

The past, present and future of digital contact tracing

The potential of digital contact tracing to slow the spread of a virus had been quietly explored for over a decade before the COVID-19 pandemic thrust the technology into the spotlight. But can it actually be effective in the hard-to-model complexity of real-world social networks?


Manuel Cebrian

It is late 2020 as I traverse Berlin's U-bahn on the way to my lab. I see large posters covering many of the windows of the train urging passengers to install the official German contact-tracing app, the Corona-Warn-App. Despite initial privacy concerns, the app has been downloaded more than 18 million times. This is a bittersweet moment for me. Around a decade ago, I had investigated whether Bluetooth could be used to trace and contain viral diffusion. Considered to be overly theoretical at the time, the technology has become one of the leading mechanisms to try to slow down the COVID-19 pandemic. However, I cannot help but wonder if the app, and the technology underpinning it, arrived too late, and whether some intricacies of the coronavirus itself limit its effectiveness. After all, a second wave of infection is currently, as I write, in full swing across Europe.

Early days

My journey into digital contact tracing began in 2009, when I joined Sandy Pentland's Human Dynamics Laboratory at the MIT Media Lab. This was my third and final postdoctoral stint, and one I thought would complete my evolution from a computer scientist to a computational social scientist¹. Sandy's lab investigated how we could potentially improve societal functioning if we knew what everyone was doing at every second. The sensing technology was not there yet, so he was building it across many experimental fronts.

The lab space was always messy, a sort of techno-punk underground corner of the increasingly shiny Media Lab. It was also being populated by beacons²: small devices that could be attached to walls to track people by capturing the signal of a WiFi or Bluetooth device emitter they'd volunteer to carry around. This was followed by sociometric badges³: small electronic lanyards that a person would wear as they interacted with others, enabling tracking and



Credit: Panther Media GmbH / Alamy Stock Photo

analysis of who they spoke to and how. And then later, mobile phone health trackers.

The mobile phone trackers looked the least futuristic of all the options, but were, of course, the ones set to become a mainstream technology. By installing special software on the mobile phones of student volunteers, it was possible to track their movement around campus, their digital communication and their face-to-face interactions, as well as to prompt them to report their emotional and physical well-being. To an uninitiated — and admittedly, naive in privacy considerations — researcher like myself, it provided a wonderful opportunity to generate rich, new datasets that could be used to better understand how information and behaviour travelled around networks.

In collaboration with the brilliant doctoral students and postdoctoral fellows in the Human Dynamics lab, we tried a

few out-of-the-box ideas on the datasets from the (often failing) devices. Most never panned out, but, to our surprise, a few did. A project led by Anmol Madan (who would later go on to found the health tracking company Ginger.io) and co-authored by David Lazer (a pioneer of modern computational social science), Sandy and I tested whether we could use these mobile phone communication data to infer actual health measurements in terms of physical symptoms, behaviour and mental health⁴. We succeeded, to a certain degree. We also investigated whether a sociometric badge could be used to model and contain viral spread in an office setting⁵. In doing so, we obtained early indications that these tracking technologies could be used to follow how germs, ideas and emotions were passed from person to person, silently traversing social networks.

Data, models and a lukewarm reception

Around this time, I began to pay increasing attention to epidemiological contact-tracing literature. I learned a lot from the public health approach to contact tracing, but I also felt that it did not capture what I knew about real-world structure and, in particular, the very large heterogeneity of social networks, which was starting to be uncovered by the labs of network scientists^{6–9}. The best paper on the topic had, I thought, been written by Ramon Huerta and Lev Timsring at the University of California, San Diego¹⁰. They had developed a model describing the dynamics of contact tracing in a complex network, and had determined large parameter spaces under which contact tracing would — and would not — be useful. It depended on the viral dynamics and the structure of the social network in which the virus spread, coupled with the contact-tracing strategies applied. The paper made me wonder what would happen if I could swap their simulated social network for a Bluetooth inferred one.

In 2010, I ran this idea by Kate Farrahi, who at the time was a visiting student at the MIT Media Lab (and is now at the University of Southampton). We began work on a project trying to evaluate whether it could be predicted, using a dataset describing social interactions and flu symptoms, who would get infected next, and when such infection would likely start. Kate made the forward projection work, and could predict the evolution of flu days before actual symptom onset¹¹. We saw the opportunity to engage in a parallel, but conceptually opposite, project in which we would no longer be predicting who would get the flu, but how the flu had gotten to someone, mimicking how public health officials traced diseases such as haemorrhagic fevers or sexually transmitted infections.

Kate took the lead of the contact-tracing project, recruiting Rémi Emonet (now at the Laboratoire Hubert Curien in France) to help develop a detailed tracing simulation model for a viral infection using real-world Bluetooth, as well as phone calls and SMS (short message service) interactions. After two years of hard work, we found that a number of contact-tracing strategies could work to reduce the peak and total size of the epidemic, if applied early enough in the course of the spread. Bluetooth seemed to be a viable way of containing an incipient pandemic in a small social network, at least under preliminary examination.

These results did not, however, seem particularly interesting to the academic community of the early 2010s, or, at least, we were failing to communicate them well.

We submitted our manuscript to a prominent social networks conference in 2012, and got a brutal rejection. We received similar, if less abrasive, feedback when discussing the work at seminars: either the model was not of practical interest, or the details of the contact-tracing strategy mattered too much for it to be considered as a straightforward application. These details, however, were critical to the future of the technique, as in many parameter settings, digital contact tracing would not be effective enough to contain a pandemic. We eventually published the paper in the journal *PLoS ONE* in May 2014 (ref. ¹²), and, as far as I can remember, it received a lukewarm reception there also.

COVID-19

I pretty much forgot about the paper, intrigued about its possibilities, but slightly disappointed about its seeming lack of application, and continued to focus on my main line of research: finding people using social networks. But on 7 March 2020, Kate received an e-mail from a senior technical executive in the Singapore government. They were rallying hundreds of software developers to develop the first COVID-19 contact-tracing app, and mentioned that the work was inspired by our 2014 paper. Shortly after, Singapore launched the very first app of this kind.

Over the next few weeks, similar requests followed, ranging from national security agencies seeking advice on how to start implementing digital contact tracing, to an eccentric group of motivated philanthropists who wanted to fund whatever technological development could stop COVID-19 and save lives. We (Kate, Rémi and I) were both surprised and overwhelmed by this response. And at the same time, and in every conversation, we urged caution. In our exploration of the parameter space, digital contact tracing could only stop a pandemic in a relatively small number of scenarios. Was COVID-19 inside that successful area of containment?

As the apps were being rolled out across Europe, summer arrived, and with it, a mirage of calm. People met up outdoors, and that alone likely reduced the spread more than any app could. Modelling studies were being published that quantified that a deadly second wave could be reduced by using contact tracing, both traditional and digital^{13,14}. The once theoretical idea was growing exponentially, like the virus itself, with hundreds of millions of dollars invested globally in its technological development. But by the autumn, and despite most European countries having deployed a widely adopted contact-tracing app, the

second wave had brought lockdowns, restrictions and high alerts to all regions of Europe, even those that had done well during the first wave.

Digital contact-tracing efficacy

Back in August, and worried about the possibility of contact tracing not being as successful as many hoped, I teamed up with Quyu Kong of the Australian National University and Ivan Dotu, co-founder and chief technology officer of Moirai Biodesign, a biomedical start-up. We began to investigate different testing strategies that could have an impact in containing a range of epidemics (from Ebola to H1N1 to measles) under constrained and limited testing resources, and exploring the limits of smart contact tracing and testing strategies. Our preliminary results confirmed some of the suspicions raised in 2014: digital contact tracing may demand a very large volume of highly targeted testing resources (testing not as you discover a potentially infected person, but testing that prioritizes those most likely to be infected and, moreover, infectious) to stop a respiratory pandemic, and for some diseases (measles, for example) it may be impossible even in the best-case scenarios.

According to our simulation model, COVID-19 is a pandemic that can only be stopped if at least 10% of the population is tested every day, and testing is conducted under a global, highly targeted contact-tracing strategy (one that allows potential suspicious cases to be prioritized, ranking them by their likelihood of having been infected and being infectious, and testing them in that order). Also, superspreading events must be avoided. This last phenomenon is governed by a variable known in the epidemiological literature as the dispersion parameter^{15,16}, which dictates the extent to which a virus spreads in a bursty, power-law fashion. The existence of large, unpredictable clusters of spread can destroy digital contact-tracing efforts, which expects a certain degree of time invariance in the viral evolution.

There is then another issue: individual incentives. In my previous crowdsourcing projects, in which I had to mobilize hundreds of thousands of participants to find information or solve a problem, symmetrical incentives that worked both for the recruiter and the recruitee had to be carefully designed, so that a successful multistep chain of viral recruitment could be formed¹⁷: being part of the recruitment brought good things to everyone involved. With most COVID-19 apps, when you are notified that you have had a high-risk contact, or multiple medium-level ones, you are supposed to notify your contacts, isolate

yourself and seek a test. This, naturally, benefits many others around you at potential risk of infection. Yet fully committing to the app's recommendations may carry a hefty cost for yourself, such as a loss of wages or an impact on your mental health. Designing incentives to ensure that people follow the recommendations of the app is a difficult problem¹⁸, and one that requires research in ethics, fairness and privacy^{19–21}.

Earlier this year, I had hope that COVID-19 would turn out to be one of the viruses for which digital contact tracing could make a difference. But now, my doubts about the technique have started to resemble those of the reviewers of our work in 2012. Only careful, large-scale inferential analyses²² of app deployment²³, as well as analysis of how usage might impact transmission dynamics²⁴, will eventually determine the usefulness of digital contact tracing for this pandemic — and perhaps more importantly still, for future even deadlier pandemics. □

Manuel Cebrian ✉

Max Planck Institute for Human Development,
Berlin, Germany.

✉e-mail: cebrian@mpib-berlin.mpg.de

Published online: 25 January 2021
<https://doi.org/10.1038/s41928-020-00535-z>

References

- Lazer, D. et al. *Science* **323**, 721–723 (2009).
- Madan, A. et al. In *Proceedings from Engaging Data: International Forum on Management of Personal Electronic Information* <https://go.nature.com/3ibNISU> (MIT Press, 2009).
- Olguin, D. O. et al. *IEEE Trans. Syst. Man Cybern. Part B* **39**, 43–55 (2008).
- Madan, A., Cebrian, M., Lazer, D. & Pentland, A. In *Proc. 12th ACM Int. Conf. Ubiquitous Computing* 291–300 (ACM, 2010).
- Waber, B. et al. Preprint at <https://arxiv.org/abs/1008.3402> (2010).
- Barabási, A.-L. & Frangos, J. *Linked: The New Science of Networks* (Basic Books, 2002).
- Newman, M., Barabási, A.-L. & Watts, D. J. *The Structure and Dynamics of Networks* (Princeton Univ. Press, 2006).
- Kleinberg, J. M. *Nature* **406**, 845 (2000).
- Christakis, N. A. & Fowler, J. H. *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives* (Little, Brown Spark, 2009).
- Huerta, R. & Tsimring, L. S. *Phys. Rev. E* **66**, 056115 (2002).
- Farrahi, K., Emonet, R. & Cebrian, M. In *Proc. 18th ACM Conf. Computer Supported Cooperative Work Social Computing* 1214–1221 <https://doi.org/10.1145/2675133.2675237> (ACM, 2015).
- Farrahi, K., Emonet, R. & Cebrian, M. *PLoS ONE* **9**, e95133 (2014).
- Aleta, A. et al. *Nat. Hum. Behav.* **4**, 964–971 (2020).
- Ferretti, L. et al. *Science* **368**, eabb6936 (2020).
- Lloyd-Smith, J. O. et al. *Nature* **438**, 355–359 (2005).
- Kupferschmidt, K. *Science* <https://doi.org/10.1126/science.abc8931> (2020).
- Cebrian, M., Coviello, L., Vattani, A. & Voulgaris, P. In *STOC '12: Proc. 44th Annu. ACM Symp. Theory Computing* 775–788 <https://doi.org/10.1145/2213977.2214047> (2012).
- Anderson, A., Huttenlocher, D., Kleinberg, J. & Leskovec, J. In *WWW '13: Proc. 22nd Int. Conf. World Wide Web* 95–106 <https://doi.org/10.1145/2488388.2488398> (2013).
- Van Bavel, J. J. et al. *Nat. Hum. Behav.* **4**, 460–471 (2020).
- de Montjoye, Y.-A. et al. Evaluating COVID-19 contact tracing apps? Here are 8 privacy questions we think you should ask. <https://go.nature.com/3hRUVqV> (Computational Privacy Group, 2020).
- Zuboff, S. *The Age of Surveillance Capitalism* (Public Affairs, 2019).
- Laxminarayan, R. et al. *Science* **370**, 691–697 (2020).
- Martin, T. et al. *Wirel. Commun. Mob. Comput.* **2020**, 8851429 (2020).
- Fowler, J. H. et al. Preprint at <https://arxiv.org/abs/2004.06098> (2020).

Competing interests

The author declares no competing interests.