

Statistical and Computational Phase Transitions in Group Testing

Amin Coja-Oghlan

AMIN.COJA-OGHLAN@TU-DORTMUND.DE

Oliver Gebhard

OLIVER.GEBHARD@TU-DORTMUND.DE

Max Hahn-Klimroth

MAXIMILIAN.HAHNKLIROTH@TU-DORTMUND.DE

TU Dortmund, Department of Computer Science, Dortmund 44227

Alexander S. Wein

AWEIN@CIMS.NYU.EDU

Algorithms and Randomness Center, Georgia Tech

Ilias Zadik

IZADIK@MIT.EDU

MIT, Department of Mathematics, Cambridge, MA 02139

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

We study the *group testing* problem where the goal is to identify a set of k infected individuals carrying a rare disease within a population of size n , based on the outcomes of pooled tests which return positive whenever there is at least one infected individual in the tested group. We consider two different simple random procedures for assigning individuals to tests: the *constant-column design* and *Bernoulli design*.

Our first set of results concerns the fundamental *statistical* limits. For the constant-column design, we give a new information-theoretic lower bound which implies that the proportion of correctly identifiable infected individuals undergoes a sharp “all-or-nothing” phase transition when the number of tests crosses a particular threshold. For the Bernoulli design, we determine the precise number of tests required to solve the associated detection problem (where the goal is to distinguish between a group testing instance and pure noise), improving both the upper and lower bounds of Truong, Aldridge, and Scarlett (2020).

For both group testing models, we also study the power of *computationally efficient* (polynomial-time) inference procedures. We determine the precise number of tests required for the class of *low-degree polynomial algorithms* to solve the detection problem. This provides evidence for an inherent *computational-statistical gap* in both the detection and recovery problems at small sparsity levels. Notably, our evidence is contrary to that of Iliopoulos and Zadik (2021), who predicted the absence of a computational-statistical gap in the Bernoulli design.¹

Keywords: Group testing, all-or-nothing phenomenon, low-degree likelihood ratio, computational-statistical gap

1. Introduction

Motivated by the ongoing COVID-19 pandemic (Mutesa et al., 2021; McMahan et al., 2012) but also a growing algorithmic and information-theoretic literature (Aldridge et al., 2019), in this work we focus on the *group (or pooled) testing model*. Introduced by Dorfman (1943), group testing is concerned with finding a subset of k individuals carrying a rare disease within a population of size n . One is equipped with a procedure that allows for testing groups of individuals such that a test returns positive if (and only if) at least one infected individual is contained in the tested group. The ultimate goal is to find a pooling procedure and an (efficient) algorithm such that inference of the

1. Full version appears as [arXiv 2206.07640, v1].

infection status of all individuals is conducted with as few tests as possible. Furthermore, group testing has found its way into various real-world applications such as DNA sequencing (Kwang-Ming and Ding-Zhu, 2006; Ngo and Du, 2000), protein interaction experiments (Mourad et al., 2013; Thierry-Mieg, 2006) and machine learning (Emad et al., 2015).

As carrying out a test is often time-consuming, many real-world applications call for fast identification schemes. As a consequence, recent research focuses on *non-adaptive* pooling schemes, i.e., all tests are conducted in parallel (Aldridge, 2019; Coja-Oghlan et al., 2020a,b; Iliopoulos and Zadik, 2021; Scarlett and Cevher, 2016). On top of this, naturally the testing scheme is required to be simple as well. Two of the most well-established and simple non-adaptive group testing designs are the *Bernoulli design* and the *constant-column design* (for a survey, see Aldridge et al. (2019)). The Bernoulli design is a randomised pooling scheme under which each individual participates in each test with a fixed probability q independently of everything else (Scarlett and Cevher, 2016). In the constant-column design (Aldridge et al., 2016; Coja-Oghlan et al., 2020a), each individual independently chooses a fixed number Δ of tests uniformly at random. We remark that the *spatially coupled design* of Coja-Oghlan et al. (2020b) may be an attractive choice in practice because it admits information-theoretically optimal inference with a computationally efficient algorithm. In this paper our focus will be on the two simpler designs (Bernoulli and constant-column), which may be favorable due to simplicity and also serve as a testbed for studying computational-statistical gaps.

In this work, we take the number of infected individuals to scale *sublinearly* in the population size, that is $k = n^{\theta+o(1)}$ for a fixed constant $\theta \in (0, 1)$. This regime is mathematically interesting and is also the one most suitable for modelling the early stages of an epidemic in the context of medical testing (Wang et al., 2011). In the two group testing models, we study two different inference tasks (defined formally in Section 2.1): (a) *approximate recovery*, where the goal is to achieve almost perfect correlation with the set of infected individuals, and (b) *weak recovery*, where the goal is to achieve positive correlation with the set of infected individuals. The task of *exact recovery* has also been studied (see Coja-Oghlan et al. (2020a)) but will not be our focus here.

Recently, there has been substantial work on the information-theoretic limits of group testing (Aldridge et al., 2014; Chan et al., 2011; Coja-Oghlan et al., 2020a,b; Truong et al., 2020). An interesting recent discovery is that for the Bernoulli group testing model there exists a critical threshold $m_{\text{inf}} := (\ln 2)^{-1} k \ln(n/k)$ such that when the number of tests m satisfies $m \geq (1 + \varepsilon)m_{\text{inf}}$ for any fixed $\varepsilon > 0$ there is a (brute-force) algorithm that can approximately recover the infected individuals, but when $m \leq (1 - \varepsilon)m_{\text{inf}}$ no algorithm (efficient or not) can even weakly recover the infected individuals. This sharp phase transition, known as the *All-or-Nothing (AoN) phenomenon*, was first proven by Truong et al. (2020) for $\theta = 0$ (that is, $k = n^{o(1)}$) and then proven for all $\theta \in [0, 1)$ by Niles-Weed and Zadik (2021). This sharp phenomenon has been established recently in many other sparse Generalized Linear Models (GLMs), starting with sparse regression (Reeves et al., 2019b). *Our first main result* (Theorem 1) establishes the AoN phenomenon in the constant-column group testing model for any $\theta \in (0, 1)$, occurring at the same information-theoretic threshold m_{inf} as in the Bernoulli model. To our knowledge, this is the first instance where AoN has been established for a GLM where the samples (tests) are not independent (see Section 1.1 for further discussion).

An emerging but less understood direction is to study the algorithmic thresholds of the group testing models. In both group testing models, the best known polynomial-time algorithm achieves approximate recovery only under the statistically suboptimal condition $m \geq (1 + \varepsilon)m_{\text{alg}}$ where $m_{\text{alg}} := (\ln 2)^{-1} m_{\text{inf}}$. For the constant-column design, the algorithm achieving this is Combinatorial Orthogonal Matching Pursuit (COMP) (Chan et al., 2011, 2014), which simply outputs all

individuals who participate in no negative tests. For the Bernoulli design, the algorithm achieving m_{alg} is called Separate Decoding (Scarlett and Cevher, 2018), which outputs all individuals who participate in no negative tests and “sufficiently many” positive tests (above some threshold). These results raise the question of whether better algorithms exist, or whether there is an inherent *computational-statistical gap*. Starting from the seminal work of Berthet and Rigollet (2013), conjectured gaps between the power of all estimators and the power of all *polynomial-time* algorithms have appeared recently throughout many high-dimensional statistical inference problems. While we do not currently have tools to prove complexity-theoretic hardness of statistical problems, there are various forms of “rigorous evidence” for hardness that can be used to justify these computational-statistical gaps, including average-case reductions (see e.g. Brennan and Bresler (2020)), sum-of-squares lower bounds (see e.g. Raghavendra et al. (2018)), and others.

In the Bernoulli group testing model, the recent work of Iliopoulos and Zadik (2021) suggested (but did not prove) that a polynomial-time Markov Chain Monte Carlo (MCMC) method can achieve approximate recovery all the way down to the information-theoretic threshold (that is, using only m_{inf} tests). The evidence for this is based on first moment Overlap Gap Property calculations and numerical simulations. The Overlap Gap Property is a landscape property originating in spin glass theory, which has been repeatedly used in inference to offer evidence for the performance of local search and MCMC methods, as initiated by Gamarnik and Zadik (2017). A significant motivation for the present work is to gain further insight into the existence or not of such a computational-statistical gap for both the constant-column and Bernoulli designs. Our approach is based on the well-studied *low-degree likelihood ratio* (discussed further in Section 2.2), which is another framework for understanding computational-statistical gaps.

In line with most existing results using the low-degree framework, we consider a *detection* (or *hypothesis testing*) formulation of the problem. In our case, this amounts to the task of deciding whether a given group testing instance was actually drawn from the group testing model with k infected individuals, or whether it was drawn from an appropriate “null” model where the test outcomes are random coin flips (containing no information about the infected individuals). *Our second set of results* is that for both the constant-column and Bernoulli designs, we pinpoint the precise low-degree detection threshold $m_{\text{LD}} = m_{\text{LD}}(k, n)$ (which is different for the two designs) in the following sense: when the number of tests exceeds this threshold, there is a polynomial-time algorithm that provably achieves *strong detection* (that is, testing with $o(1)$ error probability); on the other hand, if the number of tests lies below the threshold, all *low-degree algorithms* provably fail to *separate* the two distributions (as defined in Section 2.2). This class of low-degree algorithms captures the best known poly-time algorithms for many high-dimensional testing tasks (including those studied in this paper), and so our result suggests inherent computational hardness of detection below the threshold m_{LD} . For the exact thresholds, see Theorem 2 for the constant-column design and Theorem 3 for Bernoulli design.

Since approximate recovery is a harder problem than detection our results also suggest that approximate recovery is computationally hard below m_{LD} . Since m_{LD} exceeds m_{inf} for sufficiently small θ (see Figure 2), this suggests the presence of a computational-statistical gap for the recovery problem (in both group testing models). Notably, our evidence is contrary to that of Iliopoulos and Zadik (2021), who suggested the absence of a comp-stat gap in the Bernoulli model for all $\theta \in (0, 1)$!

Finally, *our third set of results* is to identify the precise *statistical* (information-theoretic) threshold for detection in the Bernoulli design (commonly referred to in the statistics literature as the *detection boundary*); see Theorem 4. Our main results are summarized in Figure 2.

1.1. Relation to Prior Work

Detection in the Bernoulli design To our knowledge, the only existing work on the detection boundary in group testing is [Truong et al. \(2020\)](#), which focused on the Bernoulli design. They gave a detection algorithm and an information-theoretic lower bound which did not match. In this work we pinpoint the precise information-theoretic detection boundary by improving both, the upper and lower bound (Theorem 4). The new algorithm involves counting the number of individuals who participate in no negative tests and “sufficiently many” positive tests (above some carefully chosen threshold). The lower bound of [Truong et al. \(2020\)](#) is based on a second moment calculation, and our improved lower bound uses a *conditional* second moment calculation (which conditions away a rare “bad” event).

Strictly speaking, our detection problem differs from the one studied by [Truong et al. \(2020\)](#) because our detection problem takes place on “pre-processed” graphs where the negative tests have been removed, but our results can be transferred to their setting.

All-or-Nothing phenomenon The All-or-Nothing (AoN) phenomenon was originally proven in the context of sparse regression with an i.i.d. Gaussian measurement matrix ([Gamarnik and Zadik, 2017](#); [Reeves et al., 2019a,b](#)), and was later established for (a) various other Generalized Linear Models (GLMs) such as Bernoulli group testing ([Niles-Weed and Zadik, 2021](#); [Truong et al., 2020](#)) and the Gaussian Perceptron ([Luneau et al., 2020](#); [Niles-Weed and Zadik, 2021](#)), (b) variants of sparse principal component analysis ([Barbier et al., 2020](#); [Niles-Weed and Zadik, 2020](#)), and (c) graph matching models ([Wu et al., 2021](#)). In all of the GLM cases, a key assumption behind all such proofs is that the samples (or tests in the case of Bernoulli group testing) are independent. This sample independence gives rise to properties similar to the I-MMSE formula ([Guo et al., 2005](#)), which can then be used to establish the AoN phenomenon by simply bounding the KL divergence between the planted model and an appropriate null model.

In the present work, we establish AoN for the constant-column group testing model which is a GLM where the samples (tests) are *dependent*. Despite this barrier, we manage to prove this result by following a more involved but direct argument, which employs a careful second moment argument alongside a technique from the study of random CSPs known as the “planting trick” originally used in the context of random k -SAT ([Achlioptas and Coja-Oghlan, 2008](#)).

Low-degree lower bounds Starting from the work of [Barak et al. \(2019\)](#); [Hopkins \(2018\)](#); [Hopkins et al. \(2017\)](#); [Hopkins and Steurer \(2017\)](#), lower bounds against the class of “low-degree polynomial algorithms” (defined in Section 2.2) are a common form of concrete evidence for computational hardness of statistical problems (see [Kunisky et al. \(2019\)](#) for a survey). In this paper we apply this framework to the detection problems in both group testing models, with a few key differences from prior work. For the Bernoulli design, the standard tool—the *low-degree likelihood ratio*—does not suffice to establish sharp low-degree lower bounds, and we instead need a *conditional* variant of this argument that conditions away a rare “bad” event. Along with the concurrent work of [Bandeira et al. \(2022\)](#), this is (to our knowledge) the first setting where a conditional low-degree argument has been needed. Our result for the constant-column design is (to our knowledge)

the first example of a low-degree lower bound where the null distribution does not have independent coordinates. For both group testing models, the key insight to make these calculations tractable is a “low-overlap second moment calculation.”

Comparison with Iliopoulos and Zadik (2021) Perhaps the most relevant work, in terms of studying the computational complexity of group testing, is the recent work of Iliopoulos and Zadik (2021) which focuses on the Bernoulli design. The authors provide simulations and first-moment Overlap Gap Property (OGP) evidence that a polynomial-time “local” MCMC method can approximately recover the infected individuals for any statistically possible number of tests $m \geq (1+\varepsilon)m_{\text{inf}}$ and any $\theta \in (0, 1)$. However, proving this remains open.

However, our present work shows that at least when $\theta > 0$ is small enough, no low-degree polynomial algorithm can even solve the easier detection task for some number of tests strictly above m_{inf} . Given the low-degree framework’s track record of capturing the best known algorithmic thresholds for a wide variety of statistical problems, this casts some doubts on the prediction of Iliopoulos and Zadik (2021). However, our results do not formally imply failure of the MCMC method (which is not a low-degree algorithm). Our results “raise the stakes” for proving statistical optimality of the MCMC method, as this would be a significant counterexample to optimality of low-degree algorithms for statistical problems.

Notation

We will consider the limit $n \rightarrow \infty$. Some parameters (e.g. θ, c) will be designated as “constants” (fixed, not depending on n) while others (e.g. k) will be assumed to scale with n in a prescribed way. Asymptotic notation $o(\cdot), O(\cdot), \omega(\cdot), \Omega(\cdot)$ pertains to this limit (unless stated otherwise), i.e., this notation may hide factors depending on constants such as θ, c . An event is said to occur *with high probability* (w.h.p.) if it has probability $1 - o(1)$.

2. Getting Started

2.1. Group Testing Setup and Objectives

We will consider two different group testing models. The following basic setup pertains to both.

Group testing We first fix two constants $\theta \in (0, 1)$ and $c > 0$. A group testing instance is generated as follows. There are n individuals x_1, \dots, x_n out of which exactly $k = n^{\theta+o(1)}$ are infected. There are $m = (c + o(1))k \ln(n/k)$ tests a_1, \dots, a_m .

For each test, a particular subset of the individuals is chosen to participate in that test, according to one of the two designs (constant-column or Bernoulli) described below. The assignment of individuals to tests can be expressed by a bipartite graph (see Figure 1). The *ground-truth* $\sigma \in \{0, 1\}^n$ is drawn uniformly at random among all vectors of length n and Hamming weight k . We say individual x_i is infected if and only if $\sigma_i = 1$. We denote the sequence of test results by $\hat{\sigma} \in \{0, 1\}^m$, where $\hat{\sigma}_i, i \in [m]$ is equal to one if and only if the i -th test contains at least one infected individual.

We consider two different schemes for assigning individuals to tests, which are defined below.

Constant-column design In the *constant column weight design* (also called the *random regular design*), every individual participates in exactly $\Delta = (c + o(1)) \ln(2) \ln(n/k)$ tests, chosen uniformly at random.

Bernoulli design In the *Bernoulli design*, every individual participates in each test independently with probability $q := \nu/k$ where $\nu = \ln 2 + o(1)$ is the solution to $(1 - \nu/k)^k = 1/2$ so that each test is positive with probability exactly $1/2$.

We remark that the parameter ν (in the Bernoulli design) and the constant $\ln(2)$ in the definition of Δ (in the constant-column design) could have been treated as free tuning parameters. To simplify matters, we have chosen to fix these values so that roughly half the tests are positive (maximizing the “information content” per test), but we expect our results could be readily generalized.

We will be interested in the task of recovering the ground truth σ . Two different notions of success are considered, as defined below.

Approximate recovery An algorithm is said to achieve *approximate recovery* if, given input $(\mathbf{G}_{GT}, \hat{\sigma}, k)$, it outputs a binary vector $\tau \in \{0, 1\}^n$ with the following guarantee: with probability $1 - o(1)$, $\frac{\langle \tau, \sigma \rangle}{\|\tau\|_2 \|\sigma\|_2} = 1 - o(1)$.

Equivalently, approx recovery means the number of false positive and false negatives are both $o(k)$.

Weak recovery An algorithm is said to achieve *weak recovery* if, given input $(\mathbf{G}_{GT}, \hat{\sigma}, k)$, it outputs a binary vector $\tau \in \{0, 1\}^n$ with the following guarantee: with probability $1 - o(1)$, $\frac{\langle \tau, \sigma \rangle}{\|\tau\|_2 \|\sigma\|_2} = \Omega(1)$.

Pre-processing via COMP Note that in both models we can immediately classify any individual who participates in a negative test as uninfected. Therefore, the first step in any recovery algorithm should be to pre-process the graph by removing all negative tests and their adjacent individuals. (We sometimes refer to this pre-processing step as COMP because it is the main step of the COMP algorithm of Chan et al. (2011, 2014), which simply performs this pre-processing step and then reports all remaining individuals as infected.) The resulting graph is denoted \mathbf{G}'_{GT} (see Figure 1). We let N denote the number of remaining individuals and let M denote the number of remaining tests. We use $\sigma' \in \{0, 1\}^N$ to denote the indicator vector for the infected individuals. Note that after pre-processing, all remaining tests are positive and so $\hat{\sigma}$ can be discarded.




Figure 1: The bipartite factor graph representing a group testing instance. Circles represent individuals while squares represent tests. The colour of circle/square indicates *infected / positive* in red and *uninfected / negative* in blue. The left figure shows an instance of \mathbf{G}_{GT} while the right figure shows the corresponding instance of \mathbf{G}'_{GT} where individuals in negative tests have already been classified and removed.

In addition to recovery, we will also consider an easier hypothesis testing task. Here the goal is to distinguish between a (“planted”) group testing instance and an unstructured (“null”) instance. We now define this testing model for both group testing designs. The input is an (N, M) -bipartite graph, representing a group testing instance that has already been pre-processed as described above.

Constant-column design (testing) Let $N = N_n$ and $M = M_n$ scale as $N = n^{1-(1-\theta)c(\ln 2)^2+o(1)}$ and $M = (c/2 + o(1))k \ln(n/k)$; this choice is justified below. Consider the following distributions over (N, M) -bipartite graphs (encoding adjacency between N individuals and M tests).

- Under the null distribution \mathbb{Q} , each of the N individuals participates in exactly Δ (defined above) tests, chosen uniformly at random.
- Under the planted distribution \mathbb{P} , a set of k infected individuals out of N is chosen uniformly at random. Then a graph is drawn from \mathbb{Q} conditioned on having at least one infected individual in every test.

Bernoulli design (testing) Let $N = N_n$ and $M = M_n$ scale as $N = n^{1-(1-\theta)\frac{c}{2} \ln 2 + o(1)}$ and $M = (c/2 + o(1))k \ln(n/k)$; this choice is justified below. Consider the following distributions over (N, M) -bipartite graphs (encoding adjacency between N individuals and M tests).

- Under the null distribution \mathbb{Q} , each of the N individuals participates in each of the M tests with probability q (defined above) independently.
- Under the planted distribution \mathbb{P} , a set of k infected individuals out of N is chosen uniformly at random. Then a graph is drawn from \mathbb{Q} conditioned on having at least one infected individual in every test.

Note that in the pre-processed group testing graph \mathbf{G}'_{GT} , the dimensions N, M are random variables. For the testing problems above, we will instead think of N, M as deterministic functions of n , which are allowed to vary arbitrarily within some range (due to the $o(1)$ terms). The specific scaling of N, M is chosen so that the actual dimensions of \mathbf{G}'_{GT} obey this scaling with high probability (see e.g. Coja-Oghlan et al. (2020a); Iliopoulos and Zadik (2021)). Furthermore, the planted distribution \mathbb{P} is precisely the distribution of \mathbf{G}'_{GT} conditioned on the dimensions N, M .

We now define two different criteria for success in the testing problem.

Strong detection An algorithm is said to achieve *strong detection* if, given input (\mathbf{G}, k) with \mathbf{G} drawn from either \mathbb{Q} or \mathbb{P} (each chosen with probability $1/2$), it correctly identifies the distribution (\mathbb{Q} or \mathbb{P}) with probability $1 - o(1)$.

Weak detection An algorithm is said to achieve *weak detection* if, given input (\mathbf{G}, k) with \mathbf{G} drawn from either \mathbb{Q} or \mathbb{P} (each chosen with probability $1/2$), it correctly identifies the distribution (\mathbb{Q} or \mathbb{P}) with probability $1/2 + \Omega(1)$.

We will establish a formal connection between the testing and recovery problems: any algorithm for approximate recovery can be used to solve strong detection (see the full version).

2.2. Hypothesis Testing and the Low-Degree Framework

Following Hopkins and Steurer (2017); Hopkins et al. (2017); Hopkins (2018), we will study the class of *low-degree polynomial algorithms* as a proxy for computationally-efficient algorithms (see also Kunisky et al. (2019) for a survey). Considering the hypothesis testing setting, suppose we have two (sequences of) distributions $\mathbb{P} = \mathbb{P}_n, \mathbb{Q} = \mathbb{Q}_n$ over \mathbb{R}^p for some $p = p_n$. Since our testing problems are over (N, M) -bipartite graphs, we will set $p = NM$ and take \mathbb{P}, \mathbb{Q} to be supported on $\{0, 1\}^p$ (encoding the adjacency matrix of a graph). A *degree- D polynomial algorithm*

is simply a multivariate polynomial $f : \mathbb{R}^p \rightarrow \mathbb{R}$ with real coefficients (or rather, a sequence of such polynomials: $f = f_n$). In line with prior work, we define two different notions of “success” for polynomial-based tests as follows.

Strong/weak separation A polynomial $f : \mathbb{R}^p \rightarrow \mathbb{R}$ is said to *strongly separate* \mathbb{P} and \mathbb{Q} if

$$\sqrt{\max\{\text{Var}_{\mathbb{P}}[f], \text{Var}_{\mathbb{Q}}[f]\}} = o\left(\left|\mathbb{E}_{\mathbb{P}}[f] - \mathbb{E}_{\mathbb{Q}}[f]\right|\right). \quad (2.1)$$

Also, a polynomial $f : \mathbb{R}^p \rightarrow \mathbb{R}$ is said to *weakly separate* \mathbb{P} and \mathbb{Q} if

$$\sqrt{\max\{\text{Var}_{\mathbb{P}}[f], \text{Var}_{\mathbb{Q}}[f]\}} = O\left(\left|\mathbb{E}_{\mathbb{P}}[f] - \mathbb{E}_{\mathbb{Q}}[f]\right|\right). \quad (2.2)$$

These are natural sufficient conditions for strong/weak detection: note that by Chebyshev’s inequality, strong separation immediately implies that strong detection can be achieved by thresholding the output of f ; also, by a less direct argument, weak separation implies that weak detection can be achieved using the output of f (Bandeira et al., 2022, Proposition 6.1).

Perhaps surprisingly, it has now been established that for a wide variety of “high-dimensional testing problems” (including planted clique, sparse PCA, community detection, tensor PCA, and many others), the class of degree- $O(\ln p)$ polynomial algorithms is precisely as powerful as the best known polynomial-time algorithms (e.g. Bandeira et al. (2020); Ding et al. (2019); Hopkins (2018); Hopkins et al. (2017); Hopkins and Steurer (2017); Kunisky et al. (2019)). One explanation for this is that such polynomials can capture powerful algorithmic frameworks such as spectral methods (see Kunisky et al. (2019), Theorem 4.4). Also, lower bounds against low-degree algorithms imply failure of all *statistical query algorithms* Brennan et al. (2021) and have conjectural connections to the *sum-of-squares hierarchy* (see e.g. Hopkins et al. (2017); Hopkins (2018)). While there is no guarantee that a degree- $O(\ln p)$ polynomial can be computed in polynomial time, the success of such a polynomial still tends to coincide with existence of a poly-time algorithm.

In light of the above, *low-degree lower bounds* (i.e. provable failure of all low-degree algorithms) is commonly used as a form of concrete evidence for computational hardness of statistical problems. In line with prior work, we will aim to prove hardness results of the following form.

Low-degree hardness If no degree- D polynomial achieves strong (respectively, weak) separation for some $D = \omega(\ln p)$, we say “strong (resp., weak) detection is low-degree hard”; this suggests that strong (resp., weak) detection admits no polynomial-time algorithm and furthermore requires runtime $\exp(\tilde{\Omega}(D))$ (where $\tilde{\Omega}$ hides factors of $\ln p$).

In this paper, we will establish low-degree hardness of group testing models in certain parameter regimes. While the implications for all polynomial-time algorithms are conjectural, these results identify apparent computational barriers in group testing that are analogous to those in many other problems. As a result, we feel there is unlikely to be a polynomial-time algorithm in the low-degree hard regime, at least barring a major algorithmic breakthrough.² Throughout the rest of this paper we focus on proving low-degree hardness as a goal of inherent interest, and refer the reader to the references mentioned above for further discussion on its interpretation.

2. Strictly speaking, we should perhaps only conjecture computational hardness for a slightly noisy version of group testing (say where a small constant fraction of test results are changed at random) because some “noiseless” statistical problems admit a poly-time algorithm in regimes where low-degree polynomials fail; see Section 1.3 of Zadik et al. (2021) for discussion.

3. Results

We now formally state our main results on statistical and computational thresholds in group testing, which are summarized in Figure 2. Throughout, recall that we fix the scaling regime $k = n^{\theta+o(1)}$ and $m = (c + o(1))k \ln(n/k)$ for constants $\theta \in (0, 1)$ and $c > 0$. Our objective is to characterize the values of (θ, c) for which various group testing tasks are “easy” (i.e., poly-time solvable), “hard” (in the low-degree framework), and (information-theoretically) “impossible.”

3.1. Constant-Column Design

Our first set of results pertains to the constant-column design, as defined in Section 2.1.

Weak recovery: All-or-Nothing phenomenon We start by focusing on the information-theoretic limits of weak recovery in the constant-column design. We show that the AoN phenomenon occurs at the critical constant $c_{\text{inf}} = 1/\ln 2$, i.e., at the critical number of tests $m_{\text{inf}} = (\ln 2)^{-1}k \ln(n/k)$. It was known previously that when $c > 1/\ln 2$, one can approximately recover the infected individuals via a brute-force algorithm (Coja-Oghlan et al., 2020a,b). It was also known that when $c < 1/\ln 2$, one *cannot* approximately recover the infected individuals (see Aldridge et al. (2019)). We show that in fact a much stronger lower bound holds: when $c < 1/\ln 2$, no algorithm can even achieve *weak* recovery.

Theorem 1 *Consider the constant-column design with any fixed $\theta \in (0, 1)$. If $c < c_{\text{inf}} := 1/\ln 2$ then every algorithm (efficient or not) taking input $(\mathbf{G}_{GT}, \hat{\sigma}, k)$ and returning a binary vector $\tau \in \{0, 1\}^n$ must satisfy $\frac{\langle \tau, \sigma \rangle}{\|\tau\|_2 \|\sigma\|_2} = o(1)$ with probability $1 - o(1)$. In particular, weak recovery is impossible.*

Combined with the prior work mentioned above, this establishes the All-or-Nothing phenomenon, namely:

- If $c > c_{\text{inf}}$ and $m = (c + o(1))k \ln(n/k)$ then *approximate* recovery is *possible*.
- If $c < c_{\text{inf}}$ and $m = (c + o(1))k \ln(n/k)$ then *weak* recovery is *impossible*.

As mentioned in the Introduction, the only algorithms known to achieve approximate recovery with the statistically optimal number of tests m_{inf} do not have polynomial runtime (Coja-Oghlan et al., 2020a,b). As a tool for studying this potential computational-statistical gap (and out of independent interest), we next turn our attention to the easier *detection* task. We will return to discuss the implications for hardness of the recovery problem later.

Detection boundary and low-degree methods We first pinpoint the precise “low-degree” threshold $c_{\text{LD}}^{\text{CC}} = c_{\text{LD}}^{\text{CC}}(\theta)$ (where the superscript indicates “constant-column”) for detection: above this threshold we prove that a new poly-time algorithm achieves strong detection; below this threshold we prove that all low-degree polynomial algorithms fail to achieve weak separation, giving concrete evidence for hardness (see Section 2.2). As a sanity check for the low-degree lower bound, we also verify that low-degree algorithms indeed succeed at strong separation above the threshold (specifically, this is achieved by a degree-2 polynomial that computes the empirical variance of the test degrees).

Theorem 2 Consider the constant-column design (testing variant) with parameters $\theta \in (0, 1)$ and $c > 0$. Define

$$c_{\text{LD}}^{\text{CC}} = \begin{cases} \frac{1}{\ln^2 2} \left(1 - \frac{\theta}{2(1-\theta)}\right) & \text{if } 0 < \theta < 2/3, \\ 0 & \text{if } 2/3 \leq \theta < 1. \end{cases} \quad (3.1)$$

- (a) (Easy) If $c > c_{\text{LD}}^{\text{CC}}$, there is a degree-2 polynomial achieving strong separation, and a polynomial-time algorithm achieving strong detection.
- (b) (Hard) If $c < c_{\text{LD}}^{\text{CC}}$ then any degree- $n^{\Omega(1)}$ polynomial fails to achieve weak separation. (This suggests that weak detection requires runtime $\exp(n^{\Omega(1)})$.)

We remark that when $\theta \geq 2/3$, the problem is “easy” for any constant $c > 0$ (and perhaps even for some sub-constant scalings for c , although we have not attempted to investigate this).

Hardness of recovery Above, we have given evidence for hardness of detection below the threshold $c_{\text{LD}}^{\text{CC}}$. One can show (see full version) that recovery is a formally harder problem than detection: any poly-time algorithm for approximate recovery can be made into a poly-time algorithm for strong detection, succeeding for the same parameters θ, c . These two results together give evidence for hardness of *recovery* below $c_{\text{LD}}^{\text{CC}}$, via a two-step argument: our low-degree hardness for detection leads us to conjecture that there is no poly-time algorithm for detection below $c_{\text{LD}}^{\text{CC}}$, and this conjecture (if true) formally implies that there is no poly-time algorithm for approximate recovery below $c_{\text{LD}}^{\text{CC}}$. (However, our results do not formally imply failure of *low-degree* algorithms for recovery.) Notably, it turns out that $c_{\text{LD}}^{\text{CC}}$ exceeds c_{inf} for some values of θ (namely $0 < \theta < 1 + \frac{1}{2 \ln 2 - 3} \approx 0.38$), revealing a possible-but-hard regime for recovery (Region I in Figure 2).

Since the recovery problem might be strictly harder than testing, our results do not pinpoint a precise computational threshold for recovery (even conjecturally). However, one case where we do pinpoint the computational recovery threshold is in the limit $\theta \rightarrow 0$: here, the thresholds $c_{\text{LD}}^{\text{CC}}$ and c_{alg} coincide, that is, our low-degree hardness result for detection matches the best known poly-time algorithm for recovery (COMP). This suggests that for small θ , the COMP algorithm is optimal among poly-time methods (for approximate recovery).

An interesting open question is to resolve the low-degree threshold for *recovery*, in the style of [Schramm and Wein \(2020\)](#). However, it is not clear that their techniques immediately apply here.

3.2. Bernoulli Design

Our second set of our results pertains to the Bernoulli design as defined in Section 2.1. As always, we fix the scaling regime $k = n^{\theta+o(1)}$ and $m = (c + o(1))k \ln(n/k)$ for constants $\theta \in (0, 1)$ and $c > 0$.

Detection boundary and low-degree methods Similarly to the constant-column design, we will determine both the statistical and low-degree thresholds for detection. In this case, the detection boundary is more complicated and involves the *Lambert W function*: for $x \geq -\frac{1}{e}$, $W_0(x)$ is defined to be the unique $y \geq -1$ satisfying $ye^y = x$. We begin with the low-degree threshold.

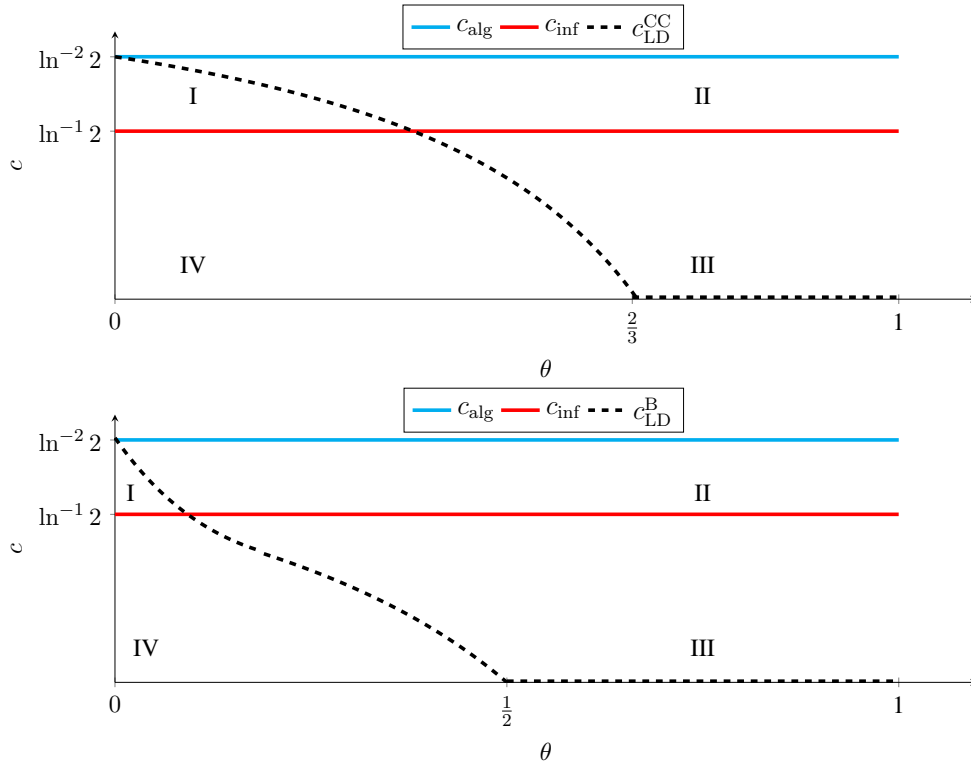


Figure 2: Phase transitions in the constant-column (left) and Bernoulli (right) designs, in (θ, c) space where $k = n^{\theta+o(1)}$ and $m = (c+o(1))k \ln(n/k)$. Recovery is possible above the red line and impossible below it. Polynomial-time recovery is only known above the blue line. Detection is achievable in polynomial time above the dotted line and (low-degree) hard below it. In Region I, detection and recovery are both possible-but-hard. In Region II, detection is easy and recovery is possible, but it is open whether recovery is easy or hard. In Region III, detection is easy and recovery is impossible. In Region IV, recovery is impossible; we expect detection is also impossible, and this is proven for the Bernoulli design only. Above the blue line, detection and recovery are both easy. See Section 3 for the formal statements.

Theorem 3 Consider the Bernoulli design (testing variant) with parameters $\theta \in (0, 1)$ and $c > 0$. Define

$$c_{\text{LD}}^{\text{B}} = \begin{cases} -\frac{1}{\ln^2 2} W_0(-\exp(-\frac{\theta}{1-\theta} \ln 2 - 1)) & \text{if } 0 < \theta < \frac{1}{2}(1 - \frac{1}{4\ln 2 - 1}), \\ \frac{1}{\ln 2} \cdot \frac{1-2\theta}{1-\theta} & \text{if } \frac{1}{2}(1 - \frac{1}{4\ln 2 - 1}) \leq \theta < \frac{1}{2}, \\ 0 & \text{if } \frac{1}{2} \leq \theta < 1. \end{cases} \quad (3.2)$$

- (a) (Easy) If $c > c_{\text{LD}}^{\text{B}}$, there is a degree- $O(\ln n)$ polynomial achieving strong separation, and a polynomial-time algorithm achieving strong detection.
- (b) (Hard) If $c < c_{\text{LD}}^{\text{B}}$, then any degree- $o(k)$ polynomial fails to achieve weak separation. (This suggests that weak detection requires runtime $\exp(\tilde{\Omega}(k))$.)

We remark that c_{LD}^{B} is a continuous function of θ (see Figure 2). The new algorithm that succeeds in the “easy” regime is based on counting the number of individuals whose degree (in the graph-theoretic sense) exceeds a particular threshold. For θ in the first case of (3.2), the low-degree hardness result requires a conditional argument that conditions away a certain rare “bad” event; for θ in the second case of (3.2), no conditioning is required and the resulting threshold matches the information-theoretic detection lower bound of [Truong et al. \(2020\)](#). We remark that the predicted runtime $\exp(\tilde{\Omega}(k))$ in the “hard” regime is essentially tight, matching the runtime of the brute-force algorithm up to log factors in the exponent.

Next, we determine the precise information-theoretic detection boundary. One (inefficient) detection algorithm is the brute-force algorithm for optimal *recovery*. Another (efficient) detection algorithm is the low-degree algorithm from Theorem 3 above. We show that for each $\theta \in (0, 1)$ statistically optimal detection is achieved by the better of these two algorithms. Brute-force is better when $\theta < 1 - \frac{\ln 2}{2\ln 2 - \ln \ln 2 - 1} \approx 0.079$, and otherwise low-degree is better.

Theorem 4 Consider the Bernoulli design (testing variant) with parameters $\theta \in (0, 1)$ and $c > 0$. Let $c_{\text{inf}} := 1/\ln 2$ and define c_{LD}^{B} as in (3.2).

- (a) (Possible) If $c > \min\{c_{\text{inf}}, c_{\text{LD}}^{\text{B}}\}$ then strong detection is possible.
- (b) (Impossible) If $c < \min\{c_{\text{inf}}, c_{\text{LD}}^{\text{B}}\}$ then weak detection is impossible.

Hardness of recovery Similarly to the constant-column design, our low-degree hardness results suggest hardness of recovery below the threshold c_{LD}^{B} (see the discussion in Section 3.1). This suggests a possible-but-hard regime for recovery (namely Region I in Figure 2) in the Bernoulli design, for sufficiently small θ (namely $\theta < 1 - \frac{\ln 2}{2\ln 2 - \ln \ln 2 - 1} \approx 0.079$). As discussed in the Introduction, this is contrary to the evidence of [Iliopoulos and Zadik \(2021\)](#), who predicted the absence of a computational-statistical gap for all $\theta \in (0, 1)$.

4. Proof Ideas

We give an overview of the proof techniques, with details deferred to the full version of this paper.

4.1. All-or-Nothing: Theorem 1

As approximate recovery is possible for $c > 1/\ln(2)$ (Coja-Oghlan et al., 2020b), we only need to prove the “nothing” part of the all-or-nothing phenomenon. Thus, our overall target is to show that for any $\delta > 0$, a randomly sampled solution τ satisfies $\mathbb{P}(\langle \sigma, \tau \rangle \geq \delta k) = o(1)$ if σ is the actual ground-truth. It suffices by standard arguments to prove that, if $Z(G)$ is the total number of solutions and $Z_\sigma(G, \alpha)$ the number of solutions with overlap αk to the ground-truth,

$$\sum_{\delta k \leq \ell \leq k} Z_\sigma(G, \ell/k) = o(Z(G)).$$

The calculations are technically demanding in a model with a *planted* infected set (the ground-truth). By an application of the planting trick, we manage to prove that it suffices to derive similar bounds for a null model. To further simplify calculus in the resulting second moment argument, we introduce a bi-regular model $\mathbb{Q}_{\Delta, \Gamma}^*$ that has fixed test degrees and individual degrees and allows an individual to participate multiple times in the same test. While the latter simplifies the underlying probability distribution, the former allows us to handle subtle local limit arguments that would not hold under fluctuating test degrees.

Therefore, the first step is to prove that this auxiliary model has comparable properties to the actual group testing instance. This highly technical proof uses the fact that with not too small probability, no infected individuals participate in the same test twice. For uninfected individuals, the calculus is much more convenient as there are no global dependencies, and the argument heavily relies on a technique similar to the *Poisson approximation of balls-into-bins*. Finally, a highly involved first moment argument that is again based on the Poisson approximation technique, proves that there is, with high probability, only a vanishing fraction of pairs of solutions σ, τ of non-trivial overlap that both satisfy a given group testing instance.

4.2. Detection Algorithms: Theorems 2(a) and 3(a)

In the constant-column design, the optimal low-degree test is given by the degree-2 polynomial $T = \sum_{j=1}^M (\Gamma_j - \frac{N\Delta}{M})^2$ where Γ_j denotes the degree of test j (i.e., the number of individuals in test j). As under \mathbb{P} every test contains at least one infected individual, there is a subset of k individuals that are distributed more equally over the whole graph, causing T to be smaller under \mathbb{P} than under \mathbb{Q} .

In the Bernoulli design, detection is easy in a larger regime of parameters because the degrees of the individuals are not fixed and these degrees give away some weak information about the planted signal: infected individuals are likely to have higher degree. The optimal low-degree test is based on counting the number of individuals whose degree exceeds a particular carefully-chosen threshold (which depends on the parameters θ, c). While this is not a polynomial, we show that it can be approximated by a logarithmic-degree polynomial up to negligible error.

In both models, the proof of strong separation proceeds by directly bounding the first two moments of the appropriate polynomials. These calculations are most challenging in the constant-column design, where the contribution from infected individuals is handled using the Poisson approximation of the balls-into-bins experiment.

4.3. Low-Degree Lower Bounds: Theorems 2(b) and 3(b)

Both low-degree lower bounds are proved using the notion of the “low-degree” χ^2 divergence between the two probability measures, as denoted by $\chi_{\leq D}^2(\mathbb{P} \parallel \mathbb{Q})$ (Hopkins, 2018). Analogous to the standard information-theoretic argument that $\chi^2(\mathbb{P}, \mathbb{Q}) = o(1)$ implies that weak detection between \mathbb{P}, \mathbb{Q} is impossible, it holds that if $\chi_{\leq D}^2(\mathbb{P} \parallel \mathbb{Q}) = o(1)$ then weak separation between \mathbb{P}, \mathbb{Q} is impossible by any D -degree polynomial. Hence, our goal is to prove $\chi_{\leq D}^2(\mathbb{P} \parallel \mathbb{Q}) = o(1)$. In the case of the Bernoulli design, we actually need a modification of this strategy where we condition away from a rare “bad” event (where individuals have unusually high degree).

In both models it seems challenging to directly bound $\chi_{\leq D}^2(\mathbb{P} \parallel \mathbb{Q})$; for the Bernoulli model this is due to the aforementioned conditioning, and for the constant-columns design this is because the null model is not a product measure and so the orthogonal polynomials are complicated. For this reason, we employ instead the recently introduced notion of the “low-overlap” χ^2 divergence between \mathbb{P}, \mathbb{Q} (Bandeira et al., 2022). This quantity corresponds a statistical-physics inspired object which (1) takes the form of an expectation of a relatively simple random variable and (2) can be proven to always upper-bound $\chi_{\leq D}^2(\mathbb{P} \parallel \mathbb{Q})$ for the group testing models. We are able to directly bound the low-overlap χ^2 divergence for the models and show that it is $o(1)$, which implies the desired result. Interestingly our lower bounds exactly match the positive algorithmic results discussed above.

4.4. Detection Boundary: Theorem 4

As discussed above, the possibility result follows from the low-degree algorithm discussed above along with the analysis of a brute-force test from prior work. It therefore remains to show the impossibility result. Similarly to the low-degree lower bounds, we achieve this by proving that the χ^2 divergence (or in some cases the conditional χ^2 divergence) is $o(1)$. This re-uses many of the calculations from the low-degree lower bounds. The main additional step is to also control the “high overlap” terms in the χ^2 divergence.

Acknowledgments

ACO was funded by DFG CO 646/3 and DFG CO 646/6. OG was funded by DFG CO 646/3. MHK was funded by DFG FOR 2975. ASW was funded by NSF awards CCF-2007443, CCF-2106444, DMS-1712730 and by the Simons Collaboration on Algorithms and Geometry; part of this work was done at the Simons Institute for the Theory of Computing and at the NYU Courant Institute. IZ is supported by the Simons-NSF grant DMS-2031883 and the Vannevar Bush Faculty Fellowship ONR-N00014-20-1-2826; part of this work was done at the Simons Institute for the Theory of Computing and at the NYU Center for Data Science. We thank Fotis Iliopoulos for helpful discussions during the first stages of this project.

References

- D. Achlioptas and A. Coja-Oghlan. Algorithmic barriers from phase transitions. *Proceedings of 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS'08)*, page 793–802, 2008.
- M. Aldridge. Individual testing is optimal for nonadaptive group testing in the linear regime. *IEEE Transactions on Information Theory*, 65:2058–2061, 2019.
- M. Aldridge, L. Baldassini, and O. Johnson. Group testing algorithms: bounds and simulations. *IEEE Transactions on Information Theory*, 60:3671–3687, 2014.
- M. Aldridge, O. Johnson, and J. Scarlett. Improved group testing rates with constant column weight designs. *Proceedings of 2016 IEEE International Symposium on Information Theory (ISIT'16)*, pages 1381–1385, 2016.
- M. Aldridge, O. Johnson, and J. Scarlett. Group testing: an information theory perspective. *Foundations and Trends in Communications and Information Theory*, 15:196–392, 2019.
- A. Bandeira, D. Kunisky, and A. Wein. Computational hardness of certifying bounds on constrained PCA problems. In *11th Innovations in Theoretical Computer Science Conference (ITCS'20)*, 2020.
- Afonso S Bandeira, Ahmed El Alaoui, Samuel B Hopkins, Tselil Schramm, Alexander S Wein, and Ilias Zadik. The Franz-Parisi criterion and computational trade-offs in high dimensional statistics. *arXiv preprint arXiv:2205.09727*, 2022.
- B. Barak, S. Hopkins, J. Kelner, P. Kothari, A. Moitra, and A. Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2): 687–735, 2019.
- J. Barbier, N. Macris, and C. Rush. All-or-nothing statistical and computational phase transitions in sparse spiked matrix estimation. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Q. Berthet and P. Rigollet. Computational lower bounds for sparse PCA. *arXiv preprint arXiv:1304.0828*, 2013.
- M. Brennan and G. Bresler. Reducibility and statistical-computational gaps from secret leakage. In *Proceedings of 33rd Conference on Learning Theory (COLT'20)*, pages 648–847, 2020.
- M. Brennan, G. Bresler, S. Hopkins, J. Li, and T. Schramm. Statistical query algorithms and low-degree tests are almost equivalent. In *Proceedings of 34th Conference on Learning Theory (COLT'21)*, 2021.
- C. Chan, P. Che, S. Jaggi, and V. Saligrama. Non-adaptive probabilistic group testing with noisy measurements: near-optimal bounds with efficient algorithms. *Proceedings of 49th Annual Allerton Conference on Communication, Control, and Computing*, 1:1832–1839, 2011.
- C. Chan, S. Jaggi, V. Saligrama, and S. Agnihotri. Non-adaptive group testing: Explicit bounds and novel algorithms. *IEEE Transactions on Information Theory*, 60(5):3019–3035, 2014.

- A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick. Information-theoretic and algorithmic thresholds for group testing. *IEEE Transactions on Information Theory*, 66(12):7911–7928, 2020a.
- A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick. Optimal group testing. *Proceedings of the 33rd Conference on Learning Theory (COLT'20)*, page 1–38, 2020b.
- Y. Ding, D. Kunisky, A. Wein, and A. Bandeira. Subexponential-time algorithms for sparse PCA. *arXiv preprint arXiv:1907.11635*, 2019.
- R. Dorfman. The detection of defective members of large populations. *Annals of Mathematical Statistics*, 14:436–440, 1943.
- A. Emad, K. Varshney, and D. Malioutov. A semiquantitative group testing approach for learning interpretable clinical prediction rules. *Signal Processing with Adaptive Sparse Structured Representations (SPARS'15)*, 2015.
- D. Gamarnik and I. Zadik. High dimensional linear regression with binary coefficients: Mean squared error and a phase transition. *Proceedings of 30th Conference on Learning Theory (COLT'17)*, 2017.
- D. Guo, S. Shamai, and S. Verdú. Mutual information and minimum mean-square error in gaussian channels. *IEEE Transactions on Information Theory*, 51(4):1261–1282, 2005.
- S. Hopkins. *Statistical Inference and the Sum of Squares Method*. PhD thesis, Cornell University, 2018.
- S. Hopkins and D. Steurer. Efficient bayesian estimation from few samples: community detection and related problems. In *IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS'17)*, pages 379–390. IEEE, 2017.
- S. Hopkins, P. Kothari, A. Potechin, P. Raghavendra, T. Schramm, and D. Steurer. The power of sum-of-squares for detecting hidden structures. In *IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS'17)*, pages 720–731. IEEE, 2017.
- F. Iliopoulos and I. Zadik. Group testing and local search: is there a computational-statistical gap? *Proceedings of the 34th Annual Conference on Learning Theory (COLT'21)*, 134:2499–2551, 2021.
- D. Kunisky, A. Wein, and A. Bandeira. Notes on computational hardness of hypothesis testing: Predictions using the low-degree likelihood ratio. *arXiv preprint arXiv:1907.11636*, 2019.
- H. Kwang-Ming and D. Ding-Zhu. Pooling designs and nonadaptive group testing: important tools for DNA sequencing. *World Scientific*, 2006.
- C. Luneau, J. Barbier, and N. Macris. Information theoretic limits of learning a sparse rule. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*, 2020.
- C. McMahan, J. Tebbs, and C. Bilder. Informative Dorfman screening. *Journal of the International Biometric Society*, 68:287–296, 2012.

- R. Mourad, Z. Dawy, and F. Morcos. Designing pooling systems for noisy high-throughput protein-protein interaction experiments using boolean compressed sensing. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 10:1478–1490, 2013.
- Leon Mutesa, Pacifique Ndishimye, Yvan Butera, Jacob Souopgui, Annette Uwineza, Robert Rutayisire, Ella Larissa Ndoricimpaye, Emile Musoni, Nadine Rujeni, Thierry Nyatanyi, et al. A pooled testing strategy for identifying SARS-CoV-2 at low prevalence. *Nature*, 589(7841):276–280, 2021.
- H. Ngo and D. Du. A survey on combinatorial group testing algorithms with applications to DNA library screening. *Discrete Mathematical Problems with Medical Applications*, 7:171–182, 2000.
- J. Niles-Weed and I. Zadik. The all-or-nothing phenomenon in sparse tensor PCA. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*, 2020.
- J. Niles-Weed and I. Zadik. It was “all” for “nothing”: sharp phase transitions for noiseless discrete channels. In *Proceedings of 34th Conference on Learning Theory (COLT’21)*, volume 134, pages 3546–3547, 2021.
- P. Raghavendra, T. Schramm, and D. Steurer. High dimensional estimation via sum-of-squares proofs. In *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018*, pages 3389–3423. World Scientific, 2018.
- G. Reeves, J. Xu, and I. Zadik. All-or-nothing phenomena: From single-letter to high dimensions. In *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 654–658, 2019a.
- G. Reeves, J. Xu, and I. Zadik. The all-or-nothing phenomenon in sparse linear regression. In *Proceedings of the Thirty-Second Conference on Learning Theory (COLT’19)*, volume 99, pages 2652–2663, 2019b.
- J. Scarlett and V. Cevher. Phase transitions in group testing. *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’16)*, 1:40–53, 2016.
- J. Scarlett and V. Cevher. Near-optimal noisy group testing via separate decoding of items. *IEEE Journal of Selected Topics in Signal Processing*, 12(5):902–915, 2018.
- T. Schramm and A. Wein. Computational barriers to estimation from low-degree polynomials. *arXiv preprint arXiv:2008.02269*, 2020.
- N. Thierry-Mieg. A new pooling strategy for high-throughput screening: the shifted transversal design. *BMC Bioinformatics*, 7:28, 2006.
- L. Truong, M. Aldridge, and J. Scarlett. On the all-or-nothing behavior of Bernoulli group testing. *IEEE Journal on Selected Areas in Information Theory*, 1(3):669–680, 2020.
- L. Wang, X. Li, Y. Zhang, Y. Zhang, and K. Zhang. Evolution of scaling emergence in large-scale spatial epidemic spreading. *PloS one*, 6(7):e21197, 2011.

- Y. Wu, J. Xu, and S. Yu. Settling the sharp reconstruction thresholds of random graph matching. *arXiv preprint arXiv:2102.00082*, 2021.
- I. Zadik, M. Song, A. Wein, and J. Bruna. Lattice-based methods surpass sum-of-squares in clustering. *arXiv preprint arXiv:2112.03898*, 2021.