i)})$, $P_Y^n(\mathbf{y}) = \prod_{i=1}^n P_Y(y^{(i)})$, and $P_{Y|X_1}^n(\mathbf{y}|\mathbf{x}_1) = \prod_{i=1}^n P_{Y|X_1}(y^{(i)}|x_1^{(i)})$.

When we specialize our results to the noiseless and symmetric noise models, we will choose

$$\nu = \nu_{\text{symm}} := \left\{ \text{unique value such that } \left(1 - \frac{\nu}{k}\right)^k = \frac{1}{2} \right\} \quad (17)$$

$$= (\log 2)(1 + o(1)). \quad (18)$$

For $k \rightarrow \infty$ (as we consider), there is essentially no difference between setting $\nu = \nu_{\text{symm}}$ or $\nu = \log 2$, but we found the former to be slightly more convenient mathematically. Either choice is known to be asymptotically optimal for maximizing I_1 in the noiseless model [17], [31]. Perhaps surprisingly, this is no longer true in general for the symmetric noise model (see Appendix A); however, it simplifies some of the analysis and does not impact the bounds significantly.

B. Initial Non-Asymptotic Bound

The following theorem provides an initial non-asymptotic upper bound on the error probability for general models. The result is proved using simple thresholding techniques that appeared in early studies of channel coding [35], [36], and have also been applied previously in the context of group testing [8], [13], [21].

Theorem 1. (Non-asymptotic, exact recovery) *For a general group testing model with Bernoulli($\frac{\nu}{k}$) testing and separate decoding of items according to (9), we have*

$$P_e \leq k\mathbb{P}[i_1^n(\mathbf{X}_1, \mathbf{Y}) \leq \gamma] + (p - k)e^{-\gamma}, \quad (19)$$

where $(\mathbf{X}_1, \mathbf{Y}) \sim P_X^n(\mathbf{x}_1)P_{Y|X_1}^n(\mathbf{y}|\mathbf{x}_1)$, and γ is given in (9).

Proof. For the exact recovery criterion, correct decoding requires the k defective items to pass the threshold test, and the $p - k$ non-defective items to fail the threshold test. Hence, by the union bound, we have

$$P_e \leq k\mathbb{P}[v_1^n(\mathbf{X}_1, \mathbf{Y}) \leq \gamma] + (p - k)\mathbb{P}[v_1^n(\bar{\mathbf{X}}_1, \mathbf{Y}) > \gamma], \quad (20)$$

where $(\mathbf{X}_1, \mathbf{Y}, \bar{\mathbf{X}}_1) \sim P_X^n(\bar{\mathbf{x}}_1)P_{Y|X_1}^n(\mathbf{y}|\bar{\mathbf{x}}_1)P_X^n(\bar{\mathbf{x}}_1)$. We bound the second term by writing

$$\mathbb{P}[v_1^n(\bar{\mathbf{X}}_1, \mathbf{Y}) > \gamma] = \sum_{\bar{\mathbf{x}}_1, \mathbf{y}} P_X^n(\bar{\mathbf{x}}_1)P_{Y|X_1}^n(\mathbf{y}|\bar{\mathbf{x}}_1) \mathbb{1}\left\{\log \frac{P_{Y|X_1}^n(\mathbf{y}|\bar{\mathbf{x}}_1)}{P_Y^n(\mathbf{y})} > \gamma\right\} \quad (21)$$

$$\leq \sum_{\bar{\mathbf{x}}_1, \mathbf{y}} P_X^n(\bar{\mathbf{x}}_1)P_{Y|X_1}^n(\mathbf{y}|\bar{\mathbf{x}}_1)e^{-\gamma} \quad (22)$$

$$= e^{-\gamma}, \quad (23)$$

where (21) bounds $P_Y^n(\mathbf{y})$ according to the event in the indicator function, and then bounding the indicator function by one. Combining (20) and (23) completes the proof. \square

C. Asymptotic Analysis

In order to apply Theorem 1, we need to characterize the probability appearing in the first term. The idea is to exploit the fact that $v_1^n(\mathbf{X}_1, \mathbf{Y})$ is an i.i.d. sum, and hence concentrates around its mean. While the following corollary is essentially a simple rewriting of Theorem 1, it makes the application of such concentration bounds more transparent. Here and subsequently, asymptotic notation such as $\rightarrow, o(\cdot), O(\cdot)$ is with respect to $p \rightarrow \infty$, and we assume that $k \rightarrow \infty$ with $k = o(p)$.

Theorem 2. (Asymptotic bound, exact recovery) *Under the setup of Theorem 1, suppose that the information density satisfies a concentration inequality of the following form:*

$$\mathbb{P}[v_1^n(\mathbf{X}_1, \mathbf{Y}) \leq nI_1(1 - \delta_2)] \leq \psi_n(\delta_2) \quad (24)$$

for some function $\psi_n(\delta_2)$. Moreover, suppose that the following conditions hold for some $\delta_1 \rightarrow 0$ and $\delta_2 > 0$:

$$n \geq \frac{\log\left(\frac{1}{\delta_1}(p - k)\right)}{I_1(1 - \delta_2)} \quad (25)$$

$$k \cdot \psi_n(\delta_2) \rightarrow 0. \quad (26)$$

Then $P_e \rightarrow 0$ under the decoder in (9) with $\gamma = \log \frac{p-k}{\delta_1}$.

Proof. Setting $\gamma = \log \frac{p-k}{\delta_1}$ in Theorem 1, we obtain

$$P_e \leq k\mathbb{P}\left[v_1^n(\mathbf{X}_1, \mathbf{Y}) \leq \log \frac{p-k}{\delta_1}\right] + \delta_1, \quad (27)$$

By the condition in (25), the probability in (27) is upper bounded by $\mathbb{P}[v_1^n(\mathbf{X}_1, \mathbf{Y}) \leq nI_1(1 - \delta_2)]$, which in turn is upper bounded by $\psi_n(\delta_2)$ by (24). We therefore have from (27) that $P_e \leq k\psi_n(\delta_2) + \delta_1$, and hence the theorem follows from the assumption $\delta_1 \rightarrow 0$ along with (26). \square

D. Concentration Bounds

In order to apply Theorem 2 to specific models, we need to characterize the concentration of $v_1^n(\mathbf{X}_1, \mathbf{Y})$ and attain an explicit expression for $\psi_n(\delta_2)$ therein. The following lemma brings us one step closer to attaining explicit expressions, giving a general concentration result based on Bernstein's inequality [37, Ch. 2].

Lemma 1. (Concentration via Bernstein's inequality) *Defining*

$$c_{\text{mean}} := k\mathbb{E}[v(X_1, Y)] = kI_1 \quad (28)$$

$$c_{\text{var}} := k\text{Var}[v(X_1, Y)] \quad (29)$$

$$c_{\text{max}} := \max_{x_1, y} |v(x_1, y)|, \quad (30)$$

we have for any $\delta_2 > 0$ that

$$\mathbb{P}\left[|v_1^n(\mathbf{X}_1, \mathbf{Y}) - nI_1| \leq n\delta_2\right] \leq 2 \exp\left(\frac{-\frac{1}{2} \cdot \frac{n}{k} \cdot c_{\text{mean}}^2 \delta_2^2}{c_{\text{var}} + \frac{1}{3} c_{\text{mean}} c_{\text{max}} \delta_2}\right) \quad (31)$$

Proof. Since $v_1^n(\mathbf{X}_1, \mathbf{Y})$ is a sum of i.i.d. random variables with mean I_1 , this lemma is a direct application of Bernstein's inequality for bounded random variables [37, Ch. 2], with the values $c_{\text{mean}}, c_{\text{var}}$ and c_{max} conveniently defined so that they behave as $\Theta(1)$ in typical examples (see Section II-E). \square

We will use Lemma 1 to establish the results shown for the symmetric noise model in Figure 1 (Right). While we could also use Lemma 1 for the noiseless model, it turns out that we can in fact do better via the following.

Lemma 2. (Concentration for noiseless model) *Under the noiseless model with $\nu = \nu_{\text{symm}}$ (cf., (17)), we have for any $\delta_2 \in (0, 1)$ that*

$$\mathbb{P}[v_1^n(\mathbf{X}_1, \mathbf{Y}) \leq nI_1(1 - \delta_2)] \leq \exp\left(-\frac{n \log 2}{k} \times \left((1 - \delta_2) \log(1 - \delta_2) + \delta_2\right)(1 + o(1))\right) \quad (32)$$

as $p \rightarrow \infty$ and $k \rightarrow \infty$ simultaneously.

Proof. We begin by characterizing at the various outcomes of (X_1, Y) and their probabilities, as well as the resulting values of $v_1(X_1, Y)$. Since $P_Y(0) = P_Y(1) = \frac{1}{2}$ by the definition of ν_{symm} , the information density simplifies to $v_1(x_1, y) = \log(2P_{Y|X_1}(y|x_1))$, and we have the following:

- $X_1 = 1$ with probability $\frac{\nu}{k}$, and in this case we deterministically have $Y = 1$, yielding $v_1 = \log 2$.
- $X_1 = 0$ with probability $1 - \frac{\nu}{k}$, and conditioned on this event we have the following:
 - $Y = 0$ with probability $(1 - \frac{\nu}{k})^{k-1} = \frac{1}{2} \cdot \frac{1}{1 - \nu/k} = \frac{1}{2} \cdot (1 + \frac{\nu/k}{1 - \nu/k}) = \frac{1}{2}(1 + o(1))$, where the first equality follows from the definition of ν_{symm} . Hence, in this case we have $v_1 = \log\left(1 + \frac{\nu/k}{1 - \nu/k}\right) = \frac{\nu}{k}(1 + o(1))$.
 - $Y = 1$ with probability $1 - \frac{1}{2} \cdot (1 + \frac{\nu/k}{1 - \nu/k}) = \frac{1}{2} \cdot (1 - \frac{\nu/k}{1 - \nu/k}) = \frac{1}{2}(1 + o(1))$, and in this case we have $v_1 = \log\left(1 - \frac{\nu/k}{1 - \nu/k}\right) = -\frac{\nu}{k}(1 + o(1))$.

From these calculations and (18), we immediately obtain

$$I_1 = \mathbb{E}[i(X_1, Y)] = \frac{(\log 2)^2}{k}(1 + o(1)). \quad (33)$$

Moreover, we see that we can write i_1^n (with the arguments $(\mathbf{X}_1, \mathbf{Y})$ implicit) as $i_1^n = \sum_{i=1}^n V^{(i)} + \sum_{i=1}^n W^{(i)}$, where the individual distributions of each $V^{(i)}$ and $W^{(i)}$ are as follows:

$$V = \begin{cases} \log 2 & \text{w.p. } \frac{\nu}{k} \\ 0 & \text{w.p. } 1 - \frac{\nu}{k} \end{cases} \quad (34)$$

$$W = \begin{cases} 0 & \text{w.p. } \frac{\nu}{k} \\ \log\left(1 + \frac{\nu/k}{1-\nu/k}\right) & \text{w.p. } \frac{1}{2}\left(1 + \frac{\nu/k}{1-\nu/k}\right) \\ \log\left(1 - \frac{\nu/k}{1-\nu/k}\right) & \text{w.p. } \frac{1}{2}\left(1 - \frac{\nu/k}{1-\nu/k}\right). \end{cases} \quad (35)$$

We proceed by fixing $\delta' > 0$ (later to be equated with $\delta_2(1 + o(1))$) and $\epsilon \in (0, \delta')$ (later to be taken to zero), and writing

$$\begin{aligned} & \mathbb{P}\left[i_1^n \leq \frac{n\nu \log 2}{k}(1 - \delta')\right] \\ & \leq \mathbb{P}\left[\sum_{i=1}^n V^{(i)} \leq \frac{n\nu \log 2}{k}(1 - \delta' + \epsilon)\right. \\ & \quad \left. \cup \sum_{i=1}^n W^{(i)} \leq \frac{-n\nu\epsilon \log 2}{k}\right] \end{aligned} \quad (36)$$

$$\begin{aligned} & \leq \mathbb{P}\left[\sum_{i=1}^n V^{(i)} \leq \frac{n\nu \log 2}{k}(1 - \delta' + \epsilon)\right] \\ & \quad + \mathbb{P}\left[\sum_{i=1}^n W^{(i)} \leq \frac{-n\nu\epsilon \log 2}{k}\right] \end{aligned} \quad (37)$$

$$=: T_1 + T_2, \quad (38)$$

where (36) follows since if both of the events on the right-hand side are violated then so is the event on the left-hand side, and (37) follows from the union bound.

To bound the term T_1 , we simply apply the multiplicative form of the Chernoff bound for Binomial random variables [38, Sec. 4.1] to obtain

$$T_1 \leq \exp\left(-\frac{n\nu}{k}\left((1 - \delta' + \epsilon)\log(1 - \delta' + \epsilon) + \delta' + \epsilon\right)\right). \quad (39)$$

As for T_2 , a direct calculation yields $\mathbb{E}[W] = o\left(\frac{1}{k^2}\right)$ and $\max[|W|] = O\left(\frac{1}{k}\right)$. Hence, by Hoeffding's inequality [37, Ch. 2], we have for any fixed $\epsilon > 0$ (not depending on p) that

$$T_2 \leq \exp(-nC\epsilon^2) \quad (40)$$

for some constant $C > 0$ and sufficiently large p . In particular, since $k \rightarrow \infty$, we have $T_2 = o(T_1)$.

Substituting the preceding bounds into (38), we obtain (32) with $1 - \delta' - \epsilon$ in place of δ_2 , and $\frac{\nu \log 2}{k}$ in place of I_1 . The lemma is concluded by noting that ϵ may be arbitrarily small, and noting from (18) and (33) that $\frac{\nu \log 2}{k} = I_1(1 + o(1))$. \square

E. Applications to Specific Models

Noiseless model: For the noiseless group testing model (cf., (1)), we immediately obtain the following from Theorem 2 and Lemma 2.

Corollary 1. (Noiseless, exact recovery) *For the noiseless group testing problem with $\nu = \nu_{\text{symm}}$ (cf., (17)) and $k = \Theta(p^\theta)$ for some $\theta \in (0, 1)$, we can achieve $P_e \rightarrow 0$ with separate decoding of items provided that*

$$n \geq \min_{\delta_2 > 0} \max \left\{ \frac{k \log p}{(\log 2)^2(1 - \delta_2)}, \frac{k \log k}{(\log 2)((1 - \delta_2)\log(1 - \delta_2) + \delta_2)} \right\} (1 + \eta) \quad (41)$$

for some $\eta > 0$.

Proof. We showed in the proof of Lemma 2 that $I_1 = \frac{(\log 2)^2}{k}(1 + o(1))$ (cf., (33)), and hence the first term in (41) follows from (25) with $\delta_1 \rightarrow 0$ sufficiently slowly. Moreover, by equating $\psi_n(\delta_2)$ with the right-hand side of (32) and performing simple rearranging, we find that the second term in (41) follows from (26). \square

Symmetric noise model: For the symmetric noisy model (cf., (1)), we make use of Lemma 1, with the constants c_{mean} , c_{var} and c_{max} therein characterized in the following. Recall that H_2 is the binary entropy function in nats.

Lemma 3. (Bernstein parameters for symmetric noise) *Under the symmetric noise model with a fixed parameter $\rho \in (0, \frac{1}{2})$ (not depending on p) and $\nu = \nu_{\text{symm}}$ (cf., (17)), we have*

$$k\mathbb{E}[i(X_1, Y)] = (\log 2)(\log 2 - H_2(\rho))(1 + o(1)) \quad (42)$$

$$\begin{aligned} k\text{Var}[i(X_1, Y)] & \leq (\log 2) \left((1 - \rho)\log^2(2(1 - \rho)) \right. \\ & \quad \left. + \rho\log^2(2\rho) \right) (1 + o(1)) \end{aligned} \quad (43)$$

$$\max_{x_1, y} |i(x_1, y)| = \log \frac{1}{2\rho} \quad (44)$$

as $p \rightarrow \infty$ and $k \rightarrow \infty$ simultaneously.

Proof. We begin by looking at the various possible outcomes of (X_1, Y) and their probabilities, as well as the resulting values of $i_1(X_1, Y)$. Since $P_Y(0) = P_Y(1) = \frac{1}{2}$ by the definition of ν_{symm} , the information density simplifies to $i_1(x_1, y) = \log(2P_{Y|X_1}(y|x_1))$, and we have the following:

- $X_1 = 1$ with probability $\frac{\nu}{k}$, and conditioned on this event we have the following:
 - $Y = 1$ with probability $1 - \rho$, and in this case we have $i_1 = \log(2(1 - \rho))$.
 - $Y = 0$ with probability ρ , and in this case we have $i_1 = \log(2\rho)$.
- $X_1 = 0$ with probability $1 - \frac{\nu}{k}$, and conditioned on this event we have the following:
 - $Y = 0$ with probability $(1 - \rho) \cdot \frac{1}{2}(1 + \xi) + \rho \cdot \frac{1}{2}(1 - \xi) = \frac{1}{2}(1 + (1 - 2\rho)\xi)$, where $\xi = \frac{\nu/k}{1-\nu/k}$ as derived following (32). Hence, in this case, we have $i_1 = \log(1 + (1 - 2\rho)\xi) = \frac{(1-2\rho)\nu}{k}(1 + o(1))$.
 - $Y = 1$ with probability $(1 - \rho) \cdot \frac{1}{2}(1 - \xi) + \rho \cdot \frac{1}{2}(1 + \xi) = \frac{1}{2}(1 - (1 - 2\rho)\xi)$, and in this case, we have $i_1 = \log(1 - (1 - 2\rho)\xi) = -\frac{(1-2\rho)\nu}{k}(1 + o(1))$.

With these computations in place, the lemma follows easily by evaluating the expectation $\mathbb{E}[z]$, variance $\mathbb{E}[(z - I_1)^2]$, and maximum $\max[z]$ directly, and substituting $\nu = \nu_{\text{symm}} = (\log 2)(1 + o(1))$. We briefly outline the details:

- For the mean, the contributions corresponding to $X_1 = 1$ already give the right-hand side of (42), whereas for $X_1 = 0$, the terms corresponding to $Y = 0$ and $Y = 1$ effectively cancel, i.e., their sum is $o(\frac{1}{k})$.
- For the variance, the contributions corresponding to $X_1 = 1$ already give the right-hand side of (43), whereas the contributions from the sub-cases of $X_1 = 0$ are $O(\frac{1}{k^2})$.
- For the maximum, we use the fact that $\log \frac{1}{2\rho} \geq \log(2(1 - \rho))$ for all $\rho \in (0, \frac{1}{2})$. \square

From this lemma, we immediately obtain the following.

Corollary 2. (Symmetric noise, exact recovery) *For noisy group testing with $\rho \in (0, \frac{1}{2})$ (not depending on p), $\nu = \nu_{\text{symm}}$, and $k = \Theta(p^\theta)$ for some $\theta \in (0, 1)$, we can achieve $P_e \rightarrow 0$ with separate decoding of items provided that*

$$n \geq \min_{\delta_2 > 0} \max \left\{ \frac{k \log p}{(\log 2)(\log 2 - H_2(\rho))(1 - \delta_2)}, \frac{(k \log k) \cdot (c_{\text{var}} + \frac{1}{3} c_{\text{mean}} c_{\text{max}} \delta_2)}{\frac{1}{2} \cdot c_{\text{mean}}^2 \delta_2^2} \right\} (1 + \eta) \quad (45)$$

for some $\eta > 0$, where c_{mean} , c_{var} , and c_{max} are respectively given by the right-hand sides of (42)–(44).

We have focused on the case $\nu = \nu_{\text{symm}}$ to simplify the analysis and establish an explicit $\log 2$ gap to joint decoding as $\theta \rightarrow 0$ (in which case the first term of (45) dominates for arbitrarily small δ_2). However, we show in Appendix A that this choice can be suboptimal even as $\theta \rightarrow 0$, and that more generally we obtain the sufficient condition $n \geq \frac{k \log p}{\nu D_2(\rho \| \rho^* e^{-\nu})} (1 + \eta)$ in this limit, where $a * b = a(1 - b) + b(1 - a)$, and $D_2(a \| b) = a \log \frac{a}{b} + (1 - a) \log \frac{1 - a}{1 - b}$ is the binary KL divergence function. This turns out to marginally improve on the condition $n \geq \frac{k \log p}{(\log 2)(\log 2 - H_2(\rho))} (1 + \eta)$ obtained via the specific choice $\nu = \nu_{\text{symm}} = (\log 2)(1 + o(1))$.

Other noise models: We showed above how to apply Lemma 1 to the symmetric noise model. However, it can also be applied more generally, yielding an analogous result for any model in which the quantities c_{mean} , c_{var} , and c_{max} in (28)–(30) behave as $\Theta(1)$. In particular, for any such model and any fixed $\nu > 0$, in the limit as $\theta \rightarrow 0$, it suffices to have

$$n \geq \frac{\log p}{I_1} (1 + \eta) = \frac{k \log p}{c_{\text{mean}}} (1 + \eta), \quad (46)$$

for arbitrarily small $\eta > 0$. In contrast, for θ strictly greater than zero, the conditions on n resulting from Bernstein's inequality may dominate (46), similarly to Corollary 2.

In the following section, we show that when we move to *partial recovery*, it is possible to circumvent the difficulties of the concentration bounds for $\theta > 0$, and to derive sufficient conditions of the form (46) valid for all $\theta \in (0, 1)$, even when the term $k \log p$ is improved to $k \log \frac{p}{k}$.

III. ACHIEVABILITY RESULTS WITH PARTIAL RECOVERY

In this section, we show that the analysis of the previous section can easily be adapted to provide achievability results when a certain number false positives d_{pos} and/or false negatives d_{neg} are allowed. We make use of the notation from Sections I-A and II-A.

The main tool we need is the following, whose proof is in fact implicit in our analysis for the exact recovery criterion.

Lemma 4. (Auxiliary result for partial recovery) *For any group testing model of the form (3), under the decoder in (9) with threshold $\gamma > 0$, we have the following:*

(i) *For any $j \notin S$, the probability of passing the threshold test is upper bounded by $e^{-\gamma}$.*

(ii) *Suppose that the information density satisfies a concentration inequality of the form (24) for some function $\psi_n(\delta_2)$, and that the number of tests satisfies $n \geq \frac{\gamma}{I_1(1 - \delta_2)}$. Then for any $j \in S$, the probability of failing the threshold test is upper bounded by $\psi_n(\delta_2)$.*

Proof. The first part was shown in (23), and the second part is contained in the proof of Theorem 2. \square

A. General Partial Recovery Achievability Results

We proceed by giving three variations of Theorem 2 for the case that way may tolerate false positives, false negatives, or both. We focus on the case that the number of false negatives and/or false positives is $\Theta(k)$, but we note that the implied constant in the $\Theta(\cdot)$ notation may be arbitrarily small.

We begin with the case that only false positives are allowed. This setting is closely related to that of *list decoding*, which was studied in [9], [26], [39], [40]. Recall that asymptotic notation such as \rightarrow , $o(\cdot)$, $O(\cdot)$ is with respect to $p \rightarrow \infty$, and we assume that $k \rightarrow \infty$ with $k = o(p)$.

Theorem 3. (Asymptotic bound, false positives only) *Consider the group testing problem with $d_{\text{pos}} = \alpha_{\text{pos}} k$ for some $\alpha_{\text{pos}} > 0$ (not depending on p) and $d_{\text{neg}} = 0$, and suppose that the information density satisfies a concentration inequality of the form (24) for some $\psi_n(\delta_2)$. Moreover, suppose that the following conditions hold for some $\delta_1 \rightarrow 0$ and $\delta_2 > 0$:*

$$n \geq \frac{\log\left(\frac{1}{\delta_1} \cdot \frac{p-k}{k}\right)}{I_1(1 - \delta_2)} \quad (47)$$

$$k \cdot \psi_n(\delta_2) \rightarrow 0. \quad (48)$$

Then under the decoder in (9) with $\gamma = \log\left(\frac{1}{\delta_1} \cdot \frac{p-k}{k}\right)$, we have $P_e(d_{\text{pos}}, 0) \rightarrow 0$.

Proof. The fact that the probability of one or more false negatives tends to zero follows by applying the second part of Lemma 4, and following the proof of (26) in Theorem 2.

Let N_{pos} denote the (random) number of false positives. Setting $\gamma = \log\left(\frac{1}{\delta_1} \cdot \frac{p-k}{k}\right)$ in the first part of Lemma 4, we obtain $\mathbb{E}[N_{\text{pos}}] \leq (p - k)e^{-\gamma} = k\delta_1$. By the assumption $\delta_1 \rightarrow 0$, it follows that $\mathbb{E}[N_{\text{pos}}] = o(k)$. Hence, by Markov's inequality, the probability of $N_{\text{pos}} > \alpha_{\text{pos}} k$ must vanish for any $\alpha_{\text{pos}} = \Theta(1)$. \square

The main difference in Theorem 3 compared to Theorem 1 is that $\log(p-k)$ is replaced by $\log \frac{p-k}{k}$ in the numerator. While this may not appear to be a drastic change, it can lead to visible improvements (cf., Figure 1), particularly for moderate to large values of θ .

Next, we consider the case that only false negatives are allowed, i.e., $d_{\text{pos}} = 0$ and $d_{\text{neg}} > 0$.

Theorem 4. (Asymptotic bound, false negatives only) *Consider the group testing problem with $d_{\text{neg}} = \alpha_{\text{neg}}k$ for some $\alpha_{\text{neg}} \in (0, 1)$ (not depending on p) and $d_{\text{pos}} = 0$, and suppose that the information density satisfies a concentration inequality of the form (24) for some $\psi_n(\delta_2)$. Moreover, suppose that the following conditions hold for some $\delta_1 \rightarrow 0$ and $\delta_2 > 0$:*

$$n \geq \frac{\log\left(\frac{1}{\delta_1} \cdot (p-k)\right)}{I_1(1-\delta_2)} \quad (49)$$

$$\psi_n(\delta_2) \rightarrow 0. \quad (50)$$

Then under the decoder in (9) with $\gamma = \log \frac{p-k}{\delta_1}$, we have $P_e(0, d_{\text{neg}}) \rightarrow 0$.

Proof. Applying the first part of Lemma 4 along with the union bound and the choice $\gamma = \log \frac{p-k}{\delta_1}$, we find that the probability of one or more false negatives is upper bounded by δ_1 , which vanishes by assumption.

Let N_{neg} denote the (random) number of false negatives. Setting $\gamma = \log\left(\frac{1}{\delta_1} \cdot \frac{p-k}{k}\right)$ in the second part of Lemma 4, we obtain $\mathbb{E}[N_{\text{neg}}] \leq k\psi_n(\delta_2)$. By the assumption $\psi_n(\delta_2) \rightarrow 0$, it follows that $\mathbb{E}[N_{\text{neg}}] = o(k)$. Hence, by Markov's inequality, $\mathbb{P}[N_{\text{neg}} > \alpha_{\text{neg}}k]$ must vanish for any $\alpha_{\text{neg}} = \Theta(1)$. \square

By allowing false negatives, we obtain a significantly milder condition on ψ_n in (50), corresponding to the concentration of $v_1^n(\mathbf{X}_1, \mathbf{Y})$. Specifically, all we need is concentration about the mean at an arbitrarily slow rate, whereas in Theorem 2 we needed a rate of $O\left(\frac{1}{k}\right)$. This turns out to significantly reduce the required number of tests for moderate to large values of θ ; see Corollary 3 below, as well as Figure 1.

Finally, we consider the case that both false positives and false negatives are allowed.

Theorem 5. (Asymptotic bound, false positives and false negatives) *Consider the group testing problem with $d_{\text{pos}} = \alpha_{\text{pos}}k$ and $d_{\text{neg}} = \alpha_{\text{neg}}k$ for some $\alpha_{\text{pos}} > 0$ and $\alpha_{\text{neg}} \in (0, 1)$ (not depending on p) and $d_{\text{pos}} = 0$, and suppose that the information density satisfies a concentration inequality of the form (24) for some function $\psi_n(\delta_2)$. Moreover, suppose that the following conditions hold for some $\delta_1 \rightarrow 0$ and $\delta_2 > 0$:*

$$n \geq \frac{\log\left(\frac{1}{\delta_1} \cdot \frac{p-k}{k}\right)}{I_1(1-\delta_2)} \quad (51)$$

$$\psi_n(\delta_2) \rightarrow 0. \quad (52)$$

Then under the decoder in (9) with $\gamma = \log \frac{p-k}{\delta_1}$, we have $P_e(d_{\text{pos}}, d_{\text{neg}}) \rightarrow 0$.

Proof. This result is directly deduced from the proofs of Theorems 3 and 4. \square

As we show in Corollary 5 below, this result leads to a broad class of noise models where the threshold $n^* = \frac{\log \frac{p}{k}}{I_1}(1+o(1))$ can be achieved for all $\theta \in (0, 1)$ with partial recovery.

B. Applications to Specific Models

Noiseless model: The following corollary gives three variations of the result in Corollary 1 corresponding to the three partial recovery settings considered in the previous subsection.

Corollary 3. (Noiseless model, partial recovery) *For the noiseless group testing problem with $\nu = \nu_{\text{symm}}$ (cf., (17)) and $k = \Theta(p^\theta)$ for some $\theta \in (0, 1)$, we can achieve $P_e(d_{\text{pos}}, d_{\text{neg}}) \rightarrow 0$ with separate decoding of items under any of the following conditions:*

(i) $d_{\text{pos}} = \Theta(k)$, $d_{\text{neg}} = 0$, and

$$n \geq \min_{\delta_2 > 0} \max \left\{ \frac{k \log \frac{p}{k}}{(\log 2)^2(1-\delta_2)}, \frac{k \log k}{(\log 2)((1-\delta_2) \log(1-\delta_2) + \delta_2)} \right\} (1+\eta) \quad (53)$$

for some $\eta > 0$.

(ii) $d_{\text{pos}} = 0$, $d_{\text{neg}} = \Theta(k)$, and

$$n \geq \frac{k \log p}{(\log 2)^2} (1+\eta) \quad (54)$$

for some $\eta > 0$.

(iii) $d_{\text{pos}} = \Theta(k)$, $d_{\text{neg}} = \Theta(k)$, and

$$n \geq \frac{k \log \frac{p}{k}}{(\log 2)^2} (1+\eta) \quad (55)$$

for some $\eta > 0$.

Proof. These results follow in the same way as Corollary 2, but with Theorems 3–5 in place of Theorem 2. When false negatives are allowed, δ_2 can be arbitrarily small, and is hence factored into the $1+\eta$ remainder term. \square

The above bounds are depicted visually and compared to existing bounds in Figure 1 (Left), and were discussed in Section I-D.

Symmetric noise model: Next, we provide an analog of Corollary 2 for partial recovery.

Corollary 4. (Symmetric noise model, partial recovery) *Under the symmetric noise model with a fixed parameter $\rho \in (0, \frac{1}{2})$ (not depending on p) and $\nu = \nu_{\text{symm}}$ (cf., (17)), we can achieve we can achieve $P_e(d_{\text{pos}}, d_{\text{neg}}) \rightarrow 0$ with separate decoding of items under any of the following conditions:*

(i) $d_{\text{pos}} = \Theta(k)$, $d_{\text{neg}} = 0$, and

$$n \geq \min_{\delta_2 > 0} \max \left\{ \frac{k \log \frac{p}{k}}{(\log 2)(\log 2 - H_2(\rho))(1-\delta_2)}, \frac{(k \log k) \cdot (c_{\text{var}} + \frac{1}{3}c_{\text{mean}}c_{\text{max}}\delta_2)}{\frac{1}{2} \cdot c_{\text{mean}}^2 \delta_2^2} \right\} (1+\eta) \quad (56)$$

for some $\eta > 0$, where c_{mean} , c_{var} , and c_{max} are respectively given by the right-hand sides of (42)–(44).

(ii) $d_{\text{pos}} = 0$, $d_{\text{neg}} = \Theta(k)$, and

$$n \geq \frac{k \log p}{(\log 2)(\log 2 - H_2(\rho))} (1+\eta) \quad (57)$$

for some $\eta > 0$.

(iii) $d_{\text{pos}} = \Theta(k)$, $d_{\text{neg}} = \Theta(k)$, and

$$n \geq \frac{k \log \frac{p}{k}}{(\log 2)(\log 2 - H_2(\rho))} (1 + \eta) \quad (58)$$

for some $\eta > 0$.

Proof. These bounds follow in the same way as Corollary 2, but with Theorems 3–5 in place of Theorem 2. \square

The above bounds are depicted visually and compared to existing bounds in Figure 1 (Right), and were discussed in Section I-D. Similarly to the discussion following Corollary 2, slightly improved bounds can be obtained by considering $\nu \neq \nu_{\text{symm}}$; for instance, (58) can be generalized to $n \geq \frac{k \log \frac{p}{k}}{\nu D_2(\rho \| \rho^* e^{-\nu})} (1 + \eta)$, where $a \star b = a(1 - b) + b(1 - a)$, and $D_2(a \| b) = a \log \frac{a}{b} + (1 - a) \log \frac{1 - a}{1 - b}$ (see Appendix A).

Other noise models: Analogous bounds to those in Corollary 4 can be obtained in an identical manner for any noise model such that the quantities c_{mean} , c_{var} , and c_{max} in (28)–(30) behave as $\Theta(1)$. To avoid repetition, we state this formally only for the case of both false positives and false negatives.

Corollary 5. (General noise models, partial recovery) *For any group testing model such that the quantities c_{mean} , c_{var} , and c_{max} in (28)–(30) behave as $\Theta(1)$, we can achieve $P_e(d_{\text{pos}}, d_{\text{neg}}) \rightarrow 0$ with separate decoding provided that $d_{\text{pos}} = \Theta(k)$, $d_{\text{neg}} = \Theta(k)$, and*

$$n \geq \frac{\log \frac{p}{k}}{I_1} (1 + \eta) = \frac{k \log \frac{p}{k}}{c_{\text{mean}}} (1 + \eta) \quad (59)$$

for some $\eta > 0$.

As mentioned in Section I-C, characterizations of c_{mean} (or equivalently, I_1) were given for several noise models in [31].

IV. CONVERSE RESULTS

In this section, we present lower bounds on the required number of tests to meet certain recovery criteria with separate decoding of items and Bernoulli testing. Specifically, these bounds apply to any decoders of the form (8), with functions ϕ_j that may differ from (9). Throughout the section, we again make use of the notation from Sections I-A and II-A.

For technical reasons, it will prove most convenient to consider converse bounds with respect to the *average* number of errors $\mathbb{E}[N_{\text{err}}] = \sum_{j=1}^p P_{e,j}$ (cf., (7)), where we define

$$P_{e,j} := \mathbb{P}[\hat{\beta}_j \neq \beta_j], \quad (60)$$

and where the probability is with respect to β , \mathbf{X} , and \mathbf{Y} . While this is different to the criteria in (5) and (6), our analysis in Section II for exact recovery can be interpreted as proving $\mathbb{E}[N_{\text{err}}] = o(1)$ and then using $\mathbb{P}[N_{\text{err}} \neq 0] \leq \mathbb{E}[N_{\text{err}}]$ (i.e., the first moment method). Similarly, Theorem 5 in Section III is based on showing that $\mathbb{E}[N_{\text{err}}] = o(k)$. Our main converse result, Theorem 7, shows (under mild concentration conditions) that if n is slightly below the threshold (51) in Theorem 5, then we must have $\mathbb{E}[N_{\text{err}}] \geq k(1 - o(1))$ (i.e., the number of errors is close to the trivial value of k that would be attained by declaring every item to be non-defective). Hence,

these results collectively indicate a *phase transition*: For n slightly below the threshold we have $\frac{\mathbb{E}[N_{\text{err}}]}{k} \rightarrow 1$, but for n slightly above that threshold we have $\frac{\mathbb{E}[N_{\text{err}}]}{k} \rightarrow 0$.

We emphasize that the converse bounds in this section are proved for Bernoulli testing, as opposed to arbitrary test matrices (e.g., see [41]). Converse bounds for arbitrary matrices appear to be difficult, and are left for future work.

A. Initial Non-Asymptotic Bound

We begin with a non-asymptotic lower bound on $P_{e,j}$, which holds for any given $j \in \{1, \dots, p\}$; this is due to the symmetry of the observations (3) and the test matrix \mathbf{X} . The proof is based on thresholding techniques that have been used extensively in channel coding [32], [33], though since the distribution of β_j is non-uniform here, the analysis is in fact more akin to *joint source-channel coding* [42].

Theorem 6. (Non-asymptotic converse for single item) *For a general group testing model with Bernoulli testing and any separate decoding of items rule of the form (8), we have*

$$P_{e,j} \geq \frac{k}{p} \mathbb{P} \left[v_1^n(\mathbf{X}_1, \mathbf{Y}) \leq \log \frac{p - k}{k} \right], \quad (61)$$

where $(\mathbf{X}_1, \mathbf{Y}) \sim P_X^n(\mathbf{x}_1) P_{Y|X_1}^n(\mathbf{y}|\mathbf{x}_1)$.

Proof. Any given item is defective with probability $\frac{k}{p}$. Hence, we let $P_{\beta_j}(b_j)$ be the Bernoulli($\frac{k}{p}$) probability mass function. Moreover, we let $P_{Y|X_j, \beta_j}^n(\cdot | \cdot, b_j)$ be the n -th product of $P_{Y|X_j, \beta_j}(\cdot | \cdot, b_j)$, with the latter defined following (9).

By the law of total probability, under any event \mathcal{A} ,

$$P_{e,j} \geq \mathbb{P}[\mathcal{A}] - \mathbb{P}[\mathcal{A} \cap \text{no error}], \quad (62)$$

where “no error” refers only to item j . Choosing $\mathcal{A} = \left\{ \log \frac{P_{\beta_j}(\beta_j) P_{Y|X_j, \beta_j}^n(\mathbf{Y}|\mathbf{X}_j, \beta_j)}{P_Y^n(\mathbf{Y})} \leq \gamma \right\}$ for some $\gamma > 0$, we can bound the second term in (62) as follows:

$$\begin{aligned} & \mathbb{P}[\mathcal{A} \cap \text{no error}] \\ &= \sum_{b_j} \sum_{(\mathbf{x}_j, \mathbf{y}) \in \mathcal{D}(b_j)} P_{\beta_j}(b_j) P_X^n(\mathbf{x}_j) P_{Y|X_j, \beta_j}^n(\mathbf{y}|\mathbf{x}_j, b_j) \\ & \quad \times \mathbb{1} \left\{ \log \frac{P_{\beta_j}(b_j) P_{Y|X_j, \beta_j}^n(\mathbf{y}|\mathbf{x}_j, b_j)}{P_Y^n(\mathbf{y})} \leq \gamma \right\} \end{aligned} \quad (63)$$

$$\leq \sum_{b_j} \sum_{(\mathbf{x}_j, \mathbf{y}) \in \mathcal{D}(b_j)} P_X^n(\mathbf{x}_j) P_Y^n(\mathbf{y}) e^\gamma \quad (64)$$

$$= e^\gamma, \quad (65)$$

where

- (63) follows by defining $\mathcal{D}(b_j)$ to be the set of $(\mathbf{x}_j, \mathbf{y})$ pairs that are decoded to b_j by the decoder of item j ;
- (64) follows by upper bounding $P_{\beta_j}(b_j) P_{Y|X_j, \beta_j}^n(\mathbf{y}|\mathbf{x}_j, b_j)$ according to the event in the indicator function, and then upper bounding the indicator function by one;
- (65) follows since by definition, the union of the sets $\mathcal{D}(b_j)$ over all b_j is precisely the set of all $(\mathbf{x}_j, \mathbf{y})$ pairs.

Substituting (65) and the choice of \mathcal{A} into (62), we obtain

$$\begin{aligned} P_{e,j} &\geq \mathbb{P} \left[\log \frac{P_{\beta_j}(\beta_j) P_{Y|X_j, \beta_j}^n(\mathbf{Y}|\mathbf{X}_j, \beta_j)}{P_Y^n(\mathbf{Y})} \leq \gamma \right] - e^\gamma \quad (66) \\ &= \frac{k}{p} \mathbb{P} \left[\log \frac{P_{Y|X_j, \beta_j}^n(\mathbf{Y}|\mathbf{X}_j, 1)}{P_Y^n(\mathbf{Y})} \leq \log \frac{p}{k} + \gamma \mid \beta_j = 1 \right] \\ &\quad + \left(1 - \frac{k}{p}\right) \mathbb{P} \left[\log \left(1 - \frac{k}{p}\right) \leq \gamma \right] - e^\gamma, \quad (67) \end{aligned}$$

where we have applied $\beta_j \sim \text{Bernoulli}(\frac{k}{p})$ and used $\log \frac{P_{Y|X_j, \beta_j}^n(\mathbf{Y}|\mathbf{X}_j, 0)}{P_Y^n(\mathbf{Y})} = 0$. The latter claim follows since conditioned on $\beta_j = 0$, the output is independent of \mathbf{X}_j and is randomly generated according to k different columns of \mathbf{X} , meaning $P_{Y|X_j, \beta_j}^n(\mathbf{Y}|\mathbf{X}_j, 0) = P_Y^n(\mathbf{Y})$.

Setting $\gamma = \log \left(1 - \frac{k}{p}\right)$ in (67), the second term becomes $1 - \frac{k}{p}$, and the third term exactly cancels out the second term. Finally, the first term equals the right-hand side of (61), since by definition $(\mathbf{X}_1, \mathbf{Y})$ in (61) has the same distribution as $(\mathbf{X}_j, \mathbf{Y})$ conditioned on $\beta_j = 1$ in (67). \square

B. Asymptotic Analysis

In order to apply Theorem 6, we need to characterize the probability appearing in the first term. Similarly to the achievability analysis in Section II, the idea is to exploit the fact that $v_1^n(\mathbf{X}_1, \mathbf{Y})$ is an i.i.d. sum, and hence concentrates around its mean.

Theorem 7. (Converse for average number of errors) *Under the setup of Theorem 6, suppose that the information density satisfies a concentration inequality of the following form:*

$$\mathbb{P}[v_1^n(\mathbf{X}_1, \mathbf{Y}) \geq nI_1(1 + \delta_2)] \leq \psi'_n(\delta_2) \quad (68)$$

for some function $\psi'_n(\delta_2)$. Moreover, suppose that the following conditions hold for some $\delta_2 > 0$:

$$n \leq \frac{\log \frac{p-k}{k}}{I_1(1 + \delta_2)} \quad (69)$$

$$\psi'_n(\delta_2) \rightarrow 0. \quad (70)$$

Then any decoder based on separate decoding of items must have a number of errors N_{err} satisfying $\mathbb{E}[N_{\text{err}}] \geq k(1 - o(1))$.

Proof. Substituting (69) into (61) gives $P_{e,j} \geq \frac{k}{p} \mathbb{P}[v_1^n(\mathbf{X}_1, \mathbf{Y}) \leq nI_1(1 + \delta_2)]$, which is lower bounded by $\frac{k}{p}(1 - \psi'_n(\delta_2))$ by the definition of ψ'_n . Hence, we have from (70) that $P_{e,j} \geq \frac{k}{p}(1 - o(1))$, and the theorem follows since $\mathbb{E}[N_{\text{err}}] = \sum_{j=1}^p P_{e,j}$ by definition. \square

C. Application to Specific Models

The following corollary applies Theorem 7 to the noiseless setting, the symmetric noise model, and more general noise models where Bernstein's inequality can be applied.

Corollary 6. (Converses for specific models) *For the group testing problem under Bernoulli testing with parameter $\nu > 0$, and $k = \Theta(p^\theta)$ for some $\theta \in (0, 1)$, any decoding rule based on separate decoding of items must satisfy $\mathbb{E}[N_{\text{err}}] \geq k(1 - o(1))$ in any of the following settings:*

(i) *The tests are noiseless, and the number of tests satisfies*

$$n \leq \frac{k \log \frac{p}{k}}{(\log 2)^2} (1 - \eta) \quad (71)$$

for arbitrarily small $\eta > 0$.

(ii) *The noise is symmetric with parameter $\rho \in (0, \frac{1}{2})$ (not depending on p), and the number of tests satisfies*

$$n \leq \frac{k \log \frac{p}{k}}{\nu D_2(\rho \| \rho \star e^{-\nu})} (1 - \eta) \quad (72)$$

for arbitrarily small $\eta > 0$, where $a \star b = a(1 - b) + b(1 - a)$, and $D_2(a \| b) = a \log \frac{a}{b} + (1 - a) \log \frac{1-a}{1-b}$.

(iii) *The noise follows an arbitrary distribution such that c_{mean} , c_{var} , and c_{max} in (28)–(30) all behave as $\Theta(1)$, and the number of tests satisfies*

$$n \leq \frac{\log \frac{p}{k}}{I_1} (1 - \eta) = \frac{k \log \frac{p}{k}}{c_{\text{mean}}} (1 - \eta) \quad (73)$$

for arbitrarily small $\eta > 0$.

Proof. Without loss of generality, we assume that the upper bounds on n in the corollary statement hold with equality, since the decoders could always choose to ignore some tests. Once n is set in this way, we simply apply Theorem 7, using Lemma 1 with $c_{\text{mean}} = \Theta(1)$, $c_{\text{var}} = \Theta(1)$, and $c_{\text{max}} = \Theta(1)$ to establish that $\psi'_n(\delta_2) \rightarrow 1$ for arbitrarily small δ_2 . For part (i), we use the fact that $\nu = \nu_{\text{symm}}$ (cf., (17)) is asymptotically optimal in the noiseless setting [17], [31], and use the corresponding characterization of I_1 in (33). For part (ii), we use the fact that $I_1 = (\frac{\nu}{k} D_2(\rho \| \rho \star e^{-\nu}))(1 + o(1))$, as shown in Appendix A. \square

Note that when $\nu = \log 2$ or $\nu = \nu_{\text{symm}}$ in (72), the denominator simplifies to $(\log 2)(\log 2 - H_2(\rho))(1 + o(1))$, thus matching the achievability bounds in the previous sections. Conversely, as discussed in Appendix A, the achievability arguments can be adapted to match (72) more generally.

V. NUMERICAL EXPERIMENTS

In this section, we complement our theoretical findings with numerical experiments, comparing separate decoding of items against state-of-the-art alternatives in both the noiseless and noisy settings. We will see that although the linear programming (LP) methods of [18] tend to be superior, the experiments are consistent with the fact that separate decoding of items is optimal up to small constant factors.

In all experiments below, we let \mathbf{X} be an i.i.d. Bernoulli matrix with parameter $\nu = \log 2$, and we estimate the success probabilities by averaging over 500 independent trials. All LP based methods are solved using Gurobi [43]. For separate decoding of items as in (9), in accordance with our theoretical analysis, we choose

$$\gamma = nI_1(1 - \delta) \quad (74)$$

for some $\delta \in (0, 1)$. Specifically, we let $\delta = 0.5$, as we found this to work well in the examples considered below.

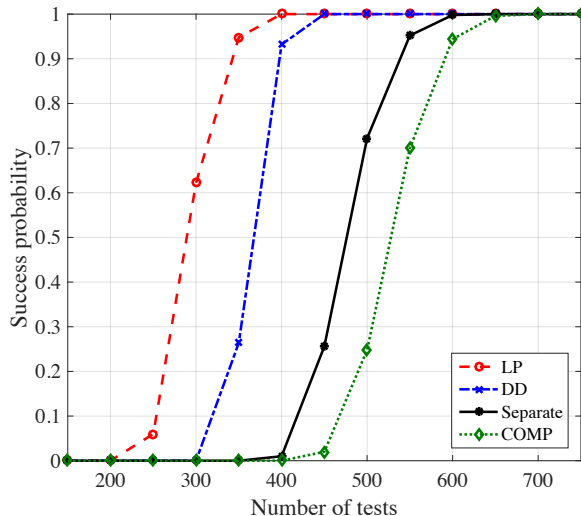


Figure 2: Experimental evaluation of practical algorithms for noiseless group testing, with $p = 3000$ items and $k = 30$ defectives, allowing up to $d_{\text{pos}} = 3$ false positives and $d_{\text{neg}} = 3$ false negatives.

A. Noiseless Setting

In the noiseless setting, we compare against the following existing algorithms:

- *Combinatorial Orthogonal Matching Pursuit* (COMP) [14], which declares an item as negative if and only if it appears in some negative test;
- *Definite Defectives* (DD) [12], which uses COMP to construct a set of *positive defectives*, and declares an item to be defective if and only if there exists a positive test in which it is the unique possibly defective item.
- The *linear programming* (LP) relaxation of [18], with non-integer returned values of the defective status rounded to $\{0, 1\}$.

Figure 2 plots the success probability as a function of the number of tests with $p = 3000$, $k = 30$, and $d_{\text{pos}} = d_{\text{neg}} = 3$. We observe that separate decoding slightly outperforms COMP, whereas both DD and LP provide better performance, with LP performing best. This is consistent with the fact that DD and LP have better theoretical guarantees in the noiseless setting for most $\theta < \frac{1}{2}$ (cf., Figure 1). Despite this, separate decoding does provide good performance, with the gap to LP amounting to a relatively small constant factor.

B. Symmetric Noise Setting

In the symmetric noise setting, we compare against the following existing algorithms:

- The *Noisy Combinatorial Orthogonal Matching Pursuit* (NCOMP) algorithm declares the j -th item to be defective if and only if $\frac{\sum_{i=1}^n \mathbb{1}\{X_j^{(i)}=1 \cap Y^{(i)}=1\}}{\sum_{i=1}^n \mathbb{1}\{X_j^{(i)}=1\}} \geq 1 - \rho(1 + \Delta)$, for some threshold $\Delta > 0$. We set $\Delta = 1.5$ based on manual tuning.
- The (noisy) *Linear Programming* (LP) of [18], which introduces slack variables indicating tests where errors

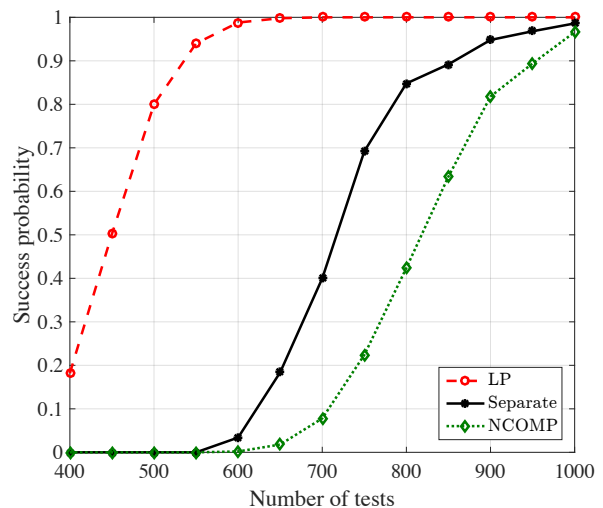


Figure 3: Experimental evaluation of practical algorithms for symmetric noisy group testing, with $p = 3000$ items, $k = 30$ defectives, and crossover probability $\rho = 0.05$, allowing up to $d_{\text{pos}} = 3$ false positives and $d_{\text{neg}} = 3$ false negatives.

occurred. We let the regularization parameter in this formulation be $\lambda = 0.5$, based on manual tuning.

Figure 3 plots the success probability as a function of the number of tests with $p = 3000$, $k = 30$, $d_{\text{pos}} = d_{\text{neg}} = 3$, and $\rho = 0.05$. We observe that separate decoding slightly outperforms NCOMP, with LP again providing the best performance, and the differences amounting to constant factors.

C. Discussion

In the above examples, we see that the LP method of [18] provides state-of-the-art performance. However, we recall from Section I-C that LP methods currently only have theoretical guarantees in the noiseless setting. Moreover, it is non-trivial to develop large-scale distributed algorithms, as was done for separate decoding of items in [44], [45]. These limitations present important directions for future research.

It may appear to be slightly surprising that the performance of NCOMP is close to the separate decoding rule of (9), since the former only depends on a small fraction of the tests (namely, those with $X_j^{(i)} = 1$). The intuition is that although most tests do not contain item j (i.e., $X_j^{(i)} = 0$), those that do contain item j are considerably more informative for determining whether j is defective. Hence, focusing only on these tests does not degrade the performance too much.

VI. CONCLUSION

We have provided an information-theoretic framework for studying group testing under the separate decoding of items method proposed in [8]. Our bounds are the best known for any practical group testing algorithm in several cases of interest, including (i) the noiseless model with partial recovery and θ sufficiently close to one, and (ii) the symmetric noise model with exact recovery or partial recovery, and any $\theta \in (0, 1)$. Overall, the results of this paper establish separate decoding of

items as a technique that is not only computationally efficient, but also comes with near-optimal theoretical guarantees.

An interesting direction for future work is to extend the achievability bounds to the case that the empirical mutual information based decoder [30] is used, and hence knowledge of the noise model is not required. Moreover, a key open challenge regarding the converse bounds of Section IV is to move beyond i.i.d. testing towards arbitrary test matrices.

APPENDIX

A. Optimizing ν for the Symmetric Noise Model

In this section, we characterize the mutual information I_1 in (13) for the symmetric noise model described in (4), with a given crossover probability $\rho \in (0, \frac{1}{2})$. In contrast to the noiseless model, we will see that the choices $\nu = \log 2$ and $\nu = \nu_{\text{symm}}$ (cf., (17)) do not always maximize I_1 .

Throughout this section, asymptotic notation is with respect to the limit as $k \rightarrow \infty$. We consider the case that the parameter ν (such that \mathbf{X} is i.i.d. on Bernoulli($\frac{\nu}{k}$)) behaves as $\Theta(1)$, since the regimes $\nu \rightarrow 0$ and $\nu \rightarrow \infty$ are easily shown to be strictly suboptimal regardless of the decoding rule.

Preliminary calculations. We have the following under the symmetric noise model (4):

$$\mathbb{P}[Y = 0] = (1 - \rho) \left(1 - \frac{\nu}{k}\right)^k + \rho \left(1 - \left(1 - \frac{\nu}{k}\right)^k\right) \quad (75)$$

$$= ((1 - \rho)e^{-\nu} + \rho(1 - e^{-\nu}))(1 + o(1)). \quad (76)$$

In addition, when item 1 is defective, we have

$$\begin{aligned} \mathbb{P}[Y = 0|X_1 = 0] \\ = (1 - \rho) \left(1 - \frac{\nu}{k}\right)^{k-1} + \rho \left(1 - \left(1 - \frac{\nu}{k}\right)^{k-1}\right) \end{aligned} \quad (77)$$

$$\begin{aligned} = (1 - \rho) \left(1 - \frac{\nu}{k}\right)^k \cdot \frac{1}{1 - \frac{\nu}{k}} \\ + \rho \left(1 - \left(1 - \frac{\nu}{k}\right)^k \cdot \frac{1}{1 - \frac{\nu}{k}}\right) \end{aligned} \quad (78)$$

$$\begin{aligned} = (1 - \rho) \left(1 - \frac{\nu}{k}\right)^k \cdot \left(1 + \frac{\nu}{k} + O\left(\frac{1}{k^2}\right)\right) \\ + \rho \left(1 - \left(1 - \frac{\nu}{k}\right)^k \cdot \left(1 + \frac{\nu}{k} + O\left(\frac{1}{k^2}\right)\right)\right) \end{aligned} \quad (79)$$

$$= \underbrace{\mathbb{P}[Y = 0]}_{=: \zeta} + \underbrace{\frac{\nu}{k} \left((1 - 2\rho) \left(1 - \frac{\nu}{k}\right)^k \right)}_{=: \Delta} + O\left(\frac{1}{k^2}\right), \quad (80)$$

where we have applied a Taylor expansion to $\frac{1}{1 - \frac{\nu}{k}}$, as well as substituting (75).

Characterizing I_1 . Using the above definitions of ζ and Δ , we simplify the mutual information I_1 in (13) as follows:

$$I_1 = H(Y) - H(Y|X_1) \quad (81)$$

$$= H_2(\zeta) - \frac{\nu}{k} H(Y|X_1 = 1) - \left(1 - \frac{\nu}{k}\right) H(Y|X_1 = 0) \quad (82)$$

$$= H_2(\zeta) - \frac{\nu}{k} H_2(\rho) - \left(1 - \frac{\nu}{k}\right) H_2(\zeta + \Delta), \quad (83)$$

where (83) follows since given $X_1 = 1$, the only uncertainty in Y is the additive noise, whereas given $X_1 = 0$, the probability that $Y = 0$ is $\zeta + \Delta$.

The derivative of $H_2(\alpha)$ is $\log \frac{1-\alpha}{\alpha}$, and hence, applying a first-order Taylor expansion in (83) yields

$$I_1 = \frac{\nu}{k} (H_2(\zeta) - H_2(\rho)) - \left(1 - \frac{\nu}{k}\right) \Delta \log \frac{1-\zeta}{\zeta} + O(\Delta^2). \quad (84)$$

Special cases $\nu = \log 2$ and $\nu = \nu_{\text{symm}}$. Setting $\nu = \log 2$, we obtain from (76) that $\zeta \rightarrow \frac{1}{2}$. Combined with the fact that $\Delta = O(\frac{1}{k})$, we deduce from (84) that

$$I_1 = \frac{\log 2}{k} (\log 2 - H_2(\rho))(1 + o(1)). \quad (85)$$

Here we have used $H_2(\frac{1}{2}) = \log 2$ and the continuity of entropy. Since $\nu_{\text{symm}} \rightarrow \log 2$ (cf., (18)), we deduce that (85) also holds when $\nu = \nu_{\text{symm}}$.

While (85) can be used to establish the optimality of separate decoding of inputs to within a factor of $\log 2$ in certain cases, it turns out that we can in fact do better via different choices of ν .

Improvements for general ν . Observing from (75) and (80) that $\Delta = \frac{\nu}{k} (\zeta - \rho) + O(\frac{1}{k^2})$, we deduce from (84) that

$$I_1 = \frac{\nu}{k} (H_2(\zeta) - H_2(\rho)) - \frac{\nu}{k} (\zeta - \rho) \log \frac{1-\zeta}{\zeta} + O\left(\frac{1}{k^2}\right) \quad (86)$$

$$= \underbrace{\frac{\nu}{k} \left(H_2(\zeta) - (\zeta - \rho) \log \frac{1-\zeta}{\zeta} - H_2(\rho) \right)}_{=: I'_1} + O\left(\frac{1}{k^2}\right). \quad (87)$$

We proceed by simplifying the leading term I'_1 :

$$\begin{aligned} I'_1 = \frac{\nu}{k} \left(-\zeta \log \zeta - (1 - \zeta) \log(1 - \zeta) \right. \\ \left. - (\zeta - \rho) \log(1 - \zeta) + (\zeta - \rho) \log \zeta - H_2(\rho) \right) \end{aligned} \quad (88)$$

$$= \frac{\nu}{k} \left(-(1 - \rho) \log(1 - \zeta) - \rho \log \zeta \right. \\ \left. + (1 - \rho) \log(1 - \rho) + \rho \log \rho \right) \quad (89)$$

$$= \frac{\nu}{k} D_2(\rho \| \xi), \quad (90)$$

where $D_2(a \| b) = a \log \frac{a}{b} + (1 - a) \log \frac{1-a}{1-b}$ is the binary KL divergence function.

Combining (90) with (76) and (87), we obtain

$$I_1 = \left(\frac{\nu}{k} D_2(\rho \| \rho \star e^{-\nu}) \right) (1 + o(1)), \quad (91)$$

where $a \star b = a(1 - b) + b(1 - a)$. A simple numerical computation reveals that when $\rho > 0$ (e.g., $\rho = 0.11$), (91) need not be optimized at $\nu = \log 2$; however, we found the gap to be quite small, and thus focused on the latter choice.

ACKNOWLEDGMENT

The authors would like to thank Dmitry Malioutov for sharing his code from [18], and Mikhail Malyutov for pointing us to [44], [45].

This work was supported in part by the European Commission under Grant ERC Future Proof, SNF 200021-146750 and SNF CRSII2-147633, and the ‘EPFL Fellows’ program (Horizon2020 grant 665667).

REFERENCES

- [1] R. Dorfman, “The detection of defective members of large populations,” *Ann. Math. Stats.*, vol. 14, no. 4, pp. 436–440, 1943.
- [2] A. Fernández Anta, M. A. Mosteiro, and J. Ramón Muñoz, “Unbounded contention resolution in multiple-access channels,” in *Distributed Computing*. Springer Berlin Heidelberg, 2011, vol. 6950, pp. 225–236.
- [3] R. Clifford, K. Efremenko, E. Porat, and A. Rothschild, “Pattern matching with don’t cares and few errors,” *J. Comp. Sys. Sci.*, vol. 76, no. 2, pp. 115–124, 2010.
- [4] G. Cormode and S. Muthukrishnan, “What’s hot and what’s not: Tracking most frequent items dynamically,” *ACM Trans. Database Sys.*, vol. 30, no. 1, pp. 249–278, March 2005.
- [5] A. Gilbert, M. Iwen, and M. Strauss, “Group testing and sparse signal recovery,” in *Asilomar Conf. Sig., Sys. and Comp.*, Oct. 2008, pp. 1059–1063.
- [6] A. C. Gilbert, M. J. Strauss, J. A. Tropp, and R. Vershynin, “One sketch for all: Fast algorithms for compressed sensing,” in *Proc. ACM-SIAM Symp. Disc. Alg. (SODA)*, New York, 2007, pp. 237–246.
- [7] M. Malyutov, “The separating property of random matrices,” *Math. Notes Acad. Sci. USSR*, vol. 23, no. 1, pp. 84–91, 1978.
- [8] M. B. Malyutov and P. S. Mateev, “Screening designs for non-symmetric response function,” *Mat. Zametki*, vol. 29, pp. 109–127, 1980.
- [9] A. G. D’yachkov, “Error probability bounds for the symmetrical model of the design of screening experiments,” *Prob. Inf. Transm.*, 1982.
- [10] A. G. D’yachkov and V. V. Rykov, “A survey of superimposed code theory,” *Prob. Contr. Inf.*, vol. 12, no. 4, pp. 1–13, 1983.
- [11] G. Atia and V. Saligrama, “Boolean compressed sensing and noisy group testing,” *IEEE Trans. Inf. Theory*, vol. 58, no. 3, pp. 1880–1901, March 2012.
- [12] M. Aldridge, L. Baldassini, and O. Johnson, “Group testing algorithms: Bounds and simulations,” *IEEE Trans. Inf. Theory*, vol. 60, no. 6, pp. 3671–3687, June 2014.
- [13] J. Scarlett and V. Cevher, “Phase transitions in group testing,” in *Proc. ACM-SIAM Symp. Disc. Alg. (SODA)*, 2016.
- [14] C. L. Chan, P. H. Che, S. Jaggi, and V. Saligrama, “Non-adaptive probabilistic group testing with noisy measurements: Near-optimal bounds with efficient algorithms,” in *Allerton Conf. Comm., Ctrl., Comp.*, Sep. 2011, pp. 1832–1839.
- [15] S. Cai, M. Jahangoshahi, M. Bakshi, and S. Jaggi, “GROTESQUE: Noisy group testing (quick and efficient),” 2013, <https://arxiv.org/abs/1307.2811>.
- [16] K. Lee, R. Pedarsani, and K. Ramchandran, “SAFFRON: A fast, efficient, and robust framework for group testing based on sparse-graph codes,” 2015, <https://arxiv.org/abs/1508.04485>.
- [17] M. B. Malyutov, “Search for sparse active inputs: A review,” in *Inf. Theory, Comb. and Search Theory*, 2013, pp. 609–647.
- [18] D. Malioutov and M. Malyutov, “Boolean compressed sensing: LP relaxation for group testing,” in *IEEE Int. Conf. Acoust. Sp. Sig. Proc. (ICASSP)*, March 2012, pp. 3305–3308.
- [19] M. Aldridge, L. Baldassini, and K. Gunderson, “Almost separable matrices,” *J. Comb. Opt.*, pp. 1–22, 2015.
- [20] M. Aldridge, “The capacity of Bernoulli nonadaptive group testing,” 2015, <http://arxiv.org/abs/1511.05201>.
- [21] J. Scarlett and V. Cevher, “Limits on support recovery with probabilistic models: An information-theoretic framework,” *IEEE Trans. Inf. Theory*, vol. 63, no. 1, pp. 593–620, 2017.
- [22] M. Aldridge, “On the optimality of some group testing algorithms,” in *IEEE Int. Symp. Inf. Theory*, 2017.
- [23] O. Johnson, M. Aldridge, and J. Scarlett, “Performance of group testing algorithms with constant tests-per-item,” 2016, submitted to *IEEE Trans. Sig. Proc.*
- [24] C. L. Chan, S. Jaggi, V. Saligrama, and S. Agnihotri, “Non-adaptive group testing: Explicit bounds and novel algorithms,” *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 3019–3035, May 2014.
- [25] D. Sejdinovic and O. Johnson, “Note on noisy group testing: Asymptotic bounds and belief propagation reconstruction,” in *Allerton Conf. Comm., Control and Comp.*, 2010.
- [26] A. Sharma and C. R. Murthy, “Computationally tractable algorithms for finding a subset of non-defective items from a large population,” 2015, <http://arxiv.org/abs/1502.04169>.
- [27] A. J. Macula, “Error-correcting nonadaptive group testing with disjoint matrices,” *Disc. App. Math.*, vol. 80, no. 2-3, pp. 217–222, 1997.
- [28] H. Q. Ngo, E. Porat, and A. Rudra, “Efficiently decodable error-correcting list disjoint matrices and applications,” in *Int. Colloq. Automata, Lang., and Prog.*, 2011.
- [29] M. Cheraghchi, “Noise-resilient group testing: Limitations and constructions,” *Disc. App. Math.*, vol. 161, no. 1, pp. 81–95, 2013.
- [30] M. Malyutov and H. Sadaka, “Maximization of ESL. Jaynes principle for testing significant inputs of linear model,” *Rand. Opt. Stoch. Eq.*, vol. 6, no. 4, pp. 339–358, 1998.
- [31] T. Laarhoven, “Asymptotics of fingerprinting and group testing: Tight bounds from channel capacities,” *IEEE Trans. Inf. Forens. Sec.*, vol. 10, no. 9, pp. 1967–1980, 2015.
- [32] T. S. Han, *Information-Spectrum Methods in Information Theory*. Springer, 2003.
- [33] S. Verdú and T. S. Han, “A general formula for channel capacity,” *IEEE Trans. Inf. Theory*, vol. 40, no. 4, pp. 1147–1157, July 1994.
- [34] Y. Polyanskiy, V. Poor, and S. Verdú, “Channel coding rate in the finite blocklength regime,” *IEEE Trans. Inf. Theory*, vol. 56, no. 5, pp. 2307–2359, May 2010.
- [35] A. Feinstein, “A new basic theorem of information theory,” *IRE Prof. Group. on Inf. Theory*, vol. 4, no. 4, pp. 2–22, Sept. 1954.
- [36] C. E. Shannon, “Certain results in coding theory for noisy channels,” *Information and Control*, vol. 1, no. 1, pp. 6–25, 1957.
- [37] S. Boucheron, G. Lugosi, and P. Massart, *Concentration Inequalities: A Nonasymptotic Theory of Independence*. OUP Oxford, 2013.
- [38] R. Motwani and P. Raghavan, *Randomized Algorithms*. Chapman & Hall/CRC, 2010.
- [39] A. G. D’yachkov, I. V. Vorobev, N. Polyansky, and V. Y. Shchukin, “Almost disjunctive list-decoding codes,” *Prob. Inf. Transm.*, vol. 51, no. 2, pp. 110–131, 2015.
- [40] J. Scarlett and V. Cevher, “How little does non-exact recovery help in group testing?” in *IEEE Int. Conf. Acoust. Sp. Sig. Proc. (ICASSP)*, New Orleans, 2017.
- [41] —, “Converse bounds for noisy group testing with arbitrary measurement matrices,” in *IEEE Int. Symp. Inf. Theory*, Barcelona, 2016.
- [42] A. Tauste Campo, G. Vazquez-Vilar, A. Guillén i Fabregas, and A. Martinez, “Random-coding joint source-channel bounds,” in *IEEE Int. Symp. Inf. Theory (ISIT)*. IEEE, 2011, pp. 899–902.
- [43] “Gurobi optimization,” <https://www.gurobi.com/>.
- [44] P. Grosu, “Parallel implementation of information-theoretic separate testing of inputs,” Ph.D. dissertation, Northeastern University, 2017.
- [45] M. Malyutov, P. Grosu, and H. Sadaka, “Separate testing inputs vs. linear programming relaxation,” *Information Processes*, vol. 15, no. 3, pp. 351–376, 2015.