ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) | 978-1-7281-7605-5/20/$31.00 ©2021 IEEE | DOI: 10.1109/ICASSP39728.2021.9414034

CONTACT TRACING ENHANCES THE EFFICIENCY OF COVID-19 GROUP TESTING

Ritesh Goenka,[⋆1] *Shu-Jie Cao,*[⋆2] *Chau-Wai Wong,*[3] *Ajit Rajwade,*[1] *Dror Baron*[3]

[1] IIT Bombay, India, [2] ShanghaiTech University, China, [3] North Carolina State University, USA

ABSTRACT

Group testing can save testing resources in the context of the ongoing COVID-19 pandemic. In group testing, we are given n samples, one per individual, and arrange them into $m < n$ pooled samples, where each pool is obtained by mixing a subset of the n individual samples. Infected individuals are then identified using a group testing algorithm. In this paper, we use side information (SI) collected from contact tracing (CT) within nonadaptive/single-stage group testing algorithms. We generate data by incorporating CT SI and characteristics of disease spread between individuals. These data are fed into two signal and measurement models for group testing, where numerical results show that our algorithms provide improved sensitivity and specificity. While Nikolopoulos et al. utilized family structure to improve nonadaptive group testing, ours is the first work to explore and demonstrate how CT SI can further improve group testing performance.

***Index Terms*—** Contact tracing, nonadaptive group testing, compressed sensing, overlapping group LASSO, generalized approximate message passing (GAMP).

1. INTRODUCTION

Widespread testing has been promoted for combating the ongoing COVID-19 pandemic. Samples are typically collected from nasal or oropharyngeal swabs, and then processed by a reverse transcription polymerase chain reaction (RT-PCR) machine. However, widespread testing is hindered by supply chain constraints and long testing times.

Pooled or *group testing* has been suggested for improving testing efficiencies [2–17]. Group testing involves mixing a subset of n individual samples into $m < n$ pools. The measurement process can be expressed as $\boldsymbol{y} = \mathfrak{N}(\boldsymbol{A}\boldsymbol{x})$, where $\boldsymbol{x}$ is a vector that quantifies the health status of the n individuals, $\boldsymbol{A}$ is an $m \times n$ binary pooling matrix with $A_{ij} = 1$ if the jth individual contributes to the ith pool, else $A_{ij} = 0$, $\boldsymbol{y}$ is a vector of m noisy measurements or tests, and $\mathfrak{N}$ represents a probabilistic noise model that relates the noiseless pooled results, $\boldsymbol{A}\boldsymbol{x}$, to $\boldsymbol{y}$. We consider two signal and noise models.

Model M1: A *binary noise* model used by Zhu et al. [6], where $\boldsymbol{x}$ is binary, $\boldsymbol{w} = \boldsymbol{A}\boldsymbol{x}$ is an auxiliary vector, and the measurement $y_i \in \{0, 1\}$ depends probabilistically on w_i, where $\Pr(y_i = 1|w_i = 0)$ and $\Pr(y_i = 0|w_i > 0)$ are probabilities of erroneous tests.

Model M2: A *multiplicative noise* model of the form $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} \circ \boldsymbol{z}$ as used in Ghosh et al. [8], where $\circ$ represents element-wise multiplication, $\boldsymbol{z}$ is a vector of m noisy elements defined as $z_i = (1+q)^{\eta_i}$, $q \in (0, 1]$ is a known amplification factor for RT-PCR, $\eta_i \sim \mathcal{N}(0, \sigma^2)$, and $\sigma^2 \ll 1$ is a known parameter controlling the strength of the noise in RT-PCR. Under model **M2**, $\boldsymbol{x}$ and $\boldsymbol{y}$ represent viral loads in the n individuals and m pools, respectively. Assuming reasonably high viral loads in $\boldsymbol{x}$, Poisson effects in $\boldsymbol{y}$ can be ignored [8].

For both models, we wish to estimate $\boldsymbol{x}$ from $\boldsymbol{y}$ and $\boldsymbol{A}$. We use single-stage *nonadaptive* algorithms as in [6, 8], rather than two-stage algorithms, which employ a second stage of tests depending on results from the first stage, as in Heidarzadeh and Narayanan [10] or the classical Dorfman approach [2]. The advantage of nonadaptive algorithms is that they reduce testing time, which is high for RT-PCR. Algorithms that estimate $\boldsymbol{x}$ from $\boldsymbol{y}$ and $\boldsymbol{A}$ [8, 12] rely primarily on the *sparsity* of $\boldsymbol{x}$, which is a valid assumption for COVID-19 due to low prevalence rates [18]. Zhu et al. [6] also exploit probabilistic information such as the prevalence rate and structure in $\boldsymbol{x}$, and state the potential benefits of using *side information* (SI). Specific forms of SI include individuals' symptoms and family structure [9]. Finally, Nikolopoulos et al. independently observed that taking into account community structure can improve the performance of group testing [11, 15]; these works focused on the encoder design and used different contact models.

In this paper, we show how to estimate $\boldsymbol{x}$ while utilizing *contact tracing* (CT) SI, which allows one to analyze the spread of the pandemic [19]. Our contributions are twofold. First, we propose a generative model for a population of n individuals that characterizes the spread of COVID-19 by explicitly using CT SI. Second, we show that CT SI, when used appropriately, can help algorithms such as generalized approximate message passing (GAMP) [20] or LASSO variants [21, 22] better estimate $\boldsymbol{x}$ from $\boldsymbol{y}$ and $\boldsymbol{A}$. Our work uses more SI than Nikolopoulos et al. [11, 15], who only considered family-style structure in binary group testing.

2. DATA GENERATION MODEL

In this section, we present a generative infection model incorporating CT SI, which we later use to prepare simulated data for algorithmic evaluation. We model a population of

⋆ RG and SJC have made equal contributions. AR acknowledges support from SERB Grant #10013890, IITB-WRCB Grant #DONWR04-002, and DST-Rakshak grant #DST0000-005. CWW acknowledges support from NSF Grant #2030430. More details appear in a supplemental document [1]. The authors thank Junan Zhu for allowing them to use his implementation of GAMP with SI in their implementation of the family and CT denoisers.

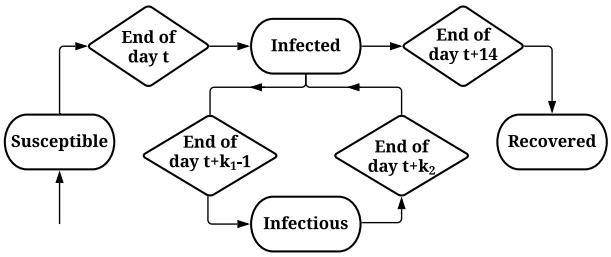

Fig. 1. State transition diagram for a node. A node is infectious only between days k_1 and k_2 (both inclusive) after getting infected. In our work, we set $(k_1, k_2) = (3, 7)$.

n individuals using a dynamical or time-varying graphical model that contains nodes $\{v_i\}_{i=1}^n$ and undirected edges $\{e_{ij}^{(t)}\}_{i,j=1}^n$. On a given day t, an edge $e_{ij}^{(t)}$ between nodes v_i and v_j encodes CT SI $\left(\tau_{ij}^{(t)}, d_{ij}^{(t)}\right)$, which can be acquired via Bluetooth-based CT applications [23]. Here, $\tau_{ij}^{(t)}$ represents the contact duration and $d_{ij}^{(t)}$ represents a measure of the physical proximity between two individuals. On day t, a node can be in one of the following states: *susceptible*, *infected*, *infectious*, and *recovered*. To keep the model simple, we assume that there are no reinfections, i.e., recovered is a terminal state, despite some reports of reinfection [24]. While our model is inspired from the classical SEIR model also considered for COVID-19 [25], our state transitions explicitly use CT information and knowledge about the pandemic [26].

We adopt a simplified infection dynamic wherein the infectious period is preceded and followed by the infected state. Our design parameters for the infection dynamics are based on a World Health Organization report on COVID-19 [26]. Specifically, a node v_i remains infected but noninfectious for $k_1 = 3$ days. On day $t + k_1$, the node becomes infectious and may transmit the disease to a susceptible neighboring node v_j with probability $p_{i,j}^{(t+k_1)}$ whose construction is described below. An infectious node can potentially transmit the infection until $k_2 = 7$ days after getting infected, and becomes noninfectious afterward. We also model a small fraction of stray infections that may occur, for example, due to sporadic contact with contaminated surfaces. Such infections only affect nodes in the susceptible state with a probability $p_1 = 2\times10^{-4}$ of our choice. A state diagram appears in Fig. 1. Regarding the viral load $x_i^{(t)}$ for node i on day t, we assume $x_i^{(t)} = 0$ if the node is susceptible or recovered. For an infected or infectious node, we make a simplified assumption that its viral load $x_i^{(t)} \sim \text{Uniform}(1, 32768)$,[1] once drawn, remains constant throughout the combined 14-day period of infection.

Next, we model the probability $p_{i,j}^{(t)}$ that the disease is transmitted from node v_i to v_j on day t. We view the infection times of the population throughout the pandemic as a nonhomogeneous Poisson process with time-varying rate function

[1]The cycle threshold for RT-PCR commonly ranges from 19 to 34 cycles [27, Fig. 3], where 34 cycles corresponds to a low initial viral load of a few molecules, and each cycle roughly doubles the viral density. Therefore, we estimate the largest possible viral load as $2^{34-19} = 2^{15} = 32768$.

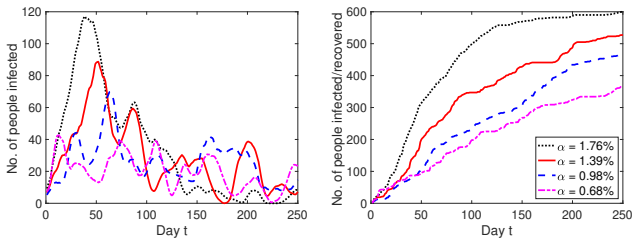

(a) (b)

Fig. 2. (a) The number of active infections, and (b) cumulative infections at different inter-clique contact levels α (defined in Sec. 2). We chose 50-day windows for testing proposed algorithms.

$\lambda(t)$. Consider a $\tau_{ij}^{(t)}$-hour contact on day t when susceptible node v_j is exposed to infectious node v_i. The average infection rate $\lambda_{ij}(t)$ for day t is assumed to be proportional to both the viral load $x_i^{(t)}$ and the physical proximity $d_{ij}^{(t)}$, namely, $\lambda_{ij}(t) = \lambda_0\, x_i^{(t)}\, d_{ij}^{(t)}$, where λ_0 is a tunable, baseline Poisson rate. The probability that v_j is infected by the end of contact period $\tau_{ij}^{(t)}$ is therefore $p_{i,j}^{(t)} = 1-\exp\left(-\lambda_0\, x_i^{(t)}\, d_{ij}^{(t)}\, \tau_{ij}^{(t)}\right)$ for $t \in [k_1, k_2] + t_i$. From the standpoint of susceptible node v_j, all its neighbors v_k that are infectious contribute to its probability of getting infected on day t, namely, $1 - \prod_k \left(1 - p_{k,j}^{(t)}\right)$.

While generating our data, we considered $n = 1000$ nodes divided into cliques based on the distribution of family sizes in India [28, pg. 18], for a duration of $t_{\max} = 250$ days. Fig. 2 shows the number of active infections and the cumulative number of infections at the end of each day. The clique structures were kept constant throughout the $t_{\max}$ days, whereas inter-clique contacts corresponding to sporadic contacts between people were dynamically added and removed. Here, a contact is defined as a pair of individuals who come in contact. We define the number of inter-clique contacts divided by the number of contacts for a given day as the inter-clique contact level, α. The varying α affects the sparsity of the underlying vector $\boldsymbol{x}$ as it brings infections to new cliques/families. Pooling of samples is performed at the beginning of each day from $t_{\text{peak}} - 24$ to $t_{\text{peak}} + 25$, where t_{peak} is the day with the maximum number of active infections.

3. PROPOSED GROUP TESTING ALGORITHMS

This section describes two classes of group testing algorithms for reconstructing the health status vector $\boldsymbol{x}$ from the pooled tests, $\boldsymbol{y}$, and pooling matrix, $\boldsymbol{A}$.

Algorithms for binary noise. For model **M1**, Zhu et al. [6] use *generalized approximate message passing* (GAMP) [20] for group testing estimation. GAMP is comprised of two components. The first component consists of an input channel that relates a prior for n individuals' viral loads, $\boldsymbol{x} = (x_i)_{i=1}^n$, and pseudo data, $\boldsymbol{v} = \boldsymbol{x} + \boldsymbol{q} \in \mathbb{R}^n$, where the n coordinates of $\boldsymbol{x}$ are correlated, and $\boldsymbol{q}$ is additive white Gaussian noise with $q_i \sim \mathcal{N}(0, \Delta)$. We estimate $\boldsymbol{x}$ from $\boldsymbol{v}$ using a denoising function (often called a denoiser):

$$\widehat{x}_i = g_{\text{in}}(\boldsymbol{v}) = \mathbb{E}\left[X_i \mid \boldsymbol{V} = \boldsymbol{v}\right], \tag{1}$$

where we use the convention that when both the upper

and lower case versions of a symbol appear, the upper case is a random variable and the lower case its realization, and $\mathbb{E}[X_i|\boldsymbol{v}]$ represents $\mathbb{E}[X_i|\boldsymbol{V}=\boldsymbol{v}]$ when the context is clear. The second component of GAMP consists of an output channel relating the auxiliary vector $\boldsymbol{w}$ to the noisy measurements $\boldsymbol{y}$ as reviewed in Sec. 1. We adopt the output channel denoiser of Zhu et al. [6], $h_i = g_{\text{out}}(y_i; k_i, \theta_i) = (\mathbb{E}[W_i \mid y_i, k_i, \theta_i] - k_i)/\theta_i$, where θ_i is the estimated variance of h_i, and k_i is the mean of our estimate for w_i. Since y_i depends probabilistically on w_i, we have $f(w_i \mid y_i, k_i, \theta_i) \propto \Pr(y_i \mid w_i) \exp\left[-\frac{(w_i-k_i)^2}{2\theta_i}\right]$, where W_i is approximated as Gaussian in the derivation of GAMP.

While Zhu et al. [6] considered Bernoulli $\boldsymbol{x}$, which implies a scalar separable denoiser g_{in} for the input channel, this paper accounts for probabilistic dependencies within $\boldsymbol{x}$. Our first probabilistic model considers groups of people, for example, members of a family. Each family is modeled as entirely healthy with probability $1-\pi_{\text{vf}}$, else each individual within the family is infected with probability π_{ind}. This model relates to our generative model of Sec. 2 by using family structure as SI. Denoting the pseudo data of family $\mathcal{F}$ by $\boldsymbol{v}_{\mathcal{F}}$, the denoiser for the ith individual of family $\mathcal{F}$ is given by

$$g_{\text{in}}^{\text{family}}(\boldsymbol{v}_{\mathcal{F}}) = \mathbb{E}\big[X_i|\mathcal{F}\text{ viral}, \boldsymbol{v}_{\mathcal{F}}\big]\ \Pr\big(\mathcal{F}\text{ viral}|\boldsymbol{v}_{\mathcal{F}}\big), \quad (2)$$

where $\mathbb{E}\big[X_i|\mathcal{F}\text{ viral}, \boldsymbol{v}_{\mathcal{F}}\big]$ and $\Pr\big(\mathcal{F}\text{ viral}|\boldsymbol{v}_{\mathcal{F}}\big)$ are parameterized by π_{vf}, π_{ind}, and Δ. For detailed expressions, we refer readers to Sec. 1.1 of the supplemental document.

Our second probabilistic model uses CT. Consider a hypothetical widespread testing program that relies on CT SI, where all individuals are tested 8 days before the group testing program begins resulting in a good estimate of their ground-truth health status. After the program begins, probability estimates from the previous group test are used as priors for the n individuals when performing the current group test. We provide detailed analysis in Secs. 2.2–2.3 of the supplemental document on the use of prior infection status. The final form of the denoiser for the CT model is as follows:

$$g_{\text{in}}^{\text{CT}}(v_i) = \Big\{1 + \big[\Pr(X_i=1)^{-1} - 1\big] \exp\Big[\big(v_i - \tfrac{1}{2}\big)/\Delta\Big]\Big\}^{-1}. \quad (3)$$

Here, $\Pr(X_i=1)$ for day $k+1$ can be estimated by aggregating CT information of individual i over the past 8 days, namely, $\widehat{\Pr}^{(k+1)}(X_i=1) = 1 - \prod_{d=k-7}^{k}\prod_{j=1}^{n}\big(1-\widehat{p}_{i,j}^{(d)}\big)$, where $\widehat{p}_{i,j}^{(d)}$ is the estimated probability of infection of individual i due to contact with individual j. This probability, $\widehat{p}_{i,j}^{(d)}$, can be determined by the CT information $(\tau_{ij}^{(d)}, d_{ij}^{(d)})$, as well as their infection status as follows:

$$\widehat{p}_{i,j}^{(d)} = \exp\Big(-\big(\lambda\, \tau_{ij}^{(d)}\, d_{ij}^{(d)}\, \Psi_{ij}^{(d)} + \epsilon\big)^{-1}\Big), \quad (4)$$

where $\Psi_{ij}^{(d)} = 1 - \widehat{\Pr}^{(d)}(X_i=0)\,\widehat{\Pr}^{(d)}(X_j=0)$, λ is a Poisson rate parameter, and ϵ is used to avoid division by zero. Note that $\widehat{p}_{i,j}^{(t)}$ depends on λ, which is unknown in practice. We estimate it using a plug-in approach by Ma et al. [29]. More details are given in Sec. 1.2 of the supplemental document.

Algorithms for multiplicative noise. For model **M2**, recall that $\boldsymbol{x}$ and $\boldsymbol{y}$ represent viral loads of individual samples and pools, respectively. The core algorithm presented in [8] uses the LASSO estimator, $\widehat{\boldsymbol{x}}^{\text{LASSO}} = \arg\min_{\boldsymbol{x}} \|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}\|_2^2 + \rho\|\boldsymbol{x}\|_1$ [30], where ρ is a smoothness parameter. LASSO exploits the sparsity of $\boldsymbol{x}$ but uses no SI. Despite the multiplicative nature of the noise, LASSO yields good estimation performance [8] for three measures: (*i*) *relative root mean squared error* (RRMSE) $= \|\boldsymbol{x} - \widehat{\boldsymbol{x}}\|_2/\|\boldsymbol{x}\|_2$; (*ii*) *false negative rate* (FNR) = #incorrectly detected negatives/#true positives; and (*iii*) *false positive rate* (FPR) = #incorrectly detected positives/#true negatives. Note that FNR $= 1-$ sensitivity and FPR $= 1-$ specificity.

In some cases, the n individuals in $\boldsymbol{x}$ can be partitioned into $n_1 \ll n$ disjoint groups of people, for example family members, who interact closely with each other and are thus likely to pass the virus between group members. This family-style structure leads to a situation where either all members of the group are uninfected, or a majority of members are infected. Note that the family-style structure also includes groups of coworkers, students taking a course together, and people sharing common accommodation. If reliable SI about how the n individuals are partitioned into families is available, and only a small portion of families, $n_2 \ll n_1$, are infected, then LASSO can be replaced by *group square-root* LASSO (SQRT-GLASSO) [21].[2] The latter is defined as

$$\widehat{\boldsymbol{x}}^{\text{SQRT-GLASSO}} = \arg\min_{\boldsymbol{x}} \|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}\|_2 + \rho\sum_{g=1}^{n_1}\|\boldsymbol{x}_g\|_2, \quad (5)$$

where $\boldsymbol{x}_g$ consists of viral loads of people from the gth family.

In other cases, accurate family SI may be unavailable or unreliable. Moreover, family SI does not account for contact between members of different families. In such cases, we use CT SI commonly available via Bluetooth [23] to directly *infer* family-style structure using clique detection algorithms; contacts between members of different families can also be considered to be small cliques. In particular, we use the Bron–Kerbosch algorithm [32] to find maximal cliques in the CT graph, and label each clique as a family. Note that one could generate these groups differently [33, Sec. 7], for example, decomposition into k-clique communities [34]. However, such a decomposition may partition the n individuals into $n_3 \ll n$ family structures that *overlap* with each other, unlike the earlier case of *disjoint* families. In a scenario with overlapping families, we use the *overlapping group square-root* LASSO (COMP-SQRT-OGLASSO) estimator [33],

$$\widehat{\boldsymbol{x}}^{\text{SQRT-OGLASSO}} = \arg\min_{\boldsymbol{x}} \|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}\|_2 + \rho\,\Omega_{\text{overlap}}(\boldsymbol{x}), \quad (6)$$

where $\Omega_{\text{overlap}}(\boldsymbol{x}) = \inf_{\boldsymbol{v}\in\mathcal{V}_G, \sum_{g\in G}\boldsymbol{v}_g=\boldsymbol{x}} \sum_{g\in G}\|\boldsymbol{v}_{\boldsymbol{g}}\|_2$, G denotes a set of possibly overlapping groups each containing a subset of the n individuals in $\boldsymbol{x}$, $\mathcal{V}_G$ is a set of $|G|$-tuples of vectors $\boldsymbol{v} = (\boldsymbol{v}_g)_{g\in G}$, g is an index for the groups in G, and $\boldsymbol{v}_g \in \mathbb{R}^n$ is a vector whose support is a subset of g. Advantages of OGLASSO over GLASSO for overlapping groups are summarized in Fig. 1 and Sec. 3 of [33].

[2] We observed that SQRT-GLASSO, which has an ℓ_2 data fidelity term instead of a squared ℓ_2 one [31], outperformed GLASSO. In contrast, conventional LASSO outperformed SQRT-LASSO.

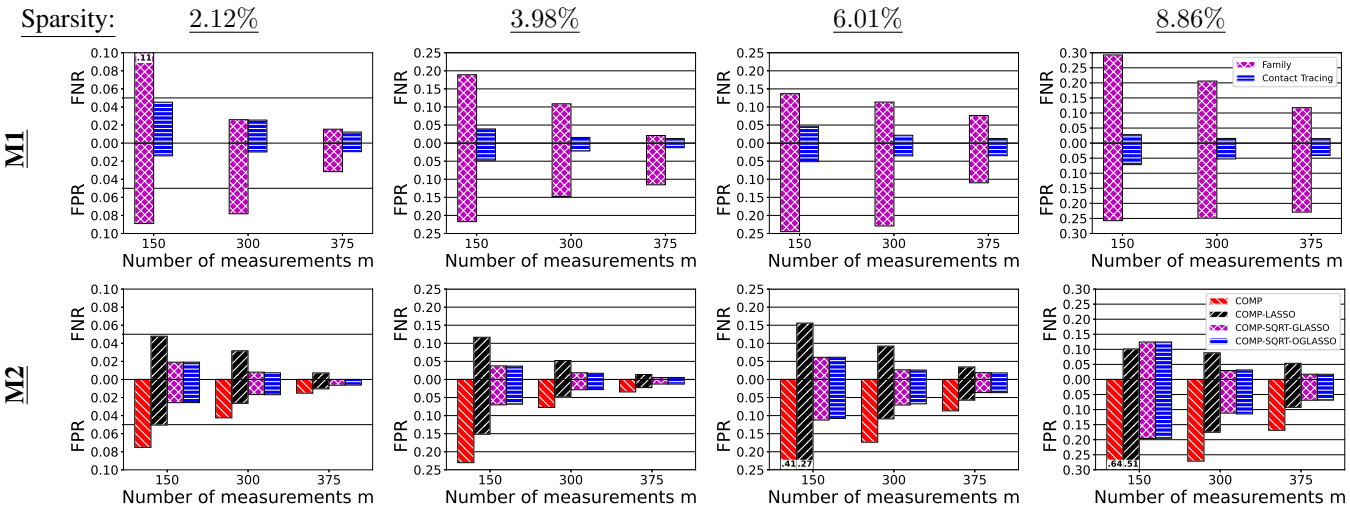

Fig. 3. Performance of the proposed group testing methods **M1** (top row) with binary noise and **M2** (bottom row) with multiplicative noise at four averaged sparsity levels and three measurement levels for a population of $n = 1000$ individuals.

All three algorithms included a non-negativity constraint on $\boldsymbol{x}$, and were preceded by a step that executed *combinatorial orthogonal matching pursuit* (COMP). COMP declares all samples that contributed to a pool with zero viral load to be negative. This COMP preprocessing step reduces the problem size and improves all three algorithms' performance, as well as speed. We refer to our algorithms as COMP-LASSO, COMP-SQRT-GLASSO and COMP-SQRT-OGLASSO.

4. NUMERICAL RESULTS

We now present numerical results obtained for models **M1** and **M2**. The data was generated based on Sec. 2, and group testing inference was performed using the algorithms proposed in Sec. 3. Since there are no publicly available datasets for pooling with associated SI, we believe that our results present strong (simulated) empirical evidence that motivates the collection of such SI for group testing in order to lower testing expenditures. We generated datasets using 4 levels of cross-clique contacts, leading to 4 averaged sparsity levels, 2.12%, 3.98%, 6.01%, and 8.86%, for $\boldsymbol{x}$. At each sparsity level, we perform pooling experiments using Kirkman triple matrices as proposed in [8] using $m \in \{150, 300, 375\}$. Measurement vectors $\boldsymbol{y}$ for M1 were generated using probabilities for erroneous binary tests, $\Pr(y_i = 1|w_i = 0) = 0.001$ and $\Pr(y_i = 0|w_i > 0) = 0.02$, per Hanel and Thurner [35]. Vectors $\boldsymbol{y}$ for M2 were generated by setting the parameter reflecting the strength of noise in RT-PCR to $\sigma^2 = 0.01$. Fig. 3 shows the performance of the proposed algorithms in terms of FNR and FPR[3] averaged across the inference results obtained for the time window of 50 days described in Sec. 2.

For model **M1**, we tested the family denoiser (2) and the CT denoiser (3). Fig. 3 shows that the CT denoiser outperforms the family denoiser in all settings. Both algorithms yield lower (better) FNR and FPR as the number of measurements, m, increases. Moreover, the CT denoiser's error rates are below 0.05, except for the challenging cases where the sparsity level is 8.86% and $m \in \{150, 300\}$.

For model **M2**, we tested four algorithms: COMP, COMP-LASSO, COMP-SQRT-GLASSO, and COMP-SQRT-OGLASSO. The results show that both COMP-SQRT-GLASSO and COMP-SQRT-OGLASSO outperform COMP-LASSO in terms of FNR and FPR, which shows the benefit of using CT SI. Note that COMP-SQRT-OGLASSO performs on par with COMP-SQRT-GLASSO, even though the former infers everything on the fly from CT SI without explicit access to family SI. The COMP algorithm by itself produces no false negatives (corresponding to FNR $= 0$), but many false positives. Further, all four algorithms yield lower (better) FNR and FPR as m increases or the averaged sparsity level decreases. Finally, we remark that COMP-SQRT-OGLASSO outperforms COMP-SQRT-GLASSO for more general CT graphs consisting of slightly incomplete cliques. We refer the readers to Sec. 3 of the supplemental document for details.

Our algorithms, when presented with SI, reduce the FNR and FPR, despite not knowing which individuals are infected within each infected family (within infected families, roughly 70% of individuals are infected, on average). Note that none of the algorithms for model **M2** make use of previous inference results, whereas the CT denoiser for model **M1** uses such information. This distinction makes the two approaches applicable in different scenarios, namely, the CT denoiser can be used for a CT and testing program where the population is tested regularly, e.g., warehouse employees, whereas COMP-SQRT-OGLASSO is useful when a population is tested once. Furthermore, while model **M1** performs well in the presence of erroneous binary tests, it does not yield viral load estimates as COMP-SQRT-OGLASSO does. Viral load estimates could prove to be useful, since there is a positive correlation between mortality and viral loads [36, 37].

[3]For M1, we chose to report the FPR and FNR pair such that the sum of the two error rates is minimized. Complete ROC curves are shown in Sec. 2.1 of the supplemental document [1]. For M2, we reported the error rates by thresholding the estimated viral load using $\tau = 0.2$. The error rates do not change much when τ varies between 0 and 1.

5. REFERENCES

[1] R. Goenka, S.-J. Cao, C.-W. Wong, A. Rajwade, and D. Baron, "Contact tracing enhances the efficiency of COVID-19 group testing," https://arxiv.org/abs/2011.14186, Nov. 2020.

[2] R. Dorfman, "The detection of defective members of large populations," *Annals of Mathematical Statistics*, vol. 14, no. 4, p. 436–440, 1943.

[3] M. Aldridge, O. Johnson, and J. Scarlett, *Group Testing: An Information Theory Perspective*, 2019.

[4] C. Hogan, M. Sahoo, and B. Pinsky, "Sample pooling as a strategy to detect community transmission of SARS-CoV-2," *Journal of the American Medical Association (JAMA)*, vol. 323, no. 19, pp. 1967–1969, Apr. 2020.

[5] B. Abdalhamid *et al.*, "Assessment of specimen pooling to conserve SARS CoV-2 testing resources," *American Journal of Clinical Pathology*, vol. 153, no. 6, pp. 715–718, May 2020.

[6] J. Zhu, K. Rivera, and D. Baron, "Noisy pooled PCR for virus testing," https://arxiv.org/abs/2004.02689, Apr. 2020.

[7] J. Yi, R. Mudumbai, and W. Xu, "Low-cost and high-throughput testing of COVID-19 viruses and antibodies via compressed sensing: System concepts and computational experiments," https://arxiv.org/abs/2004.05759, Apr. 2020.

[8] S. Ghosh, R. Agarwal, M. A. Rehan, S. Pathak, P. Agrawal, Y. Gupta, S. Consul, N. Gupta, R. Goyal, A. Rajwade, and M. Gopalkrishnan, "A compressed sensing approach to group-testing for COVID-19 detection," https://arxiv.org/abs/2005.07895, May 2020.

[9] J. Zhu, K. Rivera, C. Rush, and D. Baron, "Noisy pooled PCR for COVID-19 testing," *Paris Machine Learning Meetup*, May 2020. [Online]. Available: https://youtu.be/gYJqnXbi1Bg

[10] A. Heidarzadeh and K. Narayanan, "Two-stage adaptive pooling with RT-qPCR for COVID-19 screening," https://arxiv.org/abs/2007.02695, Jul. 2020.

[11] P. Nikolopoulos, T. Guo, C. Fragouli, and S. Diggavi, "Community aware group testing," https://arxiv.org/abs/2007.08111, Jul. 2020.

[12] N. Shental *et al.*, "Efficient high throughput SARS-CoV-2 testing to detect asymptomatic carriers," *Science Advances*, vol. 6, no. 37, Sep. 2020.

[13] Y.-J. Lin, C.-H. Yu, T.-H. Liu, C.-S. Chang, and W.-T. Chen, "Comparisons of pooling matrices for pooled testing of COVID-19," https://arxiv.org/abs/2010.00060, Sep. 2020.

[14] ——, "Positively correlated samples save pooled testing costs," https://arxiv.org/abs/2011.09794, Nov. 2020.

[15] P. Nikolopoulos, S. R. Srinivasavaradhan, T. Guo, C. Fragouli, and S. Diggavi, "Group testing for overlapping communities," https://arxiv.org/abs/2012.02804, Dec. 2020.

[16] S. Ahn, W.-N. Chen, and A. Ozgur, "Adaptive group testing on networks with community structure," https://arxiv.org/abs/2101.02405, Jan. 2021.

[17] B. Arasli and S. Ulukus, "Group testing with a graph infection spread model," https://arxiv.org/abs/2101.05792, Jan. 2021.

[18] D. Benatia, R. Godefroy, and J. Lewis, "Estimating COVID-19 prevalence in the United States: A sample selection model approach," https://doi.org/10.1101/2020.04.20.20072942.

[19] Center for Disease Control and Prevention, "Contact tracing for COVID-19," https://www.cdc.gov/coronavirus/2019-ncov/php/contact-tracing/contact-tracing-plan/contact-tracing.html.

[20] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *IEEE Int. Symp. Inf. Theory*, 2011, pp. 2168–2172.

[21] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *J. Royal Stat. Soc. Series B*, vol. 68, no. 1, 2006.

[22] D. Jacot, G. Greub, K. Jaton, and O. Opota, "Viral load of SARS-CoV-2 across patients and compared to other respiratory viruses," *Microbes and Infection*, 2020.

[23] A. Hekmati, G. Ramachandran, and B. Krishnamachari, "CONTAIN: Privacy-oriented contact tracing protocols for epidemics," https://arxiv.org/abs/2004.05251.

[24] W. Haseltine, "What COVID-19 reinfection means for vaccines," https://www.scientificamerican.com/article/what-covid-19-reinfection-means-for-vaccines/.

[25] J. M. Carcione, J. E. Santos, C. Bagaini, and J. Ba, "A simulation of a COVID-19 epidemic based on a deterministic SEIR model," *Frontiers in Public Health*, vol. 8, p. 230, 2020.

[26] World Health Organization, "Coronavirus disease 2019 (COVID-19) situation report–73," https://tinyurl.com/ybnbky8m.

[27] B. Buchan *et al.*, "Distribution of SARS-CoV-2 PCR cycle threshold values provide practical insight into overall and target-specific sensitivity among symptomatic patients," *American Journal of Clinical Pathology*, 2020.

[28] United Nations, "Household size and composition around the world 2017," https://tinyurl.com/vo7hrlv.

[29] Y. Ma, J. Tan, N. Krishnan, and D. Baron, "Empirical Bayes and full Bayes for signal estimation," in *Inf. Theory App. Workshop*, San Diego, CA, Feb. 2014, pp. 994–1001.

[30] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical Learning with Sparsity: The LASSO and Generalizations*. CRC Press, 2015.

[31] A. Belloni, V. Chernuzhukov, and L. Wang, "Square-root LASSO: Pivotal recovery of sparse signals via conic programming," *Biometrika*, vol. 98, no. 4, pp. 791–806, 2011.

[32] C. Bron and J. Kerbosch, "Algorithm 457: Finding all cliques of an undirected graph," *Commun. ACM*, vol. 16, no. 9, p. 575–577, 1973.

[33] L. Jacob, G. Obozinski, and J.-P. Vert, "Group LASSO with overlap and graph LASSO," in *ICML*, 2009.

[34] G. Palla, I. Derényi, I. Farkas, and T. Vicsek, "Uncovering the overlapping community structure of complex networks in nature and society," *Nature*, vol. 435, pp. 814–818, 2005.

[35] R. Hanel and S. Thurner, "Boosting test-efficiency by pooled testing strategies for SARS-CoV-2," https://arxiv.org/abs/2003.09944, Mar. 2020.

[36] L. F. Westblade *et al.*, "SARS-CoV-2 viral load predicts mortality in patients with and without cancer who are hospitalized with COVID-19," *Cancer Cell*, Nov. 2020.

[37] E. Pujadas *et al.*, "SARS-CoV-2 viral load predicts COVID-19 mortality," *Lancet Respiratory Medicine*, vol. 8, no. 9, p. e70, Sep. 2020.