

Perspective

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

Raj Dandekar,¹ Shane G. Henderson,² Hermanus M. Jansen,³ Joshua McDonald,⁴ Sarat Moka,⁴ Yoni Nazarathy,^{4,*} Christopher Rackauckas,⁵ Peter G. Taylor,⁶ and Aapeli Vuorinen⁷

¹Department of Computational Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139-4307, USA

²School of Operations Research and Information Engineering, Cornell University, Rhodes Hall, Ithaca, NY 14853, USA

³Department of Applied Mathematics, Delft University of Technology, Mekelweg 4, 2628CD Delft, The Netherlands

⁴School of Mathematics and Physics, The University of Queensland, Brisbane, QLD 4072, Australia

⁵Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139-4307, USA

⁶School of Mathematics and Statistics, the University of Melbourne, Melbourne, VIC 3010, Australia

⁷Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, USA

*Correspondence: y.nazarathy@uq.edu.au

<https://doi.org/10.1016/j.patter.2021.100220>

THE BIGGER PICTURE Viral spread may sometimes be controlled via social-distancing directives. However, with a pandemic such as COVID-19, the time between a patient being infected and being recorded as positive can be over a week. This creates a time lag of the order of several weeks between the initiation of a regulatory measure and its observed effect. Safe Blues offers a solution for real-time population-level estimates of an epidemic's response to government directives and near-future projections. Safe Blues strands are safe virtual "virus-like" tokens that respond to social-distancing directives similarly to the actual virus. However, they are spread using Bluetooth and are measured online. The relationship between strand counts and the progress of the actual epidemic can be determined using machine learning techniques applied to delayed measurements of the actual epidemic. This then allows real-time data on the Safe Blues tokens to be used for estimation of the epidemic's current and near-future state.



Proof-of-Concept: Data science output has been formulated, implemented, and tested for one domain/problem

SUMMARY

Viral spread is a complicated function of biological properties, the environment, preventative measures such as sanitation and masks, and the rate at which individuals come within physical proximity. It is these last two elements that governments can control through social-distancing directives. However, infection measurements are almost always delayed, making real-time estimation nearly impossible. Safe Blues is one way of addressing the problem caused by this time lag via online measurements combined with machine learning methods that exploit the relationship between counts of multiple forms of the Safe Blues strands and the progress of the actual epidemic. The Safe Blues protocols and techniques have been developed together with an experimental minimal viable product, presented as an app on Android devices with a server backend. Following initial exploration via simulation experiments, we are now preparing for a university-wide experiment of Safe Blues.

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 50 million confirmed cases and over 1.5 million deaths. In trying to mitigate the effects of this virus, economies are crumbling and whole societies are undergoing transformation due to the new norms dictated by the virus and the 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 measurement and modeling. Viral spread is a complicated function of multiple elements, including biological properties, the environment, preventative measures, such as sanitation and masks, and the level of physical proximity. It is these last two elements, preventative measures and 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 are always lagging and biased: the time between a patient being infected and being recorded as positive can be 1 or 2 weeks. A






Figure 1. The 2020 outbreak in Victoria

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

consequence is that the typical time between the initiation of a regulatory measure and our observation of its effect can be 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 have struggled to deliver an effective response to the pandemic.

Biological properties of SARS-CoV-2 have been studied since the start of the outbreak¹ and are being better understood as time progresses.² 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 $R_{\text{eff}}(t)$, the expected number of individuals infected by an infectious person at time t , is a difficult task.³

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, which presents 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 crises like these, governments face extreme difficulty in determining the optimal level of social distancing that should be imposed. Such decisions often involve guesswork because it is difficult to get immediate feedback on how various implementations of social-distancing regulation affect the level of physical proximity and, in turn, viral spread.

While 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. Can the response be improved in future outbreaks or waves? For example, 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, without real-time estimates of the epidemic’s state, it is unreasonable to expect

such a refined response. Essentially, government has very little real-time information of viral spread to inform its decision-making process.

To help address this problem, we propose the use of Safe Blues.^{4,5} Safe Blues uses Bluetooth signals to transmit tokens between mobile devices that mimic virus spread in a real-time privacy-preserving

manner. The wireless technology that it uses is similar to existing and emerging contact-tracing frameworks such as, Proximity,⁶ Blue Trace,⁷ and the Privacy-Preserving Contact Tracing framework developed by Apple and Google.⁸ However, its purpose is completely different. It does not record and store information about individuals and their interactions with the intention of mapping specific contacts. Instead, various Safe Blues virus strands are periodically created and their spread through the (mobile device) population is tracked. Aggregated counts for each strand are reported to a server without recording private information. An analysis of the counts can be used to produce aggregate estimates of population contact. The result is a real-time evaluation of the effect of any social-distancing rules that are in place. The effects of other mass contact events, such as public demonstrations and large gatherings, are also automatically measured.

This information can be thought of as a proxy measure for aggregate physical proximity. Along with retrospective information about actual (not Safe Blues) SARS-CoV-2 case numbers, it can be used to train sophisticated machine learning (ML) models to estimate SARS-CoV-2 infection numbers as a function of the prevalence of Safe Blues strands. Real-time information on Safe Blues data thus provides, via the ML projections, live near-future estimates of SARS-CoV-2 infection levels and can feed directly into policy decisions.

One may wonder how a safe virtual virus spread solution, such as Safe Blues, differs from “contact counting”-based measurements. Such contact counting could potentially be obtained via contact-tracing apps, although we are not aware of an implementation where aggregate contact-tracing data are downloaded and tracked in real time. A number of technology companies^{9,10} collect and provide data that can be used to help with estimating behavioral patterns in the community. For example, Apple provides daily updated mobile data segregated into three streams: driving, walking, and transport. Google supplies data on GPS-derived indices of the amount of time spent in workplaces, residential areas, grocery stores, transit stations, and the like. These data have been used to provide estimates of population contact rates that have been input into epidemiological analyses.^{3,11} We believe that Safe Blues can provide complementary data, with the social mixing information

provided by Apple, Google, and similar companies being integrated with Safe Blues data to provide powerful ML predictive models. Safe Blues has the potential to be a substitute, and possibly an improvement, for the micro-level distancing information derived from such surveys.

Safe Blues shares some characteristics with agent-based simulation^{12–14} in the sense that its outcomes are driven by the behavior of individuals. Both approaches 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 data is driven by actual population behavior, whereas agent-based models simulate this behavior. The latter have the advantage that they can be used to explore potential future scenarios and can model the effect of interventions, such as mask wearing and surface cleaning. However, they rely on assumptions about how agents interact rather than the actual interactions.

Ultimately policy makers should use a mix of agent-based simulations as in Victoria,¹⁵ aggregate mobility data, such as the Google and Apple data, and virtual safe virus spread solutions, such as Safe Blues, which we propose here. These tools complement one another and, if used effectively, may empower policy makers with a much clearer view on the state of the pandemic. Our focus in this paper is to present the potential of a virtual safe virus spread solution, such as Safe Blues.

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 2–4.¹⁶ Its magnitude depends on a combination of biological and behavioral factors. Key biological factors include the propensity of the pathogen to infiltrate human hosts, the duration of infectiousness in an infected individual, and the susceptibility of different age groups. Key behavioral 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 $R_{\text{eff}}(t)$ 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 behavioral 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 were outlined in a high-impact report,¹⁷ are attempts to control the behavioral component of $R_{\text{eff}}(t)$.¹⁸ 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 $R_{\text{eff}}(t)$ and the course of the epidemic.

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 $R_{\text{eff}}(t)$. Above, we mentioned papers^{3,11} that incorporated data provided by technology companies into a process for estimating $R_{\text{eff}}(t)$ in real time. Other work has involved attempts to quantify the level of human-to-human interaction either at broad scales^{19,20} (such as social gatherings) or at finer scales^{17,21,22} (such as handshaking practices). A notable effort to obtain such a quantification²³ used survey sampling of the UK population to estimate that, in March of 2020, $R_{\text{eff}}(t)$ shifted from around 2.6 before 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. Other approaches to measuring $R_{\text{eff}}(t)$ ^{24,25} use up-to-date count data, such as the now-famous dashboard by the CSSE at Johns Hopkins.²⁶ 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. Lags lead to undesirably slow control updates, which hinder policy makers’ ability to effectively reduce transmission while minimizing 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 to yield estimates of the spread of a virus, such as SARS-CoV-2. Near-future projections and estimates of the $R_{\text{eff}}(t)$ 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 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 neighboring device. Similarly, if the mobile device is in relative isolation, the strand is not likely to spread.




Figure 2. The Safe Blues concept

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. The Safe Blues system operates independently of the health status of individuals.

complete specification of the system is provided by the Safe Blues protocol.⁴

EVALUATION OF SAFE BLUES VIA SIMULATION

For initial evaluation of Safe Blues we created three different simulation models, each of a different nature. This allowed us to test the robustness of the system to different realities. In these simulation models, we tracked both the biological virus and Safe Blues strands. Full details of the models are given in the Computational

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, at least up to the stage where SARS-CoV-2-infected individuals are recognized and isolated. 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. As far as we know, data collected by existing contact-tracing apps around the world are accessed only when an individual tests positive. However, the contacts themselves are recorded without knowledge of the disease status of individuals. This is a characteristic that is shared by Safe Blues transmission. 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 those holding a mobile device in the figure) and others do not. Some level of population participation is required, but Safe Blues does not require all individuals to participate. A

procedures section in the [Supplemental information](#). We now review the key ideas.

All models feature a population comprising a fixed number of individuals. Some of these individuals have Safe Blues-enabled devices, while others do not. At each point in time, the simulated state of an individual registers whether they are susceptible, infected, or removed with respect to the actual (biological) 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 stylized discrete-time stochastic model. One of its appealing features is that the numbers of susceptible, infectious, and removed individuals converge to the well-known SIR difference equations as the population size 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, and so forth). Model III is a spatial movement model with location attraction




Figure 3. Simulation models

Model I: at every time point, each of the 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 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 who are present. Model III: a spatial model where each individual diffuses either around their base (e.g., their home) or around a center (e.g., a supermarket).

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 toward 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 were increasingly imposed over a period of 2 weeks and then fixed to prevent the vast majority of social contact. As these rules went into effect, the infected proportion began a slow but steady decline. All the while, Safe Blues strands were being transmitted on mobile devices. As a consequence, the proportion of Safe Blues infections mirrors the decline in the true infected proportion, driven by a corresponding reduction in physical proximity between Safe Blues-enabled devices. In all the runs, at around 100 simulated days after the start, and after months of a promising decline in case numbers, the social-distancing rules were mostly lifted, resulting in a second wave. The goal of our numerical experiments was to see whether Safe Blues could 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 model²⁷ from the ensemble of Safe Blues strands and historical SARS-CoV-2 information. Full details of this model and the Dynamic Deep Safe Blues model of the next section are specified in the Computational procedures section in the [Supplemental 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 2 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 past social-distancing measures is also

available as a categorical input for projections, e.g., “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 toward a second peak in the number of infected individuals a significant time before such data are 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.

SAFE BLUES FEEDBACK TO POLICY

In addition to being a tool for estimating the current number of infected individuals before such data are 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.²⁸ This approach mixes neural networks with epidemiological models 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 to help decision makers determine the levels of social distancing necessary to contain the outbreak




Figure 4. Prediction with Safe Blues

Deep Safe Blues: Safe Blues detection of a second wave applied to data generated from three different simulation models. The light-colored lines indicate counts of various Safe Blues strands that are inputs to our predictions. 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.

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.

ON THE PENETRATION PROPORTION AND STRAND PARAMETERS

An important question deals with the relationship between the level of adoption of Safe Blues in a population and its prediction accuracy. We call the former the *penetration proportion* and denote it by η . Our initial simulation experiments indicate that Safe Blues can be effective with a small penetration proportion. For example, in a relatively small simulated population of 100,000 people, an η between 0.1 and 0.2 already gives high-quality data. We hypothesize that, with a population size in the order of millions, a penetration level η between 0.05 and 0.1 may be sufficient for successful estimation using Safe Blues. It is important to note that the required penetration level for Safe Blues to be successful is much lower than is needed for contact-tracing apps to be successful.

A potential difficulty is systematic bias. For example, the mobility and interaction profiles of individuals that adopt Safe Blues may differ from those of the general population. Nevertheless, the prediction system will probably work well if this bias is consistent over time, as long as there is sufficient correlation between the proximity patterns of Safe Blues users and other cohorts of the population.

The penetration proportion η has an impact on the way in which Safe Blues strand infections should be configured. Assume that, at a given point in time, COVID-19 has an estimated infection rate of β and a removal rate of γ . We can then set the infection rate and recovery rate of a Safe Blues strand s to approximately follow, $\beta_s \approx \frac{\beta}{\eta}$ and $\gamma_s \approx \gamma$, where the approximation is due to

the lack of exact knowledge about γ and β , as well as due to the desire for some heterogeneity between strands. The choices of the parameters β_s and γ_s can then be incorporated into operational strand parameters. Regarding the incubation period for Safe Blues strands, we may want to configure the strands to incubate faster than in reality, to ensure quicker response time in measurements.

The motivation for this choice of strand parameters comes from basic epidemiological considerations appearing in SIR models. As an example, consider the difference equations associated with model I in Equation (S1) of the Computational procedures section in the Supplemental information. In this case, if one decreases the population size by a factor of η^{-1} , then achieving similar epidemic behavior (on the smaller population) can be achieved by setting β_s and γ_s as suggested above.

A CAMPUS EXPERIMENT FOR SAFE BLUES

We are now in the process of devising a campus-wide experiment to assess the efficacy of Safe Blues. We are aware of at least two similar experiments, each carried out for a different purpose. The Contagion project²⁹ was a large-scale social contact and data collection project in the UK, aiming to simulate epidemic spread. The experiment logged the GPS location of volunteers to the nearest square kilometer and used the data to simulate the spread of a virus. The study was labor intensive because it required self-reporting the number of people each participant interacted with each hour. The FluPhone (and associated EpiMap) project³⁰ used cellular devices to record contacts, but did not focus on virtual safe virus spread and did not restrict information collection to just counts as Safe Blues does. Our proposed experiment thus differs from both of these previous experiments.

The main purpose of the Safe Blues campus experiment is to test whether the Safe Blues system can add to the predictive power of traditional epidemic forecasting methods. The experiment will be carried out at the University of Queensland, St Lucia campus or/and at the University of Auckland's main campus. The experiment will source student participants who will run the Safe Blues app as they attend campus. A randomized reward mechanism will motivate participants to continue to use the app over a prolonged period.

We have developed an Android application based on the open source contact-tracing OpenTrace software⁷ published by the




Figure 5. Policy projection with Safe Blues
Demonstration of policy projection and refinement using Dynamic Deep Safe Blues on model III.

Singaporean government. Our app follows the Safe Blues protocol specified in Appendix A of a previous study⁴ 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 in tracking their spread. The source code for the app, the backend, and the dashboard is available from GitHub.³¹ Figure 6 illustrates the app and the control panel.

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 techniques³² 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 obtained 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 infec-

tiousness 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.⁵

OUTLOOK

A century of technological innovations passed between the 1918–1920 H1N1 pandemic and the current COVID-19

pandemic. 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 ML, 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 ML tools that are celebrated in so many other domains.

The value proposition of methods, such as Safe Blues, is to empower modelers and policy makers to utilize computational statistics and ML 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 modelers 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.³³

Based on simulation experiments, the ML 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 privacy-preserving 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.




Figure 6. Safe Blues software

The control panel, dashboard, and Android app used for the planned campus experiment.

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 contact-tracing apps.^{7,34}

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.

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.patter.2021.100220>.

ACKNOWLEDGMENTS

We thank additional members of the Safe Blues team for fruitful discussion and preparations toward the campus experiment. These include Azam Asanjarani, Keng Chew, Kirsty Short, and Ilze Ziedins. We thank James McCaw, Tom Stace, Yoav Banin, and Yun William Yu for insights. We thank Toshiaki Banerjee for help with illustrations. We also thank two anonymous referees for helping improve the quality of this paper. H.M.J. and Y.N. are supported by the Australian Research Council (ARC) under grant no. DP180101602. S.M. and P.G.T. are supported by the ARC Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS) under grant no. CE140100049. S.G.H. is supported by the Army Research Office under grant no. W911NF-17-1-0094 and the National Science Foundation under grant no. TRIPODS+X DMS-1839346.

AUTHOR CONTRIBUTIONS

In the list of authors, author order is alphabetical. S.G.H., H.M.J., S.M., Y.N., C.R., P.G.T., and A.V. were responsible for mathematical modeling and key ideas of Safe Blues. R.D. and C.R. were responsible for machine learning modeling, prediction, and data analysis. J.McD. and A.V. were responsible for software and app development. The work was coordinated by Y.N.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

- Xu, Z., Shi, L., Wang, Y., Zhang, J., Huang, L., Zhang, C., Liu, S., Zhao, P., Liu, H., Zhu, L., et al. (2020). Pathological findings of COVID-19 associated with acute respiratory distress syndrome. *Lancet Respir. Med.* 8, 420–422.
- Polak, S.B., Van Gool, I.C., Cohen, D., Jan, H., and van Paassen, J. (2020). A systematic review of pathological findings in COVID-19: a pathophysiological timeline and possible mechanisms of disease progression. *Mod. Pathol.* 33, 2128–2138.
- Golding, N., Shearer, F.M., Moss, R., Dawson, P., Gibbs, L., Alisic, E., McVernon, J., Price, D.J., and McCaw, J.M. (2020). Estimating temporal variation in transmission of COVID-19 and adherence to social distancing measures in Australia. https://www.doherty.edu.au/uploads/content_doc/Technical_report_15_Maypdf.pdf.
- Abhijit Dandekar, R., Henderson, S.G., Jansen, M., Moka, S., Nazarathy, Y., Rackauckas, C., Taylor, P.G., and Vuorinen, A. (2020). Safe Blues: a method for estimation and control in the fight against COVID-19. *medRxiv*. <https://doi.org/10.1101/2020.05.04.20090258>.
- Abhijit Dandekar, R., Henderson, S.G., Jansen, M., Moka, S., Nazarathy, Y., Rackauckas, C., Taylor, P.G., Vuorinen, A., Stace, T., and McDonald, J. (2020). Safe Blues project homepage. <https://safeblues.org/>.
- Faggian, M., Urbani, M., and Zanutto, L. (2020). Proximity: a recipe to break the outbreak. *arXiv*, 2003.10222.
- Bay, J., Kek, J., Tan, A., Hau, C.S., Yongquan, L., Tan, J., and Quy, T.A. (2020). Bluetrace: A Privacy-Preserving Protocol for Community-Driven Contact Tracing across Borders (Government Technology Agency). <https://bluetrace.io/>.
- (2020). Apple Inc, Google Inc, privacy-preserving contact tracing. <https://www.apple.com/covid19/contacttracing>.
- (2020). Apple Inc, mobility trends reports. <https://covid19.apple.com/mobility>.
- Google Inc (2020). Mobility report CSV documentation. <https://www.google.com/covid19/mobility/data-documentation.html>.
- Golding, N., Shearer, F.M., Moss, R., Dawson, P., Liu, D., Ross, J., Hyndman, R., Zachreson, C., Gear, N., McVernon, J., et al. (2020). Estimating temporal variation in transmission of SARS-CoV-2 and physical distancing behaviour in Australia. https://www.doherty.edu.au/uploads/content_doc/Technical_report_4_update_29July2020.pdf.
- Hoertel, N., Blachier, M., Blanco, C., Olsson, M., Massetti, M., Rico, M.S., Lemosin, F., and Leleu, H. (2020). A stochastic agent-based model of the SARS-CoV-2 epidemic in France. *Nat. Med.* 1–5, <https://doi.org/10.1038/s41591-020-1001-6>.
- Aleta, A., Martín-Corral, D., Piontti, A.P.Y., Ajelli, M., Litvinova, M., Chinazzi, M., Dean, N.E., Halloran, M.E., Longini, I.M., Jr., Merler, S., et al. (2020). Modeling the impact of social distancing, testing, contact tracing and household quarantine on second-wave scenarios of the COVID-19 epidemic. *medRxiv*.
- Aleta, A., Martín-Corral, D., Piontti, A.P.Y., Ajelli, M., Litvinova, M., Chinazzi, M., et al. (2020). Modelling the impact of testing, contact tracing and household quarantine on second waves of COVID-19. *Nat. Hum. Behav.* 4, 964–971.
- Blakely, T., Thompson, J., Carvalho, N., Bablani, L., Wilson, N., and Stevenson, M. (2020). VIC agent model. <https://www.mja.com.au/journal/2020/probability-6-week-lockdown-victoria-commencing-9-july-2020-achieving-elimination>.
- Liu, Y., Gayle, A., Wilder-Smith, A., and Rocklöv, J. (2020). The reproductive number of COVID-19 is higher compared to SARS coronavirus. *J. Travel Med.* 27, <https://doi.org/10.1093/jtm/taaa021>.
- Ferguson, N., Laydon, D., Nedjati Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunubá Perez, Z., Cuomo-Dannenburg, G., et al. (2020). 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/>.
- Fund, I.M. (2020). Policy tracker: policy responses to COVID-19. <https://www.imf.org/en/Topics/imf-and-covid19/Policy-Responses-to-COVID-19>.
- Price, D., Shearer, F., Meehan, M., McBryde, E., Golding, N., McVernon, J., and McCaw, J. (2020). 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>.
- Moss, R., Wood, J., Brown, D., Shearer, F., Black, A., Cheng, A., McCaw, J., and McVernon, J. (2020). 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>.
- Chang, S., Harding, N., Zachreson, C., Cliff, O., and Prokopenko, M. (2020). Modelling transmission and control of the COVID-19 pandemic in Australia. *arXiv*, 2003.10218.
- (2020). IISc-TIFR COVID-19 city-scale simulation team, city-scale epidemic simulator on github. <https://cni-iisc.github.io/epidemic-simulator/>.
- Jarvis, C.I., Van Zandvoort, K., Gimma, A., Prem, K., Klepac, P., Rubin, G.J., and Edmunds, W.J. (2020). Quantifying the impact of physical distance measures on the transmission of COVID-19 in the UK. *BMC Med.* 18, 1–10, <https://doi.org/10.1186/s12916-020-01597-8>.
- Lavielle, M. (2020). Modelling some COVID-19 data. <http://webpopix.org/covid19.html>.
- Price, D., Shearer, F., Meehan, M., McBryde, E., Moss, R., Golding, N., Conway, E., Dawso, P., Cromer, D., Wood, J., et al. (2020). Early analysis of the Australian COVID-19 epidemic. *eLife* 9, <https://doi.org/10.7554/eLife.58785>.
- Dong, E., Du, H., and Gardner, L. (2020). An interactive web-based dashboard to track COVID-19 in real time. *Lancet Infect. Dis.* 20, 534–553.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning* (The MIT Press).
- Rackauckas, C., Ma, Y., Martensen, J., Warner, C., Zubov, K., Supekar, R., Skinner, D., and Ramadhan, A. (2020). Universal differential equations for scientific machine learning. *arXiv*, 2001.04385.
- Klepac, P., Kissler, S., and Gog, J. (2018). Contagion! The BBC Four pandemic—the model behind the documentary. *Epidemics* 24, 49–59.
- Yoneki, E., and Crowcroft, J. (2014). Epimap: towards quantifying contact networks for understanding epidemiology in developing countries. *Ad Hoc Networks* 13, 83–93.
- Abhijit Dandekar, R., Henderson, S.G., Jansen, M., Moka, S., Nazarathy, Y., Rackauckas, C., Taylor, P.G., Vuorinen, A., Stace, T., and McDonald, J. (2020). Safe Blues github. <https://github.com/SafeBlues>.
- Bickel, J.E., and Small, C. (2020). COVID-19: What Is the Data Telling Us? (The University of Texas at Austin). <https://www.youtube.com/watch?v=wYi3XiWpIQY>, *TexTalks*.
- Oliver, N., Letouzé, E., Sterly, H., Delataille, S., De Nadai, M., Lepri, B., Lambiotte, R., Benjamins, R., Cattuto, C., Colizza, V., et al. (2020). Mobile phone data and COVID-19: missing an opportunity? *arXiv*, 2003.12347.
- Raskar, R., Schunemann, I., Barbar, R., Vilcans, K., Gray, J., Vepakomma, P., Kapa, S., Nuzzo, A., Gupta, R., Berke, A., et al. (2020). Apps gone rogue: maintaining personal privacy in an epidemic. *arXiv*, 2003.08567.