Pervasive Healthcare

EpiCURB: Learning to Derive Epidemic Control Policies

Andrei C. Rusu
University of Southampton

Katayoun Farrahi
University of Southampton

Mahesan Niranjan
University of Southampton

Editors:
Kate Farrahi
k.farrahi@soton.ac.uk

Gabriela Marcu
gmarcu@umich.edu

The effectiveness of an epidemic control policy relies largely on how much effort is invested in every public health measure. Unfortunately, it is seldom possible to optimally allocate funds to these measures if the isolated effect of each intervention cannot be reliably estimated. We show how this challenge can be overcome by utilizing *EpiCURB*, a simulation-control framework that enables us to measure the effect of both untargeted and prioritized interventions on the epidemic outcome, where the latter are guided by reinforcement learning routines that effectively rank eligible individuals.

The COVID-19 pandemic has prompted many countries to implement various non-pharmaceutical public health interventions (NPIs) to prevent the spread of the SARS-CoV-2 virus and protect their populations. Effective as they may have been, these policies have rarely been popular among the population, and have faced intense scrutiny ever since. Given the choice between implementing interventions that entail significant economic and social costs or allowing the virus to spread uncontrollably, authorities faced difficult decisions that could have impacted many lives. Balancing NPIs between their individual effectiveness and limitations, their co-occurrence and potentially inflicting major unintended consequences has been a persistent challenge for authorities around the globe.[1,2] This article demonstrates how simulation-based methods, relying on the recently-introduced *EpiCURB* framework[3] and the SEIR-T individual-based mean-field model,[4] along with suitable visualization tools, can support policy makers in optimizing such decisions and the budgets allocated for each intervention.

EpiCURB (Epidemic Control Using Reinforcement learning Budget allocation) is novel computational approach that advances two previous studies. The first work is by Farrahi et al., who used mobile phone communication traces as proxies for interaction networks where an infection spreads, and measured the effect of contact tracing on reducing the peak infection rate through continuous-time stochastic simulations.[5] The second one is by Meirom et al., who proposed a reinforcement learning (RL) model based on graph neural networks (GNNs) that can generate generalizable epidemic control policies, applicable to heterogeneous networks of tens of thousands of nodes, despite being trained on smaller graphs of 1000 vertices.[6]

In the present work, we introduce a simulation setup that enables policy makers to easily assess the effects of individual NPIs in conjunction with other interventions. We also recommend a suitable visualization methodology that can guide policy makers in identifying where additional resources are necessary to achieve a desired epidemic outcome. With a cost-based model of enhancing the intensity and budget of each NPI, authorities would be able to decide which action is more effective based off such simulated outcomes. Although our method utilizes an individual-based mean-field model specific for the SARS-CoV-2 virus,[3,4] our approach can be used to study any other pathogen, present or future, provided that it can be described via similar equation-based formulations.

PUBLIC HEALTH POLICY CHOICES

Among the various NPIs that were implemented in the COVID-19 pandemic, three types of interventions were the most common and widely used: mobility restrictions, mask mandates, and "test and trace" programmes.

According to the International Organization for Migration (IOM), by June 2020, 219 countries, territories or areas had implemented at least one form of mobility restriction, affecting billions of people.[7] These stringent measures represent some of the most controversial decisions that authorities have taken, because they have severely limited the public's social interactions, and had adverse effects on the economy and society.[2] Stay-at-home orders (also known as lockdowns) are NPIs that restrict the movement and interaction of people within a community, often with the exception of essential activities, such as obtaining food, healthcare, work, or physical activity. The extent and duration of these orders can widely differ depending on the context (e.g., availability of vaccines) and the authorities' objectives (e.g., bringing the effective reproduction number R under 1). For instance, some stay-at-home orders may target specific segments of the population, such as older adults or those with underlying health conditions, while others may encompass the entire population. Some lockdowns may be partial, affecting only certain areas or sectors, while others may be total, impacting the whole country or region. The efficacy and impact of stay-at-home orders have been widely debated and studied in the literature. Several works have found that these measures were pivotal for reducing the caseload of COVID-19, and thus implicitly preventing numerous deaths.[8,9,10] What is more, some studies have suggested that the timing of introducing lockdowns is crucial for their effectiveness and that delaying or relaxing them too early can lead to disproportionally more deaths.[9,10] In stark contrast, a few authors have challenged these claims and argued that stay-at-home orders had little to no effect on the COVID-19 mortality.[2,11] Claiming to have conducted the first comprehensive analysis of NPIs across 180 countries, Herby et al. have shown that lockdowns had no clear impact on COVID-19 excess deaths or mortality rates compared with countries that imposed them later, less stringently, or not at all.[2] One of the main contributing factors to this result, by the authors' claim, is the emergence of voluntary behavioral change among the population, which has significantly contributed to the reduction in mixing irrespective of the enforced measures. Despite being praised by some media outlets and commentators for challenging the conventional wisdom on stay-at-home orders and providing evidence for alternative approaches to managing viral outbreaks, the study has been criticized by several experts for its methodological flaws, data quality issues, causal inference problems, and ideological bias.[12] Most notably, the meta-analysis fails to consider the timing of mandate enactment, which is a critical factor for their success according to the literature.[9,10] A salient example is the UK government's lockdown policy, which has faced frequent criticism for its suboptimal timing, as well as its laxity and inconsistency, that may have ultimately undermined its overall effectiveness.[9] In this work, we present a simulation-based setup that can be used to identify the optimal levels and timing of interventions such as lockdowns, considering the current changes in social mixing, which are either voluntary or enforced by restrictions on gatherings.

Mask mandates have also been the subject of extensive debates throughout the COVID-19 pandemic,[13] with several international bodies or authorities recommending a more relaxed yet inconsistent approach to this issue.[14] Early studies summarized in the comment paper of Brooks et al. have indicated their use have had a significant impact on the spread, with notable mentions including a natural experiment in Germany, which found a 47% decrease in the infection growth rate after the introduction of mask mandates, and another in Canada, which estimated a 25% to 40% decline in the weekly diagnoses following these mandates.[15] Lab studies have also confirmed that masks significantly reduce the transmission of viral-infused droplets, with the common surgical and cotton masks shown to block around 50% when the spreader wears it properly, and 20%-40% when the receiver uses it correctly.[16] More recently, however, a Cochrane review that sparked public debate has noted the lack of reliable randomized control trials, concluding that existing high-certainty evidence on mask efficiency is limited, and points to an insignificant effect of wearing them.[17] The review's controversial conclusions have been met with considerable criticism, mostly targeted at the inclusion of a majority of pre-pandemic studies in the meta-analysis underpinning them.[13] Moreover, as the authors admit themselves, the relatively low adherence observed in the scrutinized trials “hampers drawing firm conclusions.”[17] Indeed, compliance with mask recommendations is known to have been strikingly low in previous viral outbreaks, and it has continued to be low during the COVID-19 pandemic in countries like the UK and Netherlands, while others have seen significant surges in their uptake.[16] These compliance rates have been significantly influenced by the level of perceived risk and trust among the population of each country, which has often been adversely affected by the lack of clear and consistent guidance from the authorities and scientific bodies in those regions, and the insufficient evidence on the isolated benefits of masking compared to other preventive measures, such as hand hygiene or reducing social contacts.[18] Undoubtedly, the task of disentangling the effect of each intervention remains challenging, and without this knowledge, consistency and compliance with mandates may be compromised. Our simulation-based method can shed some light on this issue, enabling policy makers to evaluate the impact of masking in combination with other interventions, as well as to identify the optimal balance between each intervention’s coverage or adoption. Furthermore, the visualizations we propose can facilitate more coherent public health policies that can ultimately enhance the civic compliance with them.

For an epidemic control strategy to be successful in the absence of vaccines or total lockdowns, testing and contact tracing must be efficiently executed by the authorities. Unfortunately, “test and trace” programmes often suffer from potentially crippling inadequacies.[3,4] First of all, manual tracing is upper bounded by staff numbers and their efficiency, leading to intrinsic delays,[19] with memory fallacies of positively-tested individuals often ensuing.[20] Digital tracing, on the other hand, is rarely optimally adopted by populations,[19] with issues such as smartphone access, Bluetooth reliability or privacy concerns being important determining factors.[20,21] The review of Anglemyer et al. have highlighted the importance of understanding how manual and digital contact tracing complement each other in controlling the spread of SARS-CoV-2, and they suggest further research on these processes’ combined effects is required.[21] Other studies have introduced modelling frameworks aimed at exploring how different levels of tracing efficiency affect the epidemic outcomes.[4,20] Using the same simulation framework and visualization setup as for the other interventions, we also investigate the trade-off between manual and digital tracing that policy makers can take into account when distributing resources for contact tracing. The second inadequacy stems from the associated costs of running widespread testing and tracing, both in terms of monetary and human resources.[22] In previous work, we have advocated for the use of RL agents backed by GNNs to maximize these interventions' impact under limited budgets.[3] In the present article, we take this idea a step forward, illustrating how the balancing of budgets can be optimized, in a similar fashion as for other NPIs, to achieve the containment level that is desired for a given region.

COST-BASED ASSESSMENT OF INTERVENTIONS

Outbreak simulations are a useful tool for determining the impact of different public health interventions. However, when modelling various policy options simultaneously, there is no clear indication of which interventions need more effort or the expected costs associated with enhancing those interventions. Here, we propose a simple, yet powerful, methodology that can guide the authorities in making these difficult decisions.

Our approach entails creating a visual space that allows the policy makers to easily investigate the effects of increasing or decreasing the effort and expenditure for each type of intervention. This space can be represented in different ways, such as using contour or 3D plots, illustrating the influence that combinations of interventions at different strengths have on the simulated epidemic outcome. For demonstration, we employ contour plots as an effective way to visualize the variations in these effects. Figure 1 depicts the contours of the mean proportion of individuals kept healthy across 50-100 epidemic simulations over various preferential attachment networks, under different scenarios of simultaneous interventions: reducing mixing and stay-at-home orders; reducing mixing and mask mandates; standard digital and manual tracing processes, relying on a random testing procedure; or RL-targeted testing and tracing.

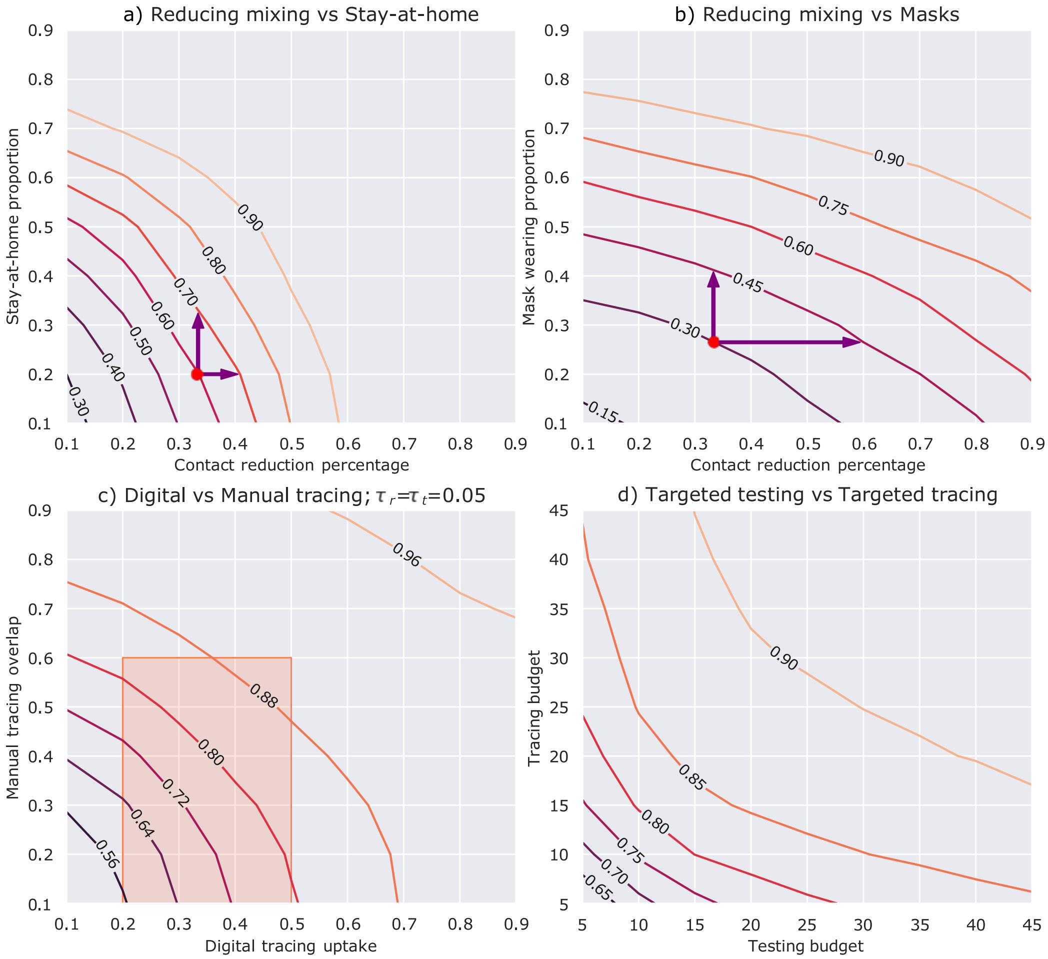

Figure 1. Contour plots of the fraction of people kept healthy across simulations by combinations of different public health interventions. This graphic illustrates how policy makers can assess and allocate resources to each measure, according to the desired level of containment. Figures 1a and 1b have marked the reference points mentioned in the text, with arrows delineating the space of actions needed to reach the next infection threshold. Figure 1c additionally highlights the most relevant region of the space, where uptakes are between 20-50%, overlaps between 10-60%, and 10-15% increases in each direction leads to around 8% less infections. Figure 1d presents the containment rates achieved by different daily budgets for a prioritized testing and tracing policy, derived by an RL+GNN agent. Contours are obtained by averaging 50-100 simulations and smoothening with a Gaussian filter. The sampling resolution used in these plots matches their corresponding axis ticks.

Through the utilization of our method, a policy maker would be presented with several choices of the form: if the goal is to achieve a 10-15% reduction in the spread, would it be preferred to move along the X axis or the Y axis? For example, assume that various factors, such as voluntary behavioral changes or closing entertainment venues and schools, have reduced the social mixing by about 33% compared to the pre-epidemic level, and that about 20% of the investigated community have been advised to stay at home (e.g., people aged over 70). Figure 1a shows that a 10% average reduction in the spread can be achieved either by further reducing the social mixing by 7.5%, or by increasing the fraction of people staying at home by 12% of the total population. Similarly, Figure 1b shows that increasing the mask wearing fraction from 27% to 41% has a comparable effect to decreasing the social mixing by more than 25%, resulting in 15% less infections overall. By applying a transparent cost-benefit analysis that accounts for the economic and social implications of scaling up interventions, authorities can use such simulated outcomes to determine which of them are most effective to implement or expand.

A key consideration for allocating resources for testing and tracing is the presence of diminishing returns: as the desired infection reach becomes more stringent, the required budget increase must be larger to achieve it. Figure 1c illustrates this point clearly. It also reveals that, for the parameter configuration under consideration, a reduction of about 8% in the pathogen attack rate is achieved by increasing either the digital tracing adoption rates within the acceptable and feasible region (20-50%) or the manual tracing coverage in the moderate range (10-60%) by approximately 10-15% (with smaller effects for the uptake, however). Here, an efficacious cost-based assessment can establish which of these actions is more advantageous in each circumstance. For example, at lower adoption levels (i.e., < 30%), the uptake can be more viable to improve, since simple usability or privacy enhancements could be enough to attract more users. In contrast, beyond a certain threshold, increasing the adoption of any application is far more challenging, while improving the staffing of the manual process could become significantly more achievable.

Finally, Figure 1d presents the daily budgeting trade-off between testing and digital tracing when an RL agent is tasked to prioritize both processes. Using such visualizations maximizes the benefits of utilizing our previously proposed targeted approach for epidemic control since policy makers can optimize the budgets allocated for each intervention in a direct manner. When tests in the studied community are insufficient, rationalizing them without compromising on the epidemic outcome is possible by increasing the number of contacts that are to be isolated accordingly. Conversely, if the testing budget is less strict, more contacts can be allowed to continue their normal behavior. The analysis reveals that the two processes have a similar dependence on the budget allocation, highlighting not only the significance of contact tracing, but also the importance of balancing the two for an effective pathogen containment.

METHODOLOGY

We employ the *EpiCURB* framework[3] to simulate several epidemics over weighted Holme-Kim networks[23] of size $N = 2000$ nodes, $m_{HK} = 3$ and $p_{\triangle} = 0.2$. The edge weights are drawn from a uniform distribution $\mathcal{U}(0.5,1)$, which represents the varying levels of transmissibility that interactions can have, depending on their duration and distancing. The pathogen is assumed to be an early variant of the SARS-CoV-2 virus, which spreads according to an individual-based mean-field model that follows an SEIR compartmental formulation with a base transmission rate $\beta = 0.0791$, average exposed rate $\epsilon = 3.7^{-1}$, and average recovery rate $\gamma = 0.05$.[3,4]

When interventions are prioritized by an RL agent, the epidemic is allowed to progress unhinged until $c_a = 5$ days and $c_i = 5\%$ infections have been fulfilled. After informing the agent about the status of $c_k = 25\%$ of infections, we train it according to our previously published routine, reusing the same parameter setting as before,[3] and study their behavior under different budgeting schemes, ranging from five to 45 daily tests or contacts traced. The rest of the studied interventions occur stochastically across the network and begin after five individuals have been exposed to the pathogen. To simulate different levels of mixing reduction, we effectively remove a varying proportion of the edges from the infection network. For mask mandates, we assume a varying percentage of compliant wearers among the population, considering their interactions to have a 50% lower infection weight than the original sampled value (or 75% if both contacts use one). Stay-at-home orders presume the targeted fraction of the nodes and their household immediately discontinue their social patterns, stopping the disease from being transmitted to or from these hubs. Finally, modelling the trade-off between digital and manual tracing is based on the SEIR-T formulation we previously introduced, where the overlap Γ and uptake r parameters are varied, while the testing rate τ_r and the tracing effort τ_t remain fixed at a moderate value of 0.05.[4]

CONCLUSION

This article demonstrates how *EpiCURB* and an appropriate visualization technique can be used to inform budgeting aspects of decision-making in epidemic control. As we exemplify, network-based simulations can reveal the suitable stringency/strength level of each public health intervention required for attaining the desired degree of containment. When standard packages of such measures are deployed, our approach can inform policy makers where additional resources should be allocated. Furthermore, when individual-level interventions, such as testing and contact tracing, are prioritized using RL or other node ranking mechanisms, the budgets under which these operate can be optimized to maximize their benefits.

We believe that balancing NPIs in terms of the associated budgets and outcomes is paramount for the effective functioning of a society during severe pathogen outbreaks. As such, implementing and assessing our proposed framework on more specific real-world scenarios represent natural extensions to this work, which could ultimately aid our preparedness for future epidemics.

ACKNOWLEDGMENTS

ACR was funded by an Engineering and Physical Sciences Research Council (EPSRC) PhD studentship, awarded at the School of Electronic and Computer Science, University of Southampton. MN's contribution to this work was funded by Grant EP/S000356/1, Artificial and Augmented Intelligence for Automated Scientific Discovery, EPSRC, UK. We also acknowledge the use of the IRIDIS High Performance Computing Facility, and its support services at the University of Southampton, in the completion of this work.

REFERENCES

1. J. Hopman and S. Mehtar, "Country level analysis of COVID-19 policies," eClinicalMedicine, vol. 25, Aug. 2020, doi: 10.1016/j.eclinm.2020.100500.
2. J. Herby, L. Jonung, and S. H. Hanke, "Did lockdowns work? The verdict on Covid restrictions," Institute of Economic Affairs, London, 2023. [Online]. Available: https://iea.org.uk/publications/did-lockdowns-work-the-verdict-on-covid-restrictions.
3. A. C. Rusu, K. Farrahi, and M. Niranjan, "Flattening the curve through reinforcement learning driven test and trace policies," presented at the 16th EAI International Conference on Pervasive Computing Technologies for Healthcare (12/12/22 - 14/12/22), Dec. 2022. Available: https://eprints.soton.ac.uk/473363.
4. A. C. Rusu, R. Emonet, and K. Farrahi, "Modelling digital and manual contact tracing for COVID-19. Are low uptakes and missed contacts deal-breakers?," PLOS ONE, vol. 16, no. 11, p. e0259969, Nov. 2021, doi: 10.1371/journal.pone.0259969.
5. K. Farrahi, R. Emonet, and M. Cebrian, "Epidemic Contact Tracing via Communication Traces," PLoS ONE, vol. 9, no. 5, p. e95133, May 2014, doi: 10.1371/journal.pone.0095133.
6. E. Meirom, H. Maron, S. Mannor, and G. Chechik, "Controlling Graph Dynamics with Reinforcement Learning and Graph Neural Networks," in Proceedings of the 38th International Conference on Machine Learning, PMLR, Jul. 2021, pp. 7565–7577. doi: 10.1126/science.286.5439.509.
7. IOM, "IOM COVID 19 Response - Situation Report 21 (26 June 2020)," International Organization for Migration, 21, Jun. 2020. [Online]. Available: https://dtm.iom.int/reports/iom-covid-19-response-situation-report-21-26-june-2020.
8. S. Flaxman et al., "Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe," Nature, vol. 584, no. 7820, Art. no. 7820, Aug. 2020, doi: 10.1038/s41586-020-2405-7.
9. N. Ferguson et al., "Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand," Imperial College London, Mar. 2020. doi: 10.25561/77482.

10. N. G. Davies et al., "Association of tiered restrictions and a second lockdown with COVID-19 deaths and hospital admissions in England: a modelling study," The Lancet Infectious Diseases, vol. 21, no. 4, pp. 482–492, Apr. 2021, doi: 10.1016/S1473-3099(20)30984-1.
11. R. Chaudhry, G. Dranitsaris, T. Mubashir, J. Bartoszko, and S. Riazi, "A country level analysis measuring the impact of government actions, country preparedness and socioeconomic factors on COVID-19 mortality and related health outcomes," eClinicalMedicine, vol. 25, Aug. 2020, doi: 10.1016/j.eclinm.2020.100464.
12. I. Sample and L. Geddes, "Revised report on impact of Covid lockdowns 'adds little insight,'" The Guardian, Jun. 05, 2023. Accessed: Jul. 05, 2023. [Online]. Available: https://www.theguardian.com/society/2023/jun/05/revised-report-on-impact-of-covid-lockdowns-leaves-unanswered-questions.
13. Z. Tufekci, "Opinion | Here's Why the Science Is Clear That Masks Work," The New York Times, Mar. 10, 2023. Accessed: Jul. 03, 2023. [Online]. Available: https://www.nytimes.com/2023/03/10/opinion/masks-work-cochrane-study.html.
14. G. D. Smith, F. Ng, and R. Watson, "'Masking the evidence': Perspectives of the COVID-19 pandemic," J Clin Nurs, vol. 29, no. 19–20, pp. 3580–3583, Oct. 2020, doi: 10.1111/jocn.15401.
15. J. T. Brooks and J. C. Butler, "Effectiveness of Mask Wearing to Control Community Spread of SARS-CoV-2," JAMA, vol. 325, no. 10, pp. 998–999, Mar. 2021, doi: 10.1001/jama.2021.1505.
16. H. Ueki et al., "Effectiveness of Face Masks in Preventing Airborne Transmission of SARS-CoV-2," mSphere, vol. 5, no. 5, pp. e00637-20, Oct. 2020, doi: 10.1128/mSphere.00637-20.
17. T. Jefferson et al., "Physical interventions to interrupt or reduce the spread of respiratory viruses," Cochrane Database Syst Rev, vol. 1, no. 1, p. CD006207, Jan. 2023, doi: 10.1002/14651858.CD006207.pub6.
18. Royal Society, "Face masks and coverings for the general public," 2020, [Online]. Available: https://royalsociety.org/-/media/policy/projects/set-c/set-c-facemasks.pdf.
19. L. Ferretti et al., "Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing," Science, vol. 368, 2020, doi: 10.1126/science.abb6936.
20. M. Mancastroppa, C. Castellano, A. Vezzani, and R. Burioni, "Stochastic sampling effects favor manual over digital contact tracing," Nat Commun, vol. 12, no. 1, Art. no. 1, Mar. 2021, doi: 10.1038/s41467-021-22082-7.
21. A. Anglemyer et al., "Digital contact tracing technologies in epidemics: a rapid review," Cochrane Database of Systematic Reviews, vol. 1, no. 8, 2020, doi: 10.1002/14651858.CD013699.
22. J. A. Pandit, J. M. Radin, G. Quer, and E. J. Topol, "Smartphone apps in the COVID-19 pandemic," Nat Biotechnol, vol. 40, no. 7, Art. no. 7, Jul. 2022, doi: 10.1038/s41587-022-01350-x.
23. P. Holme and B. J. Kim, "Growing scale-free networks with tunable clustering," Phys. Rev. E, vol. 65, no. 2, p. 026107, Jan. 2002, doi: 10.1103/PhysRevE.65.026107.

ABOUT THE AUTHORS

Andrei C. Rusu is a final year PhD student in the School of Electronics and Computer Science at the University of Southampton. Contact him at a.rusu@soton.ac.uk.

Katayoun Farrahi is an associate professor in the School of Electronics and Computer Science at the University of Southampton. Contact her at k.farrahi@soton.ac.uk.

Mahesan Niranjan is a professor in the School of Electronics and Computer Science at the University of Southampton. Contact him at mn@ecs.soton.ac.uk.