# Safe Blues: the case for virtual safe virus spread in the long-term fight against epidemics

# Raj Dandekar, Shane G. Henderson, Hermanus M. Jansen, Joshua McDonald, Sarat Moka, Yoni Nazarathy, Christopher Rackauckas, Thomas M. Stace, Peter G. Taylor, Appli Vuorinen, Appli Vuorinen, Vuorinen, Christopher Rackauckas, Karata Kar

September 19, 2020.

**THE BIGGER PICTURE:** Viral spread is a complicated function of multiple elements including biological properties, preventative measures such as sanitation and masks, the environment, and the level of physical proximity. It is this last element that governments can control through social-distancing directives. However, with a pandemic such as COVID-19, the data is always lagging and biased since the time between a patient being infected with SARS-CoV-2 and being recorded as positive can be a week or more. Consequently there can be a time lag of the order of several weeks between the initiation of a regulatory measure and its observed effect. There is thus a pressing need for real time information on the level of physical proximity while respecting personal privacy.

Safe Blues fills this need, providing real time population-level estimates of the level of physical proximity and near-future projections of the epidemic. Safe Blues strands are safe virtual 'virus-like' tokens that are spread using cellular devices and Bluetooth. Real time counts of multiple forms of the tokens are combined with delayed measurements of the actual epidemic. Then using machine-learning techniques the Safe Blues system creates more accurate projections about the current and near-future state of the epidemic.

The Safe Blues protocol and machine learning techniques have been developed together with an experimental minimal viable product, presented as an app on Android devices with a server back-end. Following initial exploration via simulation experiments, we are now preparing for a university-wide experiment of Safe Blues.

Author order is alphabetical.

<sup>\*</sup>Massachusetts Institute of Technology, rajd@mit.edu.

<sup>&</sup>lt;sup>†</sup>Cornell University, sgh9@cornell.edu.

<sup>&</sup>lt;sup>‡</sup>Delft University of Technology, h.m.jansen@tudelft.nl.

 $<sup>\</sup>fi$  The University of Queensland, joshua.mcdonald@uq.net.au.

<sup>&</sup>lt;sup>¶</sup>The University of Queensland, s.babumoka@uq.edu.au.

The University of Queensland, y.nazarathy@uq.edu.au (corresponding author).

<sup>\*\*</sup>Massachusetts Institute of Technology, contact@chrisrackauckas.com.

<sup>&</sup>lt;sup>††</sup>The University of Queensland, stace@physics.uq.edu.au.

 $<sup>^{\</sup>ddagger\ddagger}$  The University of Melbourne, taylorpg@unimelb.edu.au.

 $<sup>\</sup>$  Columbia University, contact@aapelivuorinen.com.

# Introduction

The COVID-19 pandemic is the most significant global event faced by humanity in the 21st century. In less than a year, there have been over 30 million confirmed cases, and nearly a million deaths. In trying to mitigate this virus, economies are crumbling and whole societies are undergoing transformation due to the new norms dictated by the virus and human response.

In parallel with the race to develop SARS-CoV-2 vaccines, there are many efforts in the scientific community to understand the spread of the virus through measurements and modelling. Viral spread is a complicated function of multiple elements including biological properties, preventative measures such as sanitation and masks, the environment, and very importantly the level of physical proximity. It is this last element, the level of physical proximity, that governments can control via social-distancing directives and lockdowns. However, with a pandemic such as COVID-19, the data we have is always lagging and biased: the time between a patient being infected and being recorded as positive can be one or two weeks. A consequence is that the typical time between the initiation of a lockdown measure and our observation of its effect is of the order of several weeks. This delay hinders the ability of epidemiologists, mathematical biologists, and public-health officials to make inferences about the current situation and projections about the future trajectory of viral spread. As a consequence, governments struggle to deliver an effective response to the pandemic.

Biological properties of SARS-CoV-2 have been studied since the start of the outbreak<sup>27</sup> and are being better understood as time progresses. On the other hand, population behavior is changing rapidly due to unprecedented social distancing regulation and is hard to observe, model, and predict. As a consequence, achieving tight real time estimates of time-varying parameters such as  $R_{\rm eff}(t)$ , the expected number of individuals infected by an infectious person, is a difficult task.<sup>23</sup>

The effects of this uncertainty are painfully visible in the second waves afflicting multiple countries that responded "well" with strong social-distancing measures in February-May, 2020. As an illustration, consider Figure 1 presenting the confirmed daily new cases in the state of Victoria, Australia. The initial lockdown during April-May managed to suppress COVID-19. However a second outbreak in July required a further lockdown and the magnitude of the second wave was much more significant than the first. In managing this crisis, government faces extreme difficulty in determining the optimal level of social distancing that should be imposed. Such decisions involve guesswork because there is no immediate feedback on how various levels of social-distancing regulation affect physical proximity and, in turn, viral spread.

While in general, the Victorian response appears to have suppressed the spread of SARS-CoV-2 in both waves, the social distancing measures have had huge economic and social costs. The question is then what can be improved in future outbreaks or waves? A swifter response entailing the imposition of social-distancing measures might have allowed an earlier easing of those measures with fewer overall infections and ultimately lower economic and human impact. However, due to a lack of information with current measurement techniques, such a refined response can not be expected. Essentially, government does not have real time information of viral spread to input into its decision making process.

To provide more real time information, we propose the use of "Safe Blues".<sup>1,3</sup> Safe Blues uses Bluetooth signals in a manner similar to existing and emerging contact tracing frameworks such as Proximity,<sup>11</sup> Blue Trace,<sup>5</sup> and the Privacy-Preserving Contact Tracing framework developed by Apple and Google.<sup>4</sup> In contrast to these frameworks, Safe Blues does not store information about individuals and their interactions. Instead, it mimics virus spread via the safe exchange of Bluetooth signals in a real time privacy-preserving manner. Various Safe Blues virus strands are periodically created and spread through the (mobile device) population. Aggregated counts for each strand are reported to a server without recording private information. An analysis of the counts helps to obtain aggregate estimates of population contact. The result is a real time estimate of the effect of any social distancing rules that are in place. The effect of other events such as public demonstrations and mass gatherings



Figure 1: Measured daily new cases of COVID-19 in the Australian state of Victoria, during the first half of 2020. Both the first wave and the second wave were mitigated via social distancing measures. Government imposed various social distancing directives, with severity labeled via stage 2 (lightest), stage 3, and stage 4 (complete lockdown).

is also automatically measured. This information on physical proximity, along with retrospective information about actual (not Safe Blues) SARS-CoV-2 case numbers, is then used to train sophisticated machine-learning (ML) procedures to estimate SARS-CoV-2 infection numbers as a function of Safe Blues strands. Real time information on Safe Blues strands then provides, via the ML projections, live near-future estimates of SARS-CoV-2 infection levels.

In the sense that its outcomes are driven by the behaviour of individuals, Safe Blues shares some characteristics with agent-based simulation.<sup>15</sup> Both rely on capturing realistic physical proximity behavior of individuals along with virus transmission to make predictions about population outcomes. The difference between them is that Safe Blues uses a sample from the actual population whereas agent-based models simulate this. The latter have the advantage that they can be used to explore potential future scenarios and can incorporate the effect of interventions such as mask-wearing and surface-cleaning. However, they still rely on assumptions about how agents interact rather than the actual interactions. Ultimately we believe that policy makers should use a mix of agent-based simulations as in Victoria,<sup>8</sup> and virtual safe virus spread solutions, such as Safe Blues, since these tools complement one another.

#### Measuring and Controlling an Epidemic

Epidemics are largely governed by the average number of individuals infected by each sick individual. This quantity, usually referred to as the effective reproduction number  $R_{\text{eff}}(t)$  at time t, measures the ebbs and flows of the epidemic in aggregate. For SARS-CoV-2, early estimates indicate that, without significant control measures in place,  $R_{\text{eff}}(t)$  lies in the range of 2 - 4.<sup>19</sup> Its magnitude depends on a combination of biological and behavioural factors. Key biological factors include the propensity of the pathogen to infiltrate human hosts, the duration of the disease in an infected individual, and the susceptibility of different age groups. Key behavioural factors include personal hygiene practices, cultural practices around touching such as hand shaking and, importantly, the structure of social networks and amount of time that individuals spend in close proximity to each other.

The biological factors determining  $R_{\rm eff}(t)$  tend to be uncontrollable and, with the exception of weather effects, may generally be assumed to remain constant over time as long as significant virus mutation does not occur. The behavioural factors, however, are controllable to some extent. Indeed, the social distancing measures imposed in over 150 countries during the first few months of 2020, some of which are outlined in the high-impact report,<sup>12</sup> are attempts to control the behavioural component of  $R_{\rm eff}(t)$ .<sup>13</sup> Retrospectively, it is clear that these measures were essential for slowing down the epidemic, but it is also clear that their effectiveness decreased over time as people became less cautious and were unaware of a growing second wave (see Figure 1). Unfortunately, at this early stage, it was very difficult to quantify in real time the effect that any particular social distancing measure had on population behavior or network structure, and in turn on the evolution of  $R_{\rm eff}(t)$  and the course of the epidemic.

Some proxies for social proximity are available through anonymized, passively generated mobile phone data,<sup>21</sup> and this data has shown some predictive power for outbreaks of other communicable diseases.<sup>6</sup> However, this type of information lacks high spatial resolution so may require parameter calibration that depends on epidemiological data that are only available retrospectively.

A lack of near real time information on social proximity is problematic, because all models that aid policy makers by projecting the course of the epidemic require an estimate of  $R_{\rm eff}(t)$ . This is often obtained by modelling that attempts to quantify the level of human-to-human interaction either at broad scales<sup>20, 23</sup> or at finer scales.<sup>9, 12, 16</sup> A notable recent effort to obtain such a quantification<sup>17</sup> used survey sampling of the UK population to estimate that in March of 2020  $R_{\rm eff}(t)$  shifted from around 2.6 prior to lockdown, to around 0.62 after lockdown. While impressive, such questionnaire-based surveys are difficult and expensive to execute, and do not yield real time estimates of  $R_{\rm eff}(t)$ . Other efforts to measure  $R_{\rm eff}(t)^{18, 24}$  use up-to-date count data such as the now famous dashboard by the CSSE at Johns Hopkins.<sup>10</sup> However, in such cases the problem is that reported "live data" about SARS-CoV-2 is based only on confirmed tested cases which lags infection events and does not consider the large number of asymptomatic cases or otherwise untested cases. Better estimates become known only retrospectively, after the pandemic has progressed.

In a controlled system, lags in a system's response (in this case, virus transmission) to a control intervention (such as increased social restrictions) are problematic. They lead to undesirably slow control updates, which hinder policy makers' ability to effectively reduce transmission while minimising wider social and economic disruptions. Safe Blues provides a real time proxy for the virus transmission response, thereby reducing the lag between social intervention and the epidemiological response.

#### Virtual Safe Virus Measurement

The key idea of Safe Blues is to obtain real time estimates of gross population engagement dynamics in a safe and privacy-preserving manner. Safe Blues data can be processed in real time to yield estimates of the spread of a virus such as SARS-CoV-2. Near-future projections and estimates of the effective reproductive number  $R_{\text{eff}}(t)$  can be continuously updated. At the heart of the system is a measurement framework of auxiliary variables describing the spread of Safe Blues strands which we describe below.

The system works by having personal mobile devices take part in an ongoing safe real time virus



Figure 2: Individuals of the population with Safe Blues enabled devices take part in spreading Safe Blues strands. SARS-CoV-2 infected individuals are in red and others are in green, however the Safe Blue system operates independently of the health status of individuals.

spread "simulation" where, by means of Bluetooth signals, the time that individuals spend in close proximity is a key driving factor. This is done in a way that does not compromise individual privacy, does not cause any risk to human health, and does not introduce any risk to individual software or hardware. See Figure 2 for a schematic illustration of the Safe Blues data collection system.

A Safe Blues strand is a virtual token that circulates and replicates between the mobile devices of individuals using dynamics designed to reflect the transmission of a biological virus but without any threat to safety, software, or privacy. Strands differ in their viral properties, such as incubation time and level of infectiousness. Strands are counted as "active" for a finite duration of time in each mobile device that is "infected." During that time, if the mobile device is in close proximity to another device, there is a chance for the strand to "spread" to the neighbouring device. Similarly, if the mobile device is in relative isolation, the strand is not likely to spread and the mobile device is unlikely to be "infected" by Safe Blues strands.

We allow multiple strands to be present on each mobile device. These strands evolve independently from one another, so in effect we are "simulating" multiple "epidemics," corresponding to the multiple strands, at any given time. Unlike with biological epidemics, the number of devices infected by each strand can be measured in real time.

The Safe Blues system periodically injects strands into the mobile host population and obtains real time counts of the number of infected hosts for each strand. While the population dynamics of each strand do not necessarily directly resemble the dynamics of SARS-CoV-2 or any other biological virus, the epidemics of the strands all respond to social proximity and social distancing measures in a similar way to a biological virus, because they all experience the *same* social proximity and social distancing measures. Hence, we expect the course of strand "epidemics" to be coupled with the course of the COVID-19 epidemic.

The mechanism of communication between devices is Bluetooth. This is similar to the communication protocol used by many emerging contact tracing apps. In the Safe Blues protocol, the probability of transmission increases as individuals spend time in close proximity. Conversely, as individuals maintain a higher level of social distancing, the Safe Blues strands are less likely to spread. Unlike systems that use contact tracing apps, the Safe Blues system is oblivious to the actual health status of specific individuals. For example, in Figure 2, some individuals are infected by SARS-CoV-2 (red) while others are not (green). However, Safe Blues is not aware of, and does not need, this private information. Similarly, some individuals participate in Safe Blues (as signified by a those holding a mobile device in the figure) and others do not. Clearly some level of population participation is required, but Safe Blues does not require all individuals to participate.

#### **Evaluation of Safe Blues via Simulation**

For initial evaluation of Safe Blues we created three different simulation models, each of a different nature. This allows us to test the robustness of the system to different realities. In these simulation models, we simulate both the biological virus and Safe Blues strands. Full details of the models are in Appendix C of a prior study.<sup>1</sup> We now review the key ideas.

All models feature a population comprising N individuals. Some of these individuals have Safe Blues-enabled devices, while others do not. At each point in time, the state of an individual registers whether they are susceptible, infected, or removed with respect to the actual virus. If an individual has a Safe Blues enabled-device, then the state also registers, for every Safe Blues strand, whether they are susceptible, infected, or removed. Hence an individual's simulated device maintains the state of multiple Safe Blues strands simultaneously.

The three models differ in their complexity and how they capture individual proximity. However, regardless of the model, individual proximity drives both the (simulated) SARS-CoV-2 spread and the Safe Blues spread in a coupled manner because both SARS-CoV-2 and Safe Blues strands only spread when individuals are in close proximity. This roughly approximates what one may expect to happen in reality. Importantly, all three models allow for time-varying parameters that enforce social distancing, which in turn affects both SARS-CoV-2 spread and Safe Blues spread by changing how much time individuals spend in close proximity of each other.

Figure 3 sketches the three models. Model I is a very simple and stylised discrete-time stochastic model. One of its appealing features is that it converges to the well-known SIR difference equations as the population size N becomes large, which makes this model well suited as a first test bed. Model II is a continuous-time stochastic SIR model with migration. It incorporates several social and spatial features that are ignored in Model I. In particular, Model II has a spatial component (people have to be in the same place at the same time for virus transmission to occur) as well as a notion of social levels (people have a home where they meet a selected number of other people, a workplace where they may meet a larger number of people, et cetera). Model III is a spatial movement model with location attraction in which individuals move randomly in two-dimensional space. Its distinguishing feature is that it has a notion of centrality: although individuals move around randomly, they are biased towards visiting places that are important for them, such as their home and the supermarket. This creates a form of clustering that is not present in the first two models.

In the simulations that we ran with each model (Figure 4), after an initially rapid spread of the epidemic, social distancing regulations are increasingly imposed over a period of two weeks and then fixed to prevent the vast majority of social contact. As these rules go into effect, the infected proportion begins a slow but steady decline. All the while, Safe Blues strands are being transmitted on mobile devices. As a consequence, the proportion of Safe Blues infections mirrors the decline in the true



Model III

Figure 3: Model I: At every time point, each of the N individuals selects a random number of other individuals to "invite" and this implies physical proximity. In this case orange and green individuals make invitations. Model II: All  $N = 2^k$  individuals traverse a binary tree between their private leaf and the root. At any node, infection follows a continuous-time stochastic SIR model between the individuals present. Model III: A spatial model where each individual diffuses either around their base (e.g. home) or around a center (e.g. supermarket).

infected proportion, driven by a corresponding reduction in physical proximity between Safe Bluesenabled devices. In all the runs, at around 100 simulated days after the start, and months of a promising decline in case numbers, the social distancing rules are mostly lifted, resulting in a second wave. The goal of our numerical experiments is to see whether Safe Blues can provide adequate information to predict this second wave.

# Prediction and Real Time Information with Safe Blues

For performing real time projection from the Safe Blues strand information, we created the *Deep Safe Blues* deep neural network model<sup>14</sup> from the ensemble of Safe Blues strands and historical SARS-CoV-2 information. The basic setup follows the paradigm presented in Figure 4 where accurate information of SARS-CoV-2 is available only up to a certain time, say two weeks in the past, after which only Safe Blues information is available. This lag represents the fact that SARS-CoV-2 information is not present in real time, in contrast to Safe Blues. The relative magnitude of historic social distancing measures is also available as input for projections (for example "full lockdown", "partial lockdown", etc.).

The projections displayed in Figure 4 demonstrate that Deep Safe Blues is able to accurately detect

the start of a trend towards a second peak in the number of infected individuals a significant time before such data is available in all three models. Importantly, the same neural network architecture was used for all three simulation models, meaning no tuning of architectures is required to achieve these results. Together, our results yield confidence in the ability of Safe Blues to detect a second wave before it shows up in the actual data. This can enable public health officials to respond during the essential early period before infection estimates can be updated.

This analysis demonstrates one of many methods to uncover the predictive power that can be gained by utilizing Safe Blues information. Contact tracing apps can allow public health investigators to identify specific individuals who might be infected, but the focus on the population of known infected individuals means that contact-tracing data on infection proportions is heavily biased relative to the general population. In contrast, the Safe Blues system is able to determine the status of the entire population of participating devices whether or not they are infected.



Figure 4: *Deep Safe Blues:* Safe Blues detection of a second wave applied to data generated from three different simulation models. The proportion of infected individuals is only known until the vertical black lines. After that point, only Safe Blues information is available. Nevertheless, Deep Safe Blues (trained up to the black line) is able to accurately predict a second wave of SARS-CoV-2 infections.

#### Feedback to Policy with Safe Blues

In addition to being a tool for estimating the current number of infected individuals before such data is available, Safe Blues can also help estimate the potential effect of policy decisions. For this we developed *Dynamic Deep Safe Blues* which is a tool for projecting  $R_{\text{eff}}(t)$  as a function of future levels of social distancing using Universal Ordinary Differential Equations.<sup>25</sup> This approach mixes neural networks into epidemiological models in order to directly learn how policy decisions affect the spread of Safe Blues and the actual infection. Figure 5 demonstrates data-driven projections of  $R_{\text{eff}}(t)$  under various policy levels. Such results can be used as input to decision makers for helping to determine the levels necessary to contain the outbreak and ensure that exponential growth into a second peak does not occur. By directly quantifying the effectiveness of interventions over time, this technique can be related back to historical policy decisions to determine the minimal level of social controls required to achieve declining cases and thus prevent further disease outbreak.



Figure 5: Demonstration of policy projection and refinement using *Dynamic Deep Safe Blues* on Model III. We project the effect of the social distancing strength on  $R_{\text{eff}}(t_0)$  where  $t_0$  is the end of the training period (right). We also demonstrate potential near-future trajectories as a function of policy decision (left).

# A Campus Experiment for Safe Blues

Following our simulations, we are now in the process of devising a campus-wide experiment of Safe Blues. We have developed an Android application based on the open source contact tracing OpenTrace software<sup>5</sup> published by the Singaporean government. Our app follows the Safe Blues protocol specified in Appendix A of a prior study<sup>1</sup> with additional implementation details. The app source code is written in the Kotlin language, and interfaces with a backend service using the gRPC protocol. We have developed an associated dashboard and control panel to aid in creating new Safe Blues strands and tracking their spread. The source code for the app, the backend, and the dashboard is available under the Safe Blues GitHub organisation.<sup>2</sup> See Figure 6 for an illustration of the app and the control panel.

The main purpose of the Safe Blues campus experiment is to test if the Safe Blues system can add to the predictive power of traditional epidemic modelling forecasting methods. The experiment will be carried out at The University of Queensland, St Lucia campus. The experiment will source student participants who will run the Safe Blues app as they attend campus. A randomized reward mechanism will be used to ensure participants continue to use the app over a prolonged period. As participants run the Safe Blues app, they will propagate an ensemble of *red strands* and an ensemble of *blue strands*. Each red strand simulates an epidemic that is to be tracked. With multiple red strands, the experiment can simultaneously simulate multiple epidemics. The blue strands are Safe Blues. During the experiment, they are tracked to obtain quantifiable uncertainties of the predictive power of Safe Blues versus traditional methods.

The Safe Blues predictor will have available real time Safe Blues strand information as well as delayed red strand measurements. Separately, for each red strand, the Deep Safe Blues method described above will be used to project the current state of the red epidemic, as well as a 10 day forecast. In parallel, several state-of-the-art predictive techniques<sup>7</sup> will be employed to project infection levels of the red epidemic using only suitably lagged red strand historical data. The goal of the experiment is to test whether the projections obtain with the aid of Safe Blues can yield significant improvement in the accuracy of projections that one would obtain without Safe Blues, only based on lagged red-strand data.

The experiment will take place in a small campus environment with only a few hundred active participants, so the parameters of the Safe Blues app will need to be adjusted beforehand. Such a calibration phase is needed to ensure that both red strands and Safe Blues strands can propagate efficiently, so that red/blue epidemics are possible. This phase may also involve agent-based simulations to aid in calibration.

After the calibration phase and during the experiment, social distancing will be simulated by artificially controlling the infectiousness of the red and blue strands. Such artificial control enables one to simulate varying degrees of social-distancing measures. The goal is to have a first wave (of red and blue strands), and then in the second wave, Safe Blues measurements will be employed for prediction. Further updates about the Safe Blues campus experiment will be posted on the Safe Blues website.<sup>3</sup>



Figure 6: The control panel, dashboard and Android app - used for the planned campus experiment.

### Outlook

A century of technological innovations passed between the 1918-1920 H1N1 pandemic and the current COVID-19 pandemic. One can then ask what innovations are helping humanity fight the current pandemic that were not available in 1918? In addition to many advances in the understanding of pathogens and their effect on the human body, some of the innovations used for quantifying and managing the 2020 pandemic include compartmental (SIR type) differential equation epidemic models, agent-based simulation models, contact-tracing apps, and a variety of data collection and visualization mechanisms. However, the past few decades have witnessed great advances in computational statistics and machine learning, and one may ask if tools from these areas have so far been effectively employed in the fight against COVID-19. We believe that to date, the answer is generally negative due to the lack of available coherent data associated with the pandemic. It is thus unfortunate that a pandemic hitting humanity in the midst of the "AI revolution" cannot be controlled using the machine learning tools that are celebrated in so many other domains.

The value proposition of methods such as Safe Blues is to empower modellers and policy makers to utilize computational statistics and machine learning effectively. Our method allows one to collect additional real time information at scale. Indeed, large quantities of data are often needed for effective application of deep learning and related methods. Safe Blues aims to provide policy makers and modellers with ample information that informs better decision making. To the best of our knowledge this type of framework is fundamentally different from existing solutions and other suggestions that have appeared in the literature.<sup>22</sup>

Based on simulation experiments, the machine learning principles and analysis that we use appear to be robust enough to yield immediate value from collected Safe Blues signals. The next step is to collect results using our experimental App in a campus level experiment.

A further critical attribute of any such system is privacy. The Safe Blues system is highly privacypreserving relative to many other apps that aim to fight the COVID-19 epidemic. No individual interaction information or any other private information is shared between devices or between a device and the server. This is achieved by not associating long-term identifiers with users or devices, nor collecting any information about the users themselves. The devices do not share anything other than the strands with which they are currently infected.

This is in contrast to contact-tracing apps that raise more serious concerns about personal privacy, even when engineered using novel privacy-protecting methods. Fundamentally, this is because the goal of any contact-tracing app is to observe relationships between individual people through their interactions, whereas the goal of Safe Blues is to collect only aggregate simulated epidemic signals. Thus a proper implementation of Safe Blues can provide stronger privacy guarantees than apps attempting to mitigate SARS-CoV-2 through contact-tracing methods.<sup>5, 26</sup>

Nonetheless, there are some issues that need to be addressed. For example, the case of an adversary choosing a rare strand and infecting a user with that strand to track them is averted by making sure that the seeding probability of each strand is sufficiently large to make strands common, making them meaningless for trying to identify an individual. Still, tracking might be possible using *combinations* of strands. Addressing such concerns remains a goal for future research.

An idea we have not yet pursued is to take the Safe Blues idea further, and consider gamification. For example, in a second or third generation of the app, one might consider presenting users with an up-to-date count of the strands of Safe Blues infecting their device. This information may help users to get a feel for the level of social distancing that they are practising and to stay socially responsible as advised by government. One might even use randomized rewards to further incentivize users to maintain social distancing.

# Acknowledgments

We thank James McCaw, Yoav Banin, and Yun William Yu for insights. We thank Toshali Banerjee for help with illustrations. HMJ and YN are supported by Australian Research Council (ARC) under grant number DP180101602. SM and PGT are supported by the Australian Research Council (ARC) Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS) under grant number CE140100049. SH is supported by the Army Research Office under grant number W911NF-17-1-0094 and the National Science Foundation under grant number TRIPODS+X DMS-1839346.

#### References

- R. Abhijit Dandekar, S. G. Henderson, M. Jansen, S. Moka, Y. Nazarathy, C. Rackauckas, P. G. Taylor, and A. Vuorinen. Safe Blues: A method for estimation and control in the fight against COVID-19. medRxiv, May 2020. DOI:10.1101/2020.05.04.20090258.
- [2] R. Abhijit Dandekar, S. G. Henderson, M. Jansen, S. Moka, Y. Nazarathy, C. Rackauckas, P. G. Taylor, A. Vuorinen, T. Stace, and J. McDonald. Safe Blues github. https://github.com/SafeBlues.
- [3] R. Abhijit Dandekar, S. G. Henderson, M. Jansen, S. Moka, Y. Nazarathy, C. Rackauckas, P. G. Taylor, A. Vuorinen, T. Stace, and J. McDonald. Safe Blues project homepage. https: //safeblues.org/.
- [4] Apple Inc and Google Inc. Privacy-preserving contact tracing. https://www.apple.com/covid19/ contacttracing, Apr. 2020.
- [5] J. Bay, J. Kek, A. Tan, C. S. Hau, L. Yongquan, J. Tan, and T. A. Quy. Bluetrace: A privacypreserving protocol for community-driven contact tracing across borders. https://bluetrace.io/, Apr. 2020. Government Technology Agency, Singapore.
- [6] L. Bengtsson, J. Gaudart, X. Lu, S. Moore, E. Wetter, K. Sallah, S. Rebaudet, and R. Piarroux. Using mobile phone data to predict the spatial spread of cholera. *Nature, Scientific Reports*, 5(8923), Mar. 2015. DOI:10.1038/srep08923.
- [7] J. E. Bickel and C. Small. COVID-19: What is the data telling us? https://www.youtube.com/ watch?v=wyi3XiWpIQY, Aug. 2020. TexTalks, The University of Texas at Austin.
- [8] T. Blakely, J. Thompson, Ν. Carvalho, L. Bablani, Ν. Wilson, and М. Stevenson. VIC agent model. https://www.mja.com.au/journal/2020/ probability-6-week-lockdown-victoria-commencing-9-july-2020-achieving-elimination, 2020.
- [9] S. Chang, N. Harding, C. Zachreson, O. Cliff, and M. Prokopenko. Modelling transmission and control of the COVID-19 pandemic in Australia. arXiv:2003.10218, Apr. 2020.
- [10] E. Dong, H. Du, and L. Gardner. An interactive web-based dashboard to track COVID-19 in real time. The Lancet, Infectious Diseases, 20(5):553–534, 2020.
- [11] M. Faggian, M. Urbani, and L. Zanotto. Proximity: a recipe to break the outbreak. arXiv:2003.10222, Apr. 2020.

- [12] N. Ferguson, D. Laydon, G. Nedjati Gilani, N. Imai, K. Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, Z. Cucunuba Perez, G. Cuomo-Dannenburg, et al. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. https://www.imperial.ac.uk/mrc-global-infectious-disease-analysis/covid-19/ report-9-impact-of-npis-on-covid-19/, Mar. 2020.
- [13] I. M. Fund. Policy Tracker: Policy responses to COVID-19. https://www.imf.org/en/Topics/ imf-and-covid19/Policy-Responses-to-COVID-19, Apr. 2020.
- [14] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. The MIT Press, 2016.
- [15] N. Hoertel, M. Blachier, C. Blanco, M. Olfson, M. Massetti, M. S. Rico, F. Limosin, and H. Leleu. A stochastic agent-based model of the SARS-CoV-2 epidemic in france. *Nature Medicine*, pages 1–5, 2020. DOI:10.1038/s41591-020-1001-6.
- [16] IISc-TIFR Covid-19 City-Scale Simulation Team. City-Scale Epidemic Simulator on Github. https://cni-iisc.github.io/epidemic-simulator/, 2020.
- [17] C. I. Jarvis, K. Van Zandvoort, A. Gimma, K. Prem, P. Klepac, G. J. Rubin, and W. J. Edmunds. Quantifying the impact of physical distance measures on the transmission of COVID-19 in the UK. BMC Medicine 18, 1-10. *BMC Medicine*, 18(124), 2020. DOI:10.1186/s12916-020-01597-8.
- [18] M. Lavielle. Modelling some COVID-19 data. http://webpopix.org/covidix19.html, Apr. 2020.
- [19] Y. Liu, A. Gayle, A. Wilder-Smith, and J. Rocklöv. The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine*, 27(2), 2020. DOI:10.1093/jtm/taaa021.
- [20] R. Moss, J. Wood, D. Brown, F. Shearer, A. Black, A. Cheng, J. McCaw, and J. McVernon. Modelling the impact of COVID-19 in Australia to inform transmission reducing measures and health system preparedness. https://www.doherty.edu.au/news-events/news/ covid-19-modelling-papers, Apr. 2020.
- [21] N. Oliver, B. Lepri, H. Sterly, R. Lambiotte, S. Deletaille, M. De Nadai, E. Letouzé, A. A. Salah, R. Benjamins, C. Cattuto, V. Colizza, N. de Cordes, S. P. Fraiberger, T. Koebe, S. Lehmann, J. Murillo, A. Pentland, P. N. Pham, F. Pivetta, J. Saramäki, S. V. Scarpino, M. Tizzoni, S. Verhulst, and P. Vinck. Mobile phone data for informing public health actions across the COVID-19 pandemic life cycle. American Association for the Advancement of Science, 6(23), June 2020. DOI:10.1126/sciadv.abc0764.
- [22] N. Oliver, E. Letouzé, H. Sterly, S. Delataille, M. De Nadai, B. Lepri, R. Lambiotte, R. Benjamins, C. Cattuto, V. Colizza, et al. Mobile phone data and COVID-19: Missing an opportunity? arXiv:2003.12347, Mar. 2020.
- [23] D. Price, F. Shearer, M. Meehan, E. McBryde, N. Golding, J. McVernon, and J. McCaw. Estimating the case detection rate and temporal variation in transmission of COVID-19 in Australia. https://www.doherty.edu.au/news-events/news/ covid-19-technical-modelling-report-and-press-conference, Apr. 2020.
- [24] D. Price, F. Shearer, M. Meehan, E. McBryde, R. Moss, N. Golding, E. Conway, P. Dawso, D. Cromer, J. Wood, S. J. McVernon, and J. McCaw. Early analysis of the Australian COVID-19 epidemic. *eLife*, 9(58785), Aug. 2020. DOI:10.7554/eLife.58785.

- [25] C. Rackauckas, Y. Ma, J. Martensen, C. Warner, K. Zubov, R. Supekar, D. Skinner, and A. Ramadhan. Universal differential equations for scientific machine learning. arXiv:2001.04385, Jan. 2020.
- [26] R. Raskar, I. Schunemann, R. Barbar, K. Vilcans, J. Gray, P. Vepakomma, S. Kapa, A. Nuzzo, R. Gupta, A. Berke, et al. Apps gone rogue: Maintaining personal privacy in an epidemic. arXiv:2003.08567, Mar. 2020.
- [27] Z. Xu, L. Shi, Y. Wang, J. Zhang, L. Huang, C. Zhang, S. Liu, P. Zhao, H. Liu, L. Zhu, Y. Tai, C. Bai, T. Gao, J. Song, P. Xia, J. Dong, J. Zhao, and F. S. Wang. Pathological findings of COVID-19 associated with acute respiratory distress syndrome. *The Lancet, Respiratory Medicine*, 8(4), Apr. 2020.