

Fast and Efficient Data Science Techniques for COVID-19 Group Testing

VARLAM KUTATELADZE^{1,*} AND EKATERINA SEREGINA¹

¹*Department of Economics, University of California, Riverside, CA 92521, USA*

Abstract

Researchers and public officials tend to agree that until a vaccine is readily available, stopping SARS-CoV-2 transmission is the name of the game. Testing is the key to preventing the spread, especially by asymptomatic individuals. With testing capacity restricted, group testing is an appealing alternative for comprehensive screening and has recently received FDA emergency authorization. This technique tests pools of individual samples, thereby often requiring fewer testing resources while potentially providing multiple folds of speedup. We approach group testing from a data science perspective and offer two contributions. First, we provide an extensive empirical comparison of modern group testing techniques based on simulated data. Second, we propose a simple one-round method based on ℓ_1 -norm sparse recovery, which outperforms current state-of-the-art approaches at certain disease prevalence rates.

Keywords *compressed sensing; coronavirus; lasso; pooled testing; SARS-CoV-2; sensing matrix; sparse recovery*

1 Introduction

There is broad consensus among economists and epidemiologists that massive testing is one of the key ingredients to preventing the spread of COVID-19. However, large-scale testing is not realistic due to substantial restrictions in testing kits, chemical reagents, skilled personnel and time. Group testing, also known as pooled testing or specimen pooling, is an appealing alternative to individual testing that suggests to combine a set of individual specimens into a common pool, and test the pool rather than each individual sample. As long as the disease prevalence is not too large, testing pooled samples permits to considerably reduce the total number of tests required for diagnosing the population.

First experiments with pooling samples trace back to dilution studies in 1915 (Hughes-Oliver, 2006), which attempted to determine the presence or absence of organisms in a fluid based on pooled information. Researchers cultured samples of the fluid to let the bacteria, if they were present, grow, which served as a test. The results were then gathered across the samples to infer the bacterial density in the original fluid.

Many academics, however, attribute the invention of group testing to a Harvard economist Robert Dorfman, whose influential work (Dorfman, 1943) proposed a simple pooling method for weeding out syphilitic men called up for induction. Instead of analyzing individual blood samples for the presence or absence of a “syphilitic antigen”, it is suggested to examine pooled samples combining the individual blood sera into groups of five. If the corresponding men are healthy, the pooled test should be negative. On the other hand, if at least one of the patients

*Corresponding author. Email: varlam.kutateladze@email.ucr.edu.




Figure 1: Dorfman pooling illustration. “User” icons represent individuals, red are infected and grey are healthy. In the first stage, all $N = 30$ specimen are pooled into $N/n = 6$ groups (rows) of $n = 5$, which are then tested. In the second stage, everyone in infected groups (rows two and four) is tested individually. As a result, it is possible to detect $k = 2$ positives with $6 + 10 < N = 30$ tests.

is syphilitic, the pool will contain antigen, which the test is supposed to reveal. In that case, all associated patients need to be retested individually. Putting aside possible dilution concerns, it is clear that such strategy leads to savings of chemical reagents and higher overall testing capacity in a population with low disease prevalence. The idea is illustrated in Figure 1.

Due to its simplicity, Dorfman’s two-stage approach has found widespread use in medicine. Many of its properties are also readily available. Suppose we collect N individual samples and pool them into N/n groups of size n . Given the disease prevalence rate per hundred, p (which is also the probability of a randomly selected individual being positive), the expected number of tests for diagnosing the population is

$$\mathbb{E}(T) = N/n + \underbrace{(1 - (1 - p)^n)}_{\mathbb{P}(\text{at least one positive})} n \frac{N}{n}. \quad (1)$$

The first term on the right-hand side corresponds to the number of tests in the first stage, the second term is n times the expectation of a random variable distributed as $\text{Bi}(N/n, (1 - (1 - p)^n))$, characterizing the number of positive groups in the second stage. Clearly, this method is more beneficial at a lower prevalence rate p . For fixed $\mathbb{E}(T)$ and p , one could also optimize over the pool size n to get the largest possible coverage N ,

$$n^* = \frac{2W(-\frac{1}{2}\sqrt{-\ln(1-p)})}{\ln(1-p)}, \quad (2)$$

where $W(\cdot)$ is the Lambert W function. Notice that this also is decreasing in p and, interestingly, is independent of $\mathbb{E}(T)$. The expected number of tests per person is approximately minimized when the group size is $n = 1/\sqrt{p}$ and hence the expected number of tests per person is $2\sqrt{p}$. Graphs illustrating the above relationships are provided in Appendix A.1.

Sterrett (1957) proposed a modification to Dorfman’s second stage: instead of testing every individual, one would only do so until a positive sample is found, after which continue with group testing. In that case, if the prevalence is low, it is likely that the new sub-pool will be negative. This leads to an increase in savings of tests from Dorfman’s 80% (compared to individual testing) to 86% at 1% rate. There have been other alternatives as well, e.g. Sobel and Groll (1959), halving techniques in Litvak et al. (1994) and others. These methods generally trade off higher efficiency with more complexity and longer wait times.

In group testing literature, Dorfman’s approach and its modifications are classified as adaptive (or hierarchical), in a sense that they “adapt” to the results of preceding stages. An alternative approach, known as non-adaptive (or non-hierarchical), designs a single-stage experiment, results of which should allow to infer (often in a probabilistic manner) the original assignment of positives and negatives. This should generally, although not necessarily, come at the cost of having to run more tests overall since the sequential approach takes in more information. A distinctive feature of the non-adaptive approach is in assigning a single individual to multiple groups, i.e. groups are overlapping. How to design such assignment is a crucial question that is considered later. In a scenario such as the current SARS-CoV-2 pandemic, with the disease spreading fast and standard testing kits not showing results immediately, such non-adaptive approaches would have a clear advantage.

Furthermore, with multiple stages of testing, adaptive techniques also bear the risk of running out of tests before learning the outcomes. Given a limited number of tests, Dorfman’s approach may require more test than are available. In contrast, single-round designs do not suffer from such indeterminacy.

Hence, we focus on such “fast” single-stage techniques. We consider several recent combinatorial and probabilistic algorithms. Importantly, we propose a simple method based on ℓ_1 -norm sparse recovery, which outperforms the above algorithms.

Pooling strategies help resolve two kinds of problems, namely estimation and classification. The first seeks to estimate the prevalence of positive individuals in a population. The second, which may or may not rely on the information on estimated prevalence, aims to identify the infected individuals. The performance is typically gauged by the expected number of tests required for a given specificity or sensitivity, or conversely, based on predictive accuracy for a given number of tests. We focus on classification.

1.1 RT-qPCR Test

The two main ways of determining whether an individual has a SARS-CoV-2 virus are (1) to check for the presence of antibodies to the virus, (2) to check for the presence of the virus RNA itself. The former, although capable of uncovering whether a recovered individual had the virus in the past, is less widespread; the latter includes so-called reverse transcription quantitative polymerase chain reaction (RT-qPCR), the gold standard for COVID-19 testing recommended by the Centers for Disease Control and Prevention. Though popular, massive individual PCR testing is not possible due to serious constraints in equipment, chemical reagents and skilled personnel. The resulting readouts of RT-qPCR are of key interest to this study so we briefly describe the testing process.

To perform this test, nasopharyngeal swabs from subjects are collected and diluted in a fluid medium. The first stage, reverse transcription, then transforms the virus RNA to complementary DNA (cDNA). This eventually allows to start the next stage, polymerase chain reaction, which aims to exponentially amplify the viral cDNA molecules through the process that involves up to nearly 40 cycles of heating and cooling. To trace this increase, the virus-specific sequences are marked by fluorescent. The testing machine then measures the amount of fluorescent signal in real time and displays it as a function of cycles. This information we are interested in is when (if at all) the fluorescence exceeds the critical level associated with a positive subject. This is given by a cycle threshold (C_t), the number of cycles completed before crossing the threshold. The subject is then declared positive if the threshold is exceeded before about 40 cycles. The C_t indicator is (negatively) correlated with the original viral load, with larger initial viral loads

leading to sooner crossing of the threshold and thus shorter cycle thresholds. The entire process takes up to approximately 4 hours.

The information on cycle thresholds of pooled samples is the key input to group testing algorithms. Surprisingly, many known group testing algorithms do not take this quantitative information into account and instead work with degenerate binary transformations. The algorithm proposed in this study is capable of not only incorporating the quantitative information, but also producing corresponding quantitative predictions measuring the original individual viral loads.

1.2 Biomedical Considerations

One of the major concerns with pooling approaches is dilution. Fortunately, there is growing evidence that pooling of SARS-CoV-2 with negative samples does not lead to substantial dilution of the virus DNA. In a recent study (Yelin et al., 2020), Israeli researchers discovered that it is possible to detect a single positive SARS-CoV-2 sample in pools of up to 64 samples with reasonably high accuracy. That is, the fluorescent signal of a pooled positive sample, diluted with up to 63 negative samples, amplifies sufficiently to cross the required threshold. Other investigations (Abdalhamid et al., 2020, Hogan et al., 2020, Mutesa et al., 2020) tend to agree with such claims.

Pooled testing for SARS-CoV-2 has been conducted in a number of countries, including the United States (Stanford Health Care Clinical Virology Laboratory and Nebraska's Public Health Laboratory), Germany (University Hospital Frankfurt at Goethe University), China and Israel (Rambam Health Care Campus).

Group testing has been used for detecting the HIV (Emmanuel et al., 1988); in fact, it is now a routine option in blood screening. Pooling not only decreases the cost but also the probability of making an error in low disease prevalence populations. Pooling has also been deployed against malaria (Taylor et al., 2010), influenza (Van et al., 2012) and a few other diseases. It has also found its use in non-medical settings, for example detecting defective units in manufacturing, computer fault diagnosis or testing collections of documents in data forensics.

2 Non-adaptive Group Testing

We first briefly review some of the recent non-adaptive algorithms introduced in the literature (Chan et al., 2011). As all are one-round approaches, it is required to construct an $m \times N$ pooling matrix A ,¹ which would assign each of the N individuals to one or more of the m groups. As opposed to adaptive testing, we may have the same subject sample split among several groups. Figure 2 illustrates the idea. The corresponding pooling matrix has its (i, j) th entry equal to one if an i th individual is assigned to group j , $i \leq N$, $j \leq m$, and zero otherwise. One would typically also normalize the matrix, but this does not raise any substantial challenges so we omit this issue.

Once the pooling matrix is specified, one then observes the results of m pooled tests via an $m \times 1$ vector y and the goal is to identify which of the $N \gg m$ individuals are truly positive. We briefly discuss novel algorithms and provide an intuitive explanations behind their principles.

¹This matrix is not to be confused with the $(N/n) \times n$ Dorfman “matrix” in Figure 1. Dorfman approach does not involve allocating individuals to groups.




Figure 2: Illustration of a pooling matrix assigning N individuals to m pooled tests. A black circle indicates that the corresponding individual (column) has been assigned to the given test (row). If an individual is assigned to several tests, his sample is split accordingly.

Exact algorithmic formulations can be found in our code.

Combinatorial Basis Pursuit (CBP) is a simple algorithm that is based on the following idea: declare all individuals that were included in negative pools as negative (since if at least one sample was positive, the entire group would have been positive), and declare the remaining individuals as positives. Since this strategy would identify healthy individuals for certain, it will not produce any false negatives.

Instead of looking at each test (row), one may instead try to decode the matrix “column”-wise, i.e. by going over the individuals. This is what the combinatorial orthogonal matching pursuit (COMP) algorithm does: if all tests an individual participated in turn out to be positive, then the individual is considered to be infected, and negative otherwise. This deciphering method never produces false negatives, only false positives. A false positive would only occur if a healthy individual happened to always participate in tests that contained at least one infected sample. Of course the probability of this happening decreases with m .

Definite defectives (DD) algorithm starts with COMP to leverage its ability of identifying true negatives. Once COMP is over, DD switches to “row”-wise search by looking at positive tests and seeks to determine individuals that are “definite defectives” (positives). All remaining subjects are declared negative. This reversal in the algorithm leads to DD producing only false negatives, and no false positives. This leads to greater accuracy in sparse settings: because there are a lot more healthy subjects, one should by default assume that an individual is not infected, all else equal.

Finally, sequential COMP (SCOMP) further attempts to improve DD by modifying its last step of labeling remaining subjects as negatives. The key is to observe that if the current set of individuals that are declared positive cannot explain all of the positive pooled tests, one can do better by sequentially declaring potential positives as positives until the set of positives accounts for all positive tests. From a list of potential candidates, the algorithm picks the one that would account for the largest number of unexplained tests. SCOMP has been shown to perform close to the information-theoretic bound.

3 Problem Formulation

We now turn to our algorithm that leverages recent advancements in the field of compressed sensing in engineering and statistics literature (Candes et al., 2006; Donoho, 2006). Our main goal is to efficiently infer \mathbf{x} , the N -dimensional sparse vector of individual viral loads, from $m \ll N$ available group test results stacked in a vector \mathbf{y} . In general, we have $\mathbf{y} = g(\mathbf{A}\mathbf{x}) + \epsilon$,

however we will restrict our attention to the simplest case when $y = Ax + \epsilon$. Hence, $A \in \mathbb{R}^{m \times N}$ represents a set of linear measurements on the variable of interest x . This formulation has a crucial difference with regression type of problems: in our setting one gets to choose how to design the pooling matrix A , while in regression problems A is pre-determined by the data. Thus, there are two steps to solving such problems.

The first step is to encode the sparse signal, by designing a proper pooling matrix A . This matrix provides the assignments for each individual specimen to the corresponding groups and must satisfy certain desirable conditions pertaining to the group testing problem.

The second step attempts to decipher the first step with fewest test measurements. For a large-dimensional vector x finding the corresponding sparsest vector that would be consistent with m pooled observations is an NP-hard problem. However, recent advancements in engineering allow to transition this problem to a convex domain where exact decoding is feasible with high probability.

While x is generally a vector of quantitative measurements of all individuals, one can equivalently think of it as a sparse vector, i.e. with most entries equal to 0 (associated with healthy individuals) and very few 1's, without loss of generality. What matters is that the vector needs to be sparse in some transformed coordinate system.

3.1 Pooling Matrix Design

We first focus on the design of a pooling (also known as sensing or measurement) matrix A . Due to the nature of our primary application, we only consider sparse pooling matrices with few nonzero elements. When properly formed, this should ensure there is not too much dilution: a single sample is not split into too many subsamples and any one group sample does not contain too many specimens. The simplest approach would be to generate a random Bernoulli matrix with entries, which together with a normally distributed random matrix, has been shown to satisfy desirable properties, mainly the null space condition (NSC) and the restricted isometry property (RIP) discussed later, that guarantee a precise recovery of the original vector of interest with high probability.

We instead use a pooling matrix that was proposed in a different branch of group testing literature. Known as a constant column weight design (Aldridge et al., 2016), it was shown to outperform simple Bernoulli matrices in terms of its encoding capabilities. The initial approach outlined in Aldridge et al. (2016) constructs A by inserting up to an L of ones into each column. Concretely, L indices of each column are sampled with replacement and ones are inserted in the unique positions. This complication seems to be necessary for their proofs, however the real performance does not depend on whether one bootstraps or simply permutes a fixed number of ones. Hence, we focus on a simpler, permutation version. One example of such matrix with $N = 6$, $L = 2$ and $m = 4$ is

$$\begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 \end{bmatrix}$$

As discussed earlier, the i th individual is assigned to group j only if the (i, j) th entry is 1. This design avoids too much dilution as long as $0 < L < m \ll N$, and outperforms Bernoulli design. For simplicity, in our experiments we set $L = \lceil m/2 \rceil$ (rounded up), but this value could be theoretically and practically optimized as is done in Aldridge et al. (2016) and Johnson

et al. (2019). Importantly, we prove that this design is RIP which has immediate theoretical implications, which are discussed in the next section.

Theorem 1. *A random matrix $A \in \mathbb{R}^{m \times N}$ of constant column weight design with L ones in each column satisfies the restricted isometry property with high probability, specifically there exists $\delta \in (0, 1)$ such that*

$$(1 - \delta) \|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta) \|x\|_2^2$$

holds with high probability for any $x \in \mathbb{R}^N$ and $0 < L < m$.

Proof. See Appendix A.2. □

The ℓ_p norm of a vector $v \in \mathbb{R}^d$ is $\|v\|_p = (\sum_{i=1}^d |v_i|^p)^{1/p}$, which is a norm for $1 \leq p \leq \infty$.

3.2 ℓ_1 -Norm Sparse Recovery

Once m measurements in y are formed, one can employ several strategies for decoding the original signal. A direct, brute force approach to tackle the problem would be to find the sparsest vector of viral loads x that is consistent with the linear measurements, that is

$$\min_{x \in \mathbb{R}^N} \|x\|_0 \quad \text{s.t.} \quad \|Ax - y\|_2 \leq \epsilon.$$

Unfortunately, this problem is NP-hard as its solution requires an exhaustive search over all possible combinations in x , although this may still be feasible for low-dimensional problems. Luckily, a convenient convex relaxation is available, which has been proven to yield accurate solutions as long as the sensing matrix A satisfies RIP (Candes et al., 2006, Donoho, 2006). This is a sufficient condition, which Theorem 1 shows to hold with high probability. In practice one could generate a random matrix and attempt to verify whether RIP holds for a given matrix, although doing so is by itself NP-hard (Bandeira et al., 2013). The corresponding convex alternative is

$$\min_{x \in \mathbb{R}^N} \|x\|_1 \quad \text{s.t.} \quad \|Ax - y\|_2 \leq \epsilon. \quad (3)$$

This is known as Basis Pursuit Denoising (Chen and Donoho, 1994), although many statisticians are more familiar with its equivalent formulation, Lasso,

$$\min_{x \in \mathbb{R}^N} \|Ax - y\|_2^2 + \lambda \|x\|_1.$$

The two problems are identical for certain choices of ϵ and λ . We simply add a nonnegativity constraint $x \geq 0$ which reflects the inherent characteristic of the problem and also improves the empirical performance.

This type of ℓ_1 -norm recovery uses $m = O(k \log(N))$ tests while standard group testing algorithms require $m = O(k^2 \log(N))$ tests. Another advantage of this approach is its ability to handle real-valued quantitative readouts; many group testing algorithms are only capable of dealing with binary measurements. Furthermore, the output is also a real-valued number estimating individual's viral load.




Figure 3: Identification of negative and positive positions. Only the proposed (SR) algorithm is able to successfully determine which samples correspond to positive/negative specimens.

4 Application

This section compares the performance of the above algorithms in simple numerical experiments with no noise. For clarity of exposition, the vector of interest \mathbf{x} is generated to be a binary 0-1 vector instead of a real-numbered qPCR-like measurements. More general treatments can be found in our code.

We consider a case with $N = 100$ specimens where there are $k = 2$ true positive cases. This is a conservative estimate in a sense that this share of positives is larger than the share of active cases in the United States as of December 1, 2020 (Worldometer, 2020). In Appendix A.3 we additionally report cases with $k = 1, 3, 4, 5$.

To illustrate, we generate a 100-dimensional binary vector with 2 ones and the sensing matrices as described above to obtain $m = 20$ linear measurements. We then apply the decoding algorithms to try to infer the original binary vector (both the number of positive k and their positions) with only 20 measurements. Figure 3 demonstrates a particular realization where only the proposed algorithm, denoted as SR (for sparse recovery), is capable of correctly identifying the positions. Other algorithms produce either false positives, false negatives or both.

Specifically, to obtain SR estimates we first solve

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^N} \|\mathbf{Ax} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{x}\|_1, \quad \mathbf{x} \geq \mathbf{0}, \quad (4)$$

which generally would not produce 0-1 estimates. Hence we simply round the estimates at the threshold value of $\tau = .5$, that is $\hat{x}_i = \begin{cases} 1 & \text{if } \tilde{x}_i \geq \tau \\ 0 & \text{if } \tilde{x}_i < \tau. \end{cases}$

Next, we repeat this process for 1000 iterations across different group sizes m and report the average root mean square error (RMSE) plotted against the number of test measurements m in Panel A of Figure 4. RMSE is defined as $\frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_2}{\|\hat{\mathbf{x}}\|_2}$, where \mathbf{x} is a true binary vector and $\hat{\mathbf{x}}$ is one of the estimates. We still keep $N = 100$ and $k = 2$, but Appendix A.3 reports cases for $k = 1, 3, 4, 5$.

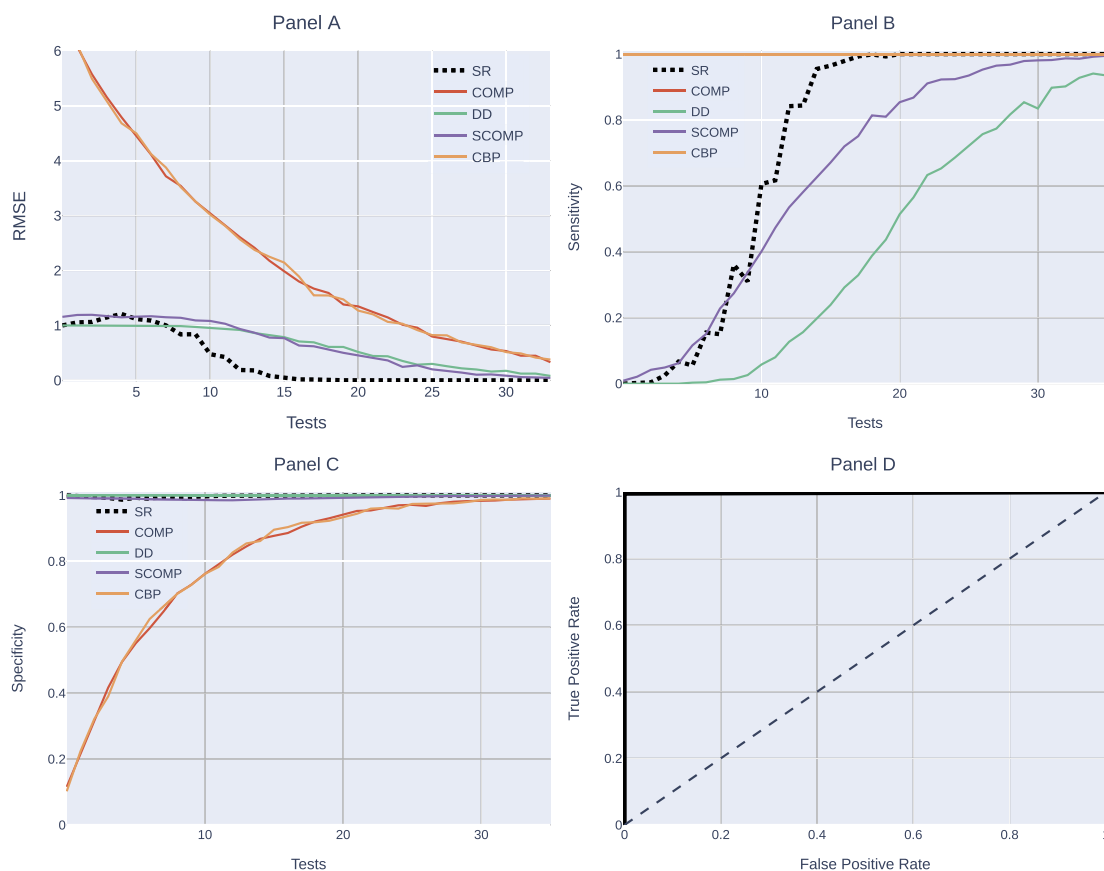


Figure 4: RMSE, Sensitivity, Specificity and ROC for $N = 100$, $k = 2$. *Panel A*: RMSE of each approach as a function of test measurements m . SR outperforms standard non-adaptive group testing algorithms. *Panel B*: Sensitivity of each approach as a function of test measurements m . *Panel C*: Specificity of each approach as a function of test measurements m . *Panel D*: ROC, thresholding SR estimates. AOC = .9995.

As can be seen, the proposed method makes approximately the same error with $m = 15$ tests as the best alternative (SCOMP) with $m = 30$ tests. For comparison, Dorfman approach would require approximately $m = 30$ tests and two testing stages. As the more detailed comparison in Appendix A.3 shows, SR is still superior for $k \geq 2$, but loses dominance to SCOMP for $k = 1$.

However, RMSE does not tell the whole story. One is also interested in sensitivity (or true positive rate) and specificity (or true negative rate). These are defined as the ratio of identified positives to all true positives and the ratio of identified negatives to all true negatives respectively, and reported in Panel B and C of Figure 4.

Notice that CBP and COMP report perfect sensitivity. This is a sanity check since these algorithms should not produce any false negatives. Among the other algorithm SR is again a clear winner. Naturally, the relationship between the two groups reverses for specificity: we have SR, COMP and DD achieving ideal (or almost ideal) specificity with a minimum number of tests, while COMP and CBP slowly catch up.

Additionally, we plot the receiver operating characteristic curve (ROC) for SR in Panel D of Figure 4, where we keep the same parameter values $N = 100$, $m = 20$, $k = 2$. The corresponding

Table 1: Improvement factors, $\frac{N}{\mathbb{E}(\# \text{ of tests})}$, for five different prevalence rates.

	$\frac{k}{N} = 1\%$	$\frac{k}{N} = 2\%$	$\frac{k}{N} = 3\%$	$\frac{k}{N} = 4\%$	$\frac{k}{N} = 5\%$
Dorfman	4.02	3.37	2.93	2.60	2.34
COMP	8.10	4.69	3.63	2.80	2.21
DD	6.25	2.86	2.39	1.99	1.99
CBP	9.42	4.49	3.60	2.81	2.19
SCOMP	9.83	3.95	3.06	2.48	2.10
SR	8.69	7.14	5.54	4.42	3.07

area under the curve (AUC) is .9995. This curve traces the true and false positive rates for different values of the threshold τ . The figure is indicative of a strong classification ability of the proposed method.

Finally, we report so-called improvement factors in Table 1, given as the ratio of the number of specimens N to the expected number of tests required for achieving at least 95% in specificity & sensitivity. When needed, the expected number is computed through Monte Carlo averaging. An improvement factor measures the effectiveness of a given method by computing how many more tests standard individual testing would need compared to a group testing algorithm. It essentially provides an estimate of how many individuals one group test effectively “covers”. For the five prevalence ratios, SR dominates both the non-adaptive algorithms and Dorfman approach for $k/N \geq 2\%$, while both CBP and SCOMP seem to be more efficient for $k/N = 1\%$.

5 Related Work

A closely related work by Yi et al. (2020) also considers the techniques based on compressed sensing. However, there are differences in both encoding and decoding steps. While their pooling matrix is also sparse, its 0-1 entries are generated from Bernoulli distribution with probability .5. Aldridge et al. (2016) and Johnson et al. (2019) compared the encoding capabilities of Bernoulli and constant column weight design matrices and documented substantial theoretical and empirical superiority of the latter in nonadaptive testing settings. The superiority is even more pronounced for sparse settings; in fact, such design “*in sparse cases is the best proven performance of any practical algorithm*” (Aldridge et al., 2016). Furthermore, Bernoulli design is not suitable for biomedical considerations as it is prone to undesirable extremes: one may have columns with all zeroes (or all ones) which corresponds to the case when the sample is not used at all (or is being split into too many subsamples); this is excluded by constant column weight design with $0 < L < m$. Another design Yi et al. (2020) to consider are expander matrices, which perform on par with Bernoulli in their experiments. However, the precise structure (for example, the number of ones per column) and theoretical fitness of such matrices are not discussed. Their decoding algorithm is non-negative BPDN (Eq. (3) coupled with a nonnegativity constraint) which has a less readily available software implementation than non-negative Lasso. We were not able to obtain the code of Yi et al. (2020) to make a direct comparison.

Another extensive study by Ghosh et al. (2020) also focuses on sparse nonadaptive methods. Researchers used so-called Kirkman triple matrices, which as in our work assign a finite number of 1’s per column. However, this design suffers from a major restriction on the dimensions of the sensing matrix, m and N , the number of groups and individuals respectively. Specifically, m

is only allowed to be an integer multiple of 3, while N needs to be an integer multiple of $m/3$. This makes direct Monte Carlo comparison with classical approaches, as is done e.g. in Figure 4, difficult, since m cannot vary freely. To resolve this by sequentially “truncating” their pooling matrix, i.e. only taking the first m rows from a fixed size 24×60 Tapestry matrix. The results are plotted in Appendix A.4. Our approach is more precise for lower disease prevalence rates, approximately for $k/N \leq 8\%$. Finally, although their matrix is claimed to satisfy RIP, we have not seen the proof.

6 Concluding Remarks

Pooled testing has been around for more than 70 years and has been successfully employed against a number of diseases. There are reasons to believe that pooling can also be effective against SARS-CoV-2. First, low prevalence of the virus is crucial to making group testing effective. Second, the recent evidence with dilution experiments suggests that pooling can be a viable method. Third, pooling is also compatible with the widely used testing kits such as RT-qPCR. Finally, group testing has been authorized by the FDA (FDA, 2020), which claimed it to be “especially important as infection rates decline and we begin testing larger portions of the population.”

To this end, we considered a simple one-stage group testing method that is able to diagnose a large number of specimens with the fewest number of tests and thus substantially increase the throughput of testing. Our approach does not require to know the number of positive samples in population to run and compares favorably based on the experiments on synthetic data. It produces very few false positives and false negatives, and is also capable of predicting viral loads. Compared to widely used adaptive strategies it minimizes latency in delivering test results, while compared with non-adaptive strategies it only requires $m \sim O(k \log n)$ tests. The numerical results suggest this approach is appealing for a wide range of prevalence rates; particularly, it outperforms standard non-adaptive methods at prevalence rates greater than 1%, and performs better than Kirkman matrices (Ghosh et al., 2020) for prevalence rates lower than 8%.

Supplementary Material

The code supplement (Kutateladze and Seregina, 2020) is available in Google Colab environment. It is written in Python and readily allows to replicate all the graphs provided, as well as produce additional exercises.

Appendix

A.1 Dorfman Group Testing Figures




Figure A.1: Dorfman group testing. *Upper left*: Dorfman theoretical optimal pool size, n^* , plotted over the prevalence rate, p . Larger pool sizes are preferred at a lower disease prevalence. *Upper right*: Dorfman expected number of tests, $\mathbb{E}(T)$, plotted over the prevalence rate, p . The benefits of classical approach diminish at higher prevalence rates. *Lower*: Number of individuals, N , Dorfman strategy can cover at 2% prevalence with $T = 1500$ tests, plotted over the number of groups, N/n . The maximum is achieved at $n = n^*$.

A.2 Proof of Theorem (1)

$A = \{a\}_{ij}$ is an $m \times N$ matrix where each column randomly permutes $0 < L < m$ ones among zeros. Without loss of generality, let us assume that each column has been demeaned and normalized to be of unit length, i.e. divided by $\sqrt{L(1 - \frac{L}{m})^2 + (m - L)(\frac{L}{m})^2} = \sqrt{L(1 - \frac{L}{m})}$. It is then evident that $\mathbb{E}(a_{ij}) = 0$ and $\mathbb{E}(a_{ij}^2) = \frac{1}{m}$.

First, we want to show that for any fixed $\mathbf{x} \in \mathbb{R}^N$, the random variable $\|\mathbf{Ax}\|_2^2$ concentrates around its mean, i.e.

$$\Pr\left(\left|\|\mathbf{Ax}\|_2^2 - \|\mathbf{x}\|_2^2\right| \geq \epsilon \|\mathbf{x}\|_2^2\right) \leq 2e^{-m(\epsilon^2/4 - \epsilon^3/6)}. \quad (5)$$

For $i = 1, \dots, m$, denote c_i the i th entry of \mathbf{Ax} , i.e. $c_i = \sum_{j=1}^N a_{ij}x_j$, then

$$\begin{aligned} \mathbb{E} c_i &= \mathbb{E} \left(\sum_{j=1}^N a_{ij}x_j \right) = \sum_{j=1}^N \mathbb{E}(a_{ij})x_j = 0, \\ \mathbb{E}(c_i^2) &= \mathbb{E} \left(\left(\sum_{j=1}^N a_{ij}x_j \right)^2 \right) = \mathbb{E} \left(\sum_{j=1}^N (a_{ij}x_j)^2 + 2 \sum_{l=1}^N \sum_{m=1}^N a_{ij}a_{mj}x_lx_m \right) \\ &= \sum_{j=1}^N \mathbb{E}(a_{ij}^2)x_j^2 + 2 \sum_{l=1}^N \sum_{m=1}^N \mathbb{E}(a_{lj}) \mathbb{E}(a_{mj})x_lx_m = \frac{1}{m} \|\mathbf{x}\|_2^2, \end{aligned}$$

and hence $\mathbb{E}(\|\mathbf{Ax}\|_2^2) = \mathbb{E}(\sum_{i=1}^m c_i^2) = \sum_{i=1}^m \mathbb{E}(c_i^2) = \|\mathbf{x}\|_2^2$.

Since $\|\mathbf{Ax}\|_2^2$ is proportional to $\|\mathbf{x}\|_2^2$, it is sufficient to demonstrate the concentration for arbitrary unit vectors. For all fixed unit vectors $\mathbf{x} \in \mathbb{R}^N$,

$$\mathbb{P}(\|\mathbf{Ax}\|_2^2 > 1 + \epsilon) = \mathbb{P}\left(e^{t\|\mathbf{Ax}\|_2^2} > e^{t(1+\epsilon)}\right) \quad (6)$$

$$< \mathbb{E}\left(e^{t\|\mathbf{Ax}\|_2^2}\right) e^{-t(1+\epsilon)}, \quad (7)$$

where (6) and (7) simply apply the Chernoff technique for $t > 0$. Now, because the columns c_i are i.i.d. we have $\mathbb{E}\left(e^{t\|\mathbf{Ax}\|_2^2}\right) = \mathbb{E}\left(e^{t\sum_{i=1}^m c_i^2}\right) = \left(\mathbb{E}\left(e^{tc_1^2}\right)\right)^m$, leading to

$$\mathbb{P}(\|\mathbf{Ax}\|_2^2 > 1 + \epsilon) < \left(\mathbb{E}\left(e^{tc_1^2}\right)\right)^m e^{-t(1+\epsilon)} \quad (8)$$

$$\leq (1 - 2t/m)^{-m/2} e^{-t(1+\epsilon)}, \quad (9)$$

where (9) follows from Lemma 1. Optimizing this bound with respect to t , $t^* = \frac{m\epsilon}{2(1+\epsilon)}$, we can write

$$\mathbb{P}(\|\mathbf{Ax}\|_2^2 > 1 + \epsilon) < ((1 + \epsilon)e^{-\epsilon})^{m/2} \quad (10)$$

$$< e^{-m(\epsilon^2/4 - \epsilon^2/6)}, \quad (11)$$

where the last inequality comes from truncating the Taylor approximation of (10). Similarly, for the other bound,

$$\mathbb{P}(\|\mathbf{Ax}\|_2^2 < 1 - \epsilon) < \left(\mathbb{E}\left(e^{-tc_1^2}\right)\right)^m e^{t(1-\epsilon)} \quad (12)$$

$$< \left(\mathbb{E}\left(1 - tc_1^2 + t^2c_1^4/2\right)\right)^m e^{t(1-\epsilon)} \quad (13)$$

$$\leq \left(1 - \frac{t}{m} + \frac{3t^2}{2m^2}\right)^m e^{t(1-\epsilon)} \quad (14)$$

$$= \left(1 - \frac{\epsilon}{2(1+\epsilon)} + \frac{3\epsilon^2}{8(1+\epsilon)^2}\right)^m e^{\frac{m\epsilon(1-\epsilon)}{2(1+\epsilon)}} \quad (15)$$

$$< e^{-m(\epsilon^2/4 - \epsilon^3/6)}, \quad (16)$$

where (13) and (16) is a Taylor approximation, (14) uses the fact $\mathbb{E}(c_1^4) = \frac{1}{m^2} \leq \frac{3}{m^2}$ and (15) plugs in the earlier value of t^* .

Lemma 1. For $m \geq 1$ and all $x \in \mathbb{R}^N$ s.t. $\|x\|_2^2 = 1$, $\mathbb{E}\left(e^{tc_1^2}\right) \leq (1 - 2t/m)$, $\forall t \in [0, m/2]$.

Proof. Let $W \sim \mathcal{N}(0, \frac{1}{m})$, then

$$\mathbb{E}\left(e^{tc_1^2}\right) = \sum_{i=1}^{\infty} \frac{t^i}{i!} \mathbb{E}(c_1^{2i}) \tag{17}$$

$$\leq \sum_{i=1}^{\infty} \frac{t^i}{i!} \mathbb{E}(W^{2i}) \tag{18}$$

$$= \mathbb{E}\left(e^{tW^2}\right) \tag{19}$$

$$= (1 - 2t/m)^{-1/2}. \tag{20}$$

Observe that for $t \in [0, m/2]$ the expectations in (17) and (19) are bounded, allowing to push the expectation inside the limiting sums in (17) and (18). Inequality in (18) holds since $\mathbb{E}(c_1^{2i}) = m^{-i} \leq \mathbb{E}(W^{2i}) = m^{-i} \frac{(2i)!}{i!2^i}$ holds for each $i = 0, 1, 2, \dots$ \square

Given the concentration of $\|Ax\|_2^2$ around its mean, RIP follows from Lemma 5.1 in Baraniuk et al. (2008), which is adapted and reiterated below for completeness.

Lemma 2. Let a random matrix $A \in \mathbb{R}^{m \times N}$ satisfy the concentration inequality in (5). Then, for any set T with $q = \#(T) < m$ and any $0 < \delta < 1$, we have

$$\Pr\left((1 - \delta) \|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta) \|x\|_2^2\right) \geq 1 - 2(12/\delta)^q e^{-(\delta/2)m(\epsilon^2/4 - \epsilon^3/6)}.$$

Proof. See Lemma 5.1 in Baraniuk et al. (2008). \square

A.3 Additional Experiments for $k = 1, 3, 4, 5$

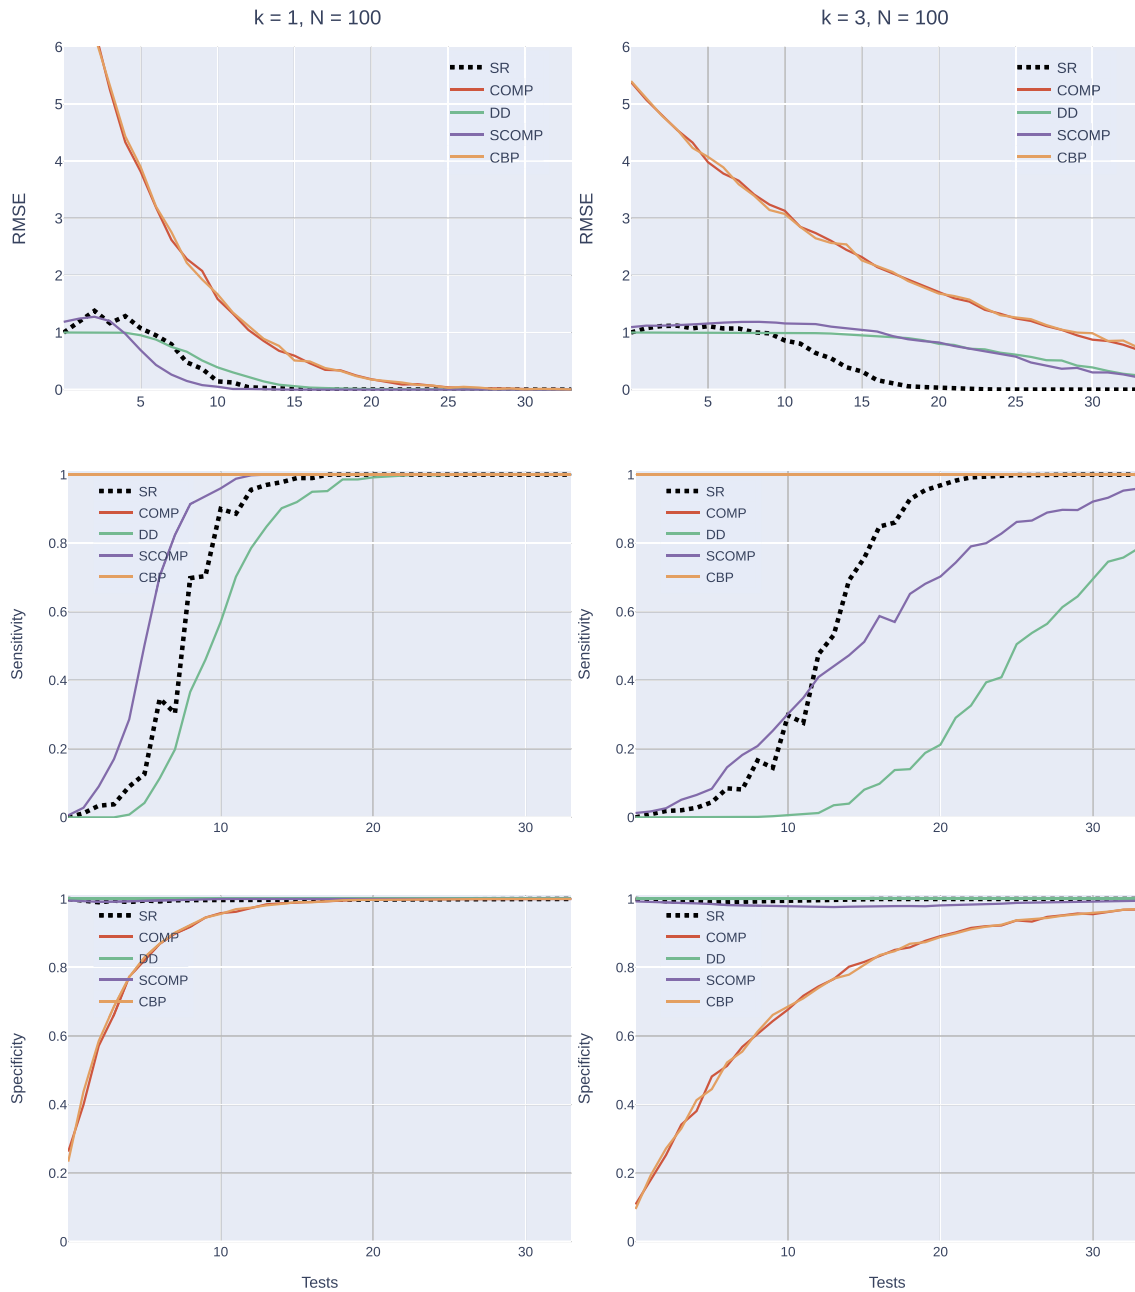


Figure A.2: Additional experiments for $N = 100$ and $k = 1, 3$. Rows represent RMSE, sensitivity and specificity (top to bottom), columns correspond to $k = 1$ and $k = 3$ (left to right) infected individuals in population of $N = 100$. The proposed approach is second-best for $k = 1$ and the best for $k = 3$.

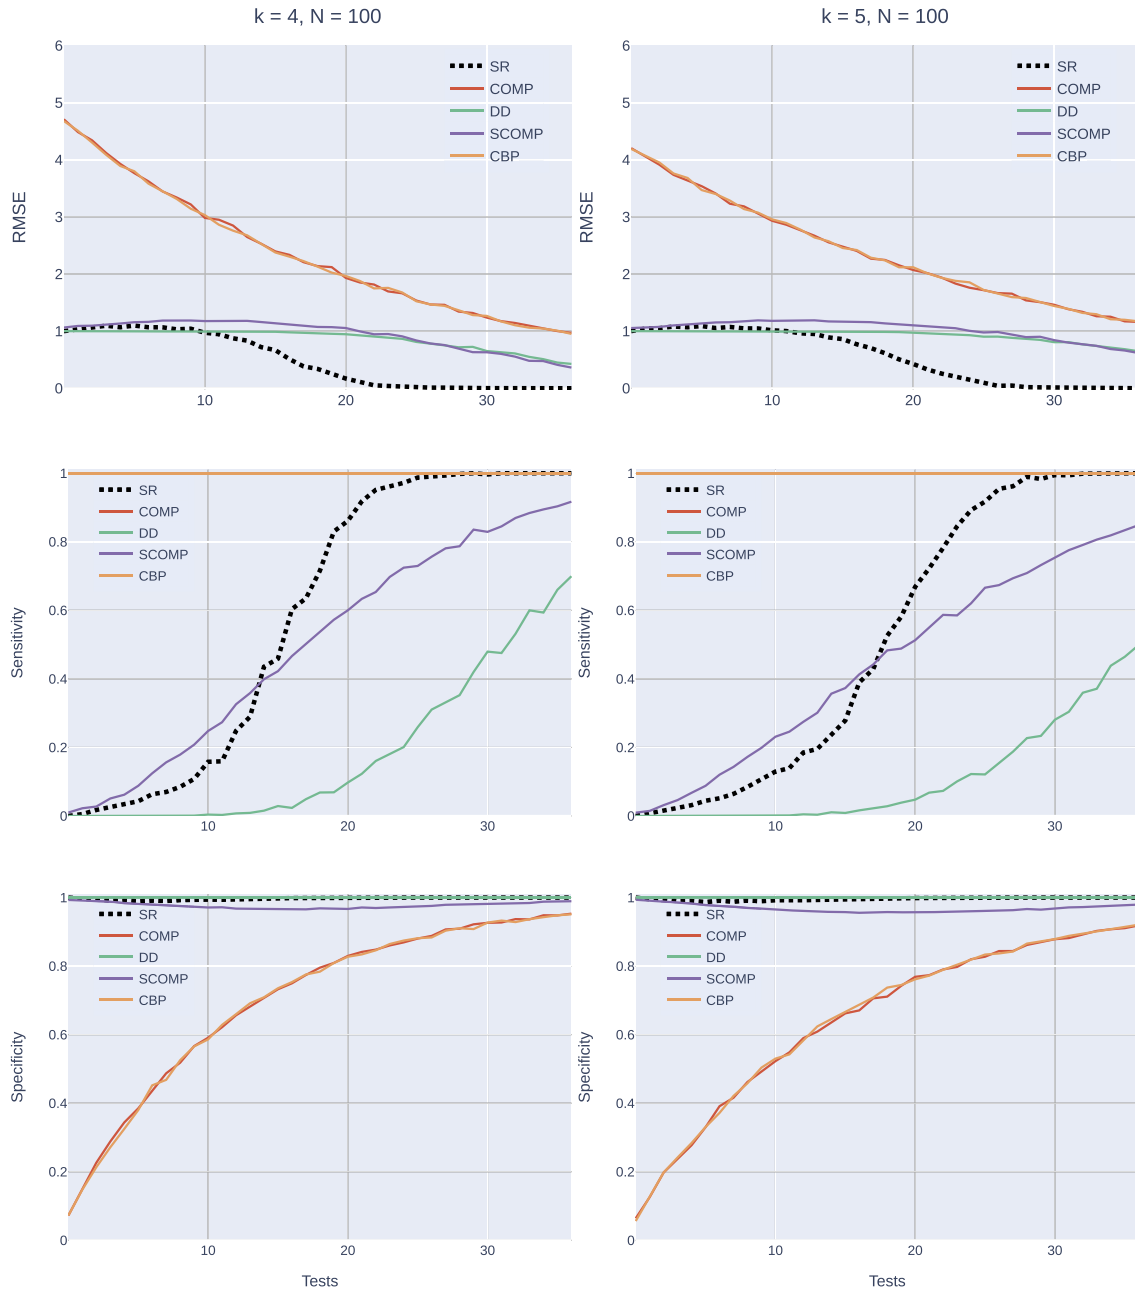


Figure A.3: Additional experiments for $N = 100$ and $k = 4, 5$. Rows represent RMSE, sensitivity and specificity (top to bottom), columns correspond to $k = 4$ and $k = 5$ (left to right) infected individuals in population of $N = 100$. The proposed approach retains its desirable properties.

A.4 Comparison with Tapestry (Ghosh et al., 2020)

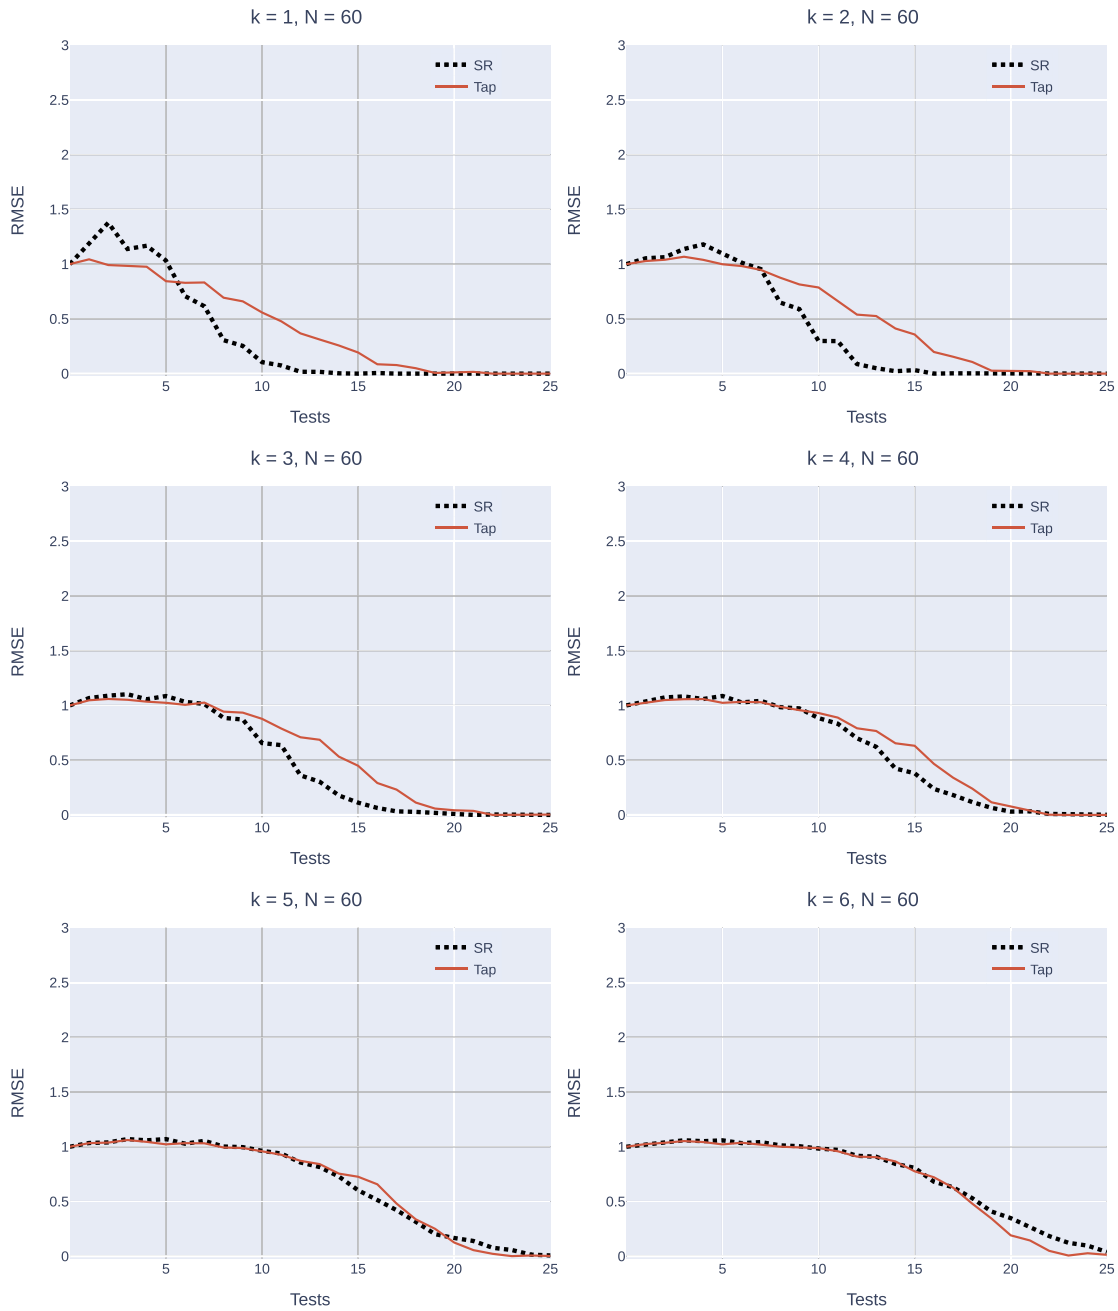


Figure A.4: RMSE comparison for $N = 60$, $k = 1, 2, 3, 4, 5, 6$. The proposed SR method dominates Tapestry (Ghosh et al., 2020) for $k = 1, 2, 3, 4$, but seems to be less precise for $k = 5, 6$.

Acknowledgement

We thank the anonymous referee and the editor, Jun Yan, for their very constructive comments, which led to a much improved presentation. In addition, the paper has benefited from the suggestions and insightful comments received from the participants of the Data Science Conference on COVID-19 (Dr. Jun Yan), OCLB American Statistical Association (Dr. Xinping Cui), Day of Data Science 2020 (Abidemi K. Adeniji), NABE Tech Economics Conference 2020 (Manuel Balmaseda).

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. Declarations of interest: none.

References

- Abdalhamid B, Bilder CR, McCutchen EL, Hinrichs SH, Koepsell SA, Iwen PC (2020). Assessment of specimen pooling to conserve SARS CoV-2 testing resources. *American Journal of Clinical Pathology*, 153(6): 715–718.
- Aldridge M, Johnson O, Scarlett J (2016). Improved group testing rates with constant column weight designs. In: *2016 IEEE International Symposium on Information Theory (ISIT)*, 1381–1385.
- Bandeira AS, Dobriban E, Mixon DG, Sawin WF (2013). Certifying the restricted isometry property is hard. *IEEE Transactions on Information Theory*, 59(6): 3448–3450.
- Baraniuk R, Davenport M, DeVore R, Wakin M (2008). A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3): 253–263.
- Candes EJ, Romberg J, Tao T (2006). Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2): 489–509.
- Chan CL, Che PH, Jaggi S, Saligrama V (2011). Non-adaptive probabilistic group testing with noisy measurements: Near-optimal bounds with efficient algorithms. In: *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 1832–1839.
- Donoho DL (2006). Compressed sensing. *IEEE Transactions on Information Theory*, 52(4): 1289–1306.
- Dorfman R (1943). The detection of defective members of large populations. *Ann. Math. Statist.*, 14(4): 436–440.
- Emmanuel JC, Bassett MT, Smith HJ, Jacobs JA (1988). Pooling of sera for human immunodeficiency virus (hiv) testing: An economical method for use in developing countries. *Journal of Clinical Pathology*, 41(5): 582–585.
- FDA (2020). Emergency Authorization for Sample Pooling. <https://www.fda.gov/news-events/press-announcements/coronavirus-covid-19-update-fda-issues-first-emergency-authorization-sample-pooling-diagnostic>.
- Ghosh S, Agarwal R, Rehan M, Pathak S, Agarwal P, Gupta Y, et al. (2020). A compressed sensing approach to group-testing for COVID-19 detection.
- Hogan CA, Sahoo MK, Pinsky BA (2020). Sample pooling as a strategy to detect community transmission of SARS-CoV-2. *JAMA*, 323(19): 1967–1969.

- Hughes-Oliver JM (2006). *Pooling Experiments for Blood Screening and Drug Discovery*. 48–68. Springer, New York, New York, NY.
- Johnson O, Aldridge M, Scarlett J (2019). Performance of group testing algorithms with near-constant tests per item. *IEEE Transactions on Information Theory*, 65(2): 707–723.
- Kutateladze V, Seregina E (2020). Code supplement to “Fast and Efficient Data Science Techniques for COVID-19 Group Testing”. <https://tinyurl.com/y4vo86sb>.
- Litvak E, Tu XM, Pagano M (1994). Screening for the presence of a disease by pooling sera samples.
- Mutesa L, Ndishimye P, Butera Y, Souopgui J, Uwineza A, Rutayisire R, et al. (2020). A strategy for finding people infected with SARS-CoV-2: optimizing pooled testing at low prevalence. medRxiv preprint: <https://doi.org/10.1101/2020.05.02.20087924>.
- Chen S, Donoho D (1994). Basis pursuit. In: *Proceedings of 1994 28th Asilomar Conference on Signals, Systems and Computers*, volume 1, 41–44. 1.
- Sobel M, Groll PA (1959). Group testing to eliminate efficiently all defectives in a binomial sample. *Bell System Technical Journal*, 38(5): 1179–1252.
- Sterrett A (1957). On the detection of defective members of large populations. *The Annals of Mathematical Statistics*, 28(4): 1033–1036.
- Taylor SM, Juliano JJ, Trottman PA, Griffin JB, Landis SH, Kitsa P, et al. (2010). High-throughput pooling and real-time pcr-based strategy for malaria detection. *Journal of Clinical Microbiology*, 48(2): 512–519.
- Van TT, Miller J, Warshauer DM, Reisdorf E, Jernigan D, Humes R, et al. (2012). Pooling nasopharyngeal/throat swab specimens to increase testing capacity for influenza viruses by pcr. *Journal of Clinical Microbiology*, 50(3): 891–896.
- Worldometer (2020). US SARS-CoV-2 cases. <https://www.worldometers.info/coronavirus/country/us/>.
- Yelin I, Aharony N, Shaer-Tamar E, Argoetti A, Messer E, Berenbaum D, et al. (2020). Evaluation of COVID-19 rt-qpcr test in multi-sample pools. medRxiv preprint: <https://doi.org/10.1101/2020.03.26.20039438>.
- Yi J, Mudumbai R, Xu W (2020). Low-cost and high-throughput testing of COVID-19 viruses and antibodies via compressed sensing: System concepts and computational experiments.