Tracking the debate on COVID-19 surveillance tools

Contact-tracing apps could help keep countries open before a vaccine is available. But do we have a sufficient understanding of their efficacy, and can we balance protecting public health with safeguarding civil rights? We interviewed five experts, with backgrounds in digital health ethics, internet law and social sciences.

Lillian Edwards is a professor of law, innovation and society. When the UK announced plans for building a coronavirus contact-tracing app in April, Edwards and colleagues drafted a model bill about basic safeguards regarding data collection, repurposing and discrimination. Effy Vayena is professor of bioethics at ETH Zurich and is developing guidance for the ethical use of digital public health tools against COVID-19. Virginia Dignum is a professor of AI and member of the European Commission High Level Expert Group on Artificial Intelligence, and Frank Dignum is a professor of socially aware AI. With many other European researchers they are working on agent-based social simulations to understand the effect of policy responses to the pandemic across different countries. David Murakami-Wood is an associate professor of sociology and faculty at the Surveillance Studies Centre.

Lilian Edwards

The debate surrounding contact-tracing apps has primarily focused on centralized (where anonymized data is aggregated on a central server) versus decentralized (where data remains distributed on individual devices) approaches. Is this framing useful?

My worry when I drafted the bill initially was that this debate was primarily becoming an argument over technical architectures, bringing to mind the term 'techno-solutionism'. A lot of really important issues are marginalized by this framing, such as how people who failed to install these apps might be discriminated against, especially those who are already vulnerable or who do not have a lot of agency. People might lose their job if installing the app was made a condition, or find it difficult to get shifts in the gig economy, because they didn't use the app. My concern is that everyone was talking about how to build the contact-tracing app — but they were not talking about the social conditions in which data got into the app and what becomes of the data that comes out of the app. And this is why I began working on the Coronavirus Safeguards Bill.

Tell us about your Coronavirus Safeguards Bill.

Very simply, the bill says that no one shall be penalized or sanctioned for not using the app. Aside from preserving people's rights, the purpose of this is to instil trust and confidence in the app, as we need a high number of smartphone users — around 80% — to install the app to ensure it works well. There may be many reasons for users to not install the apps — maybe because they have old phones or feel worried by technology or fear consequences like some immigrant groups — and Western governments are unlikely to deal out fines or punishments. But it's much easier to imagine that companies might refuse access to buildings or services unless employees have installed the app. And indeed, there are employers that are already talking about just that. This might be acceptable in certain places where you want to keep risk as low as possible, for example in a doctor's waiting room, but this could also be heavily abused.

What about the prospect of automated decision making via these apps?

I think there hasn't been enough debate about this, because we don't know enough about how these apps would actually work. With some apps that use self-reporting rather than exclusively positive tests there's a real danger — with the best will in the world — of lots of false positives. One way to avoid this is to develop what's called risk scoring, where the app might or might not notify contacts based on your self-reported symptoms, depending, for example, on your recent contacts, their self-reported symptoms or tests and the length of your

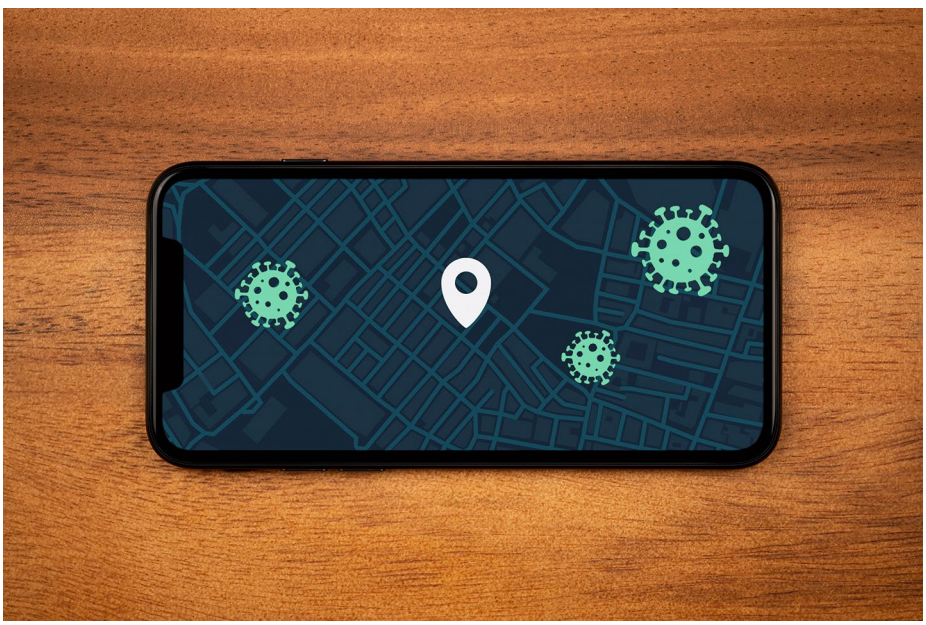

M4OS Photos/Alamy Stock Photo.

exposure to them. This risk profiling in itself is a kind of automated decision-making process. Sometimes it could go wrong, which may lead to contacts having to self-isolate when they don't actually have to. This raises many concerns. How would you, or your contacts, find out what your risk score is? How would you find out how it was calculated? How would you get it changed? Could it be used to reduce, for example, your freedom of movement?

In most high-stakes decisions, such as in the judicial or welfare system, usually there is a human involved in the decision. So now we have — possibly for the first time in some jurisdictions — a solely automated decision about sensitive personal data in an area highly significant for civil liberties. If these apps, or others using automated decision making (for example, to regulate entry to spaces by temperature readings), become popular, then I think we are going to see serious demands around transparency, rights to an explanation, rights to rectification and even the right to take my information and put it into a different model and try and make it work better.

Connectedly, are you concerned about other forms of digital surveillance technologies that might be used during the crisis?

If you imagine the work-from-home world that we're moving into, a big issue is whether work-related digital surveillance becomes more acceptable, because employers still need to somehow keep an eye on their employees, or if it becomes less acceptable, because you're working within your domestic sphere. Workplace surveillance is not in the public consciousness as much as state surveillance, which is very Orwellian, or the influence of social media, which people are sensitive to since the Cambridge Analytica scandal, for example. Workplace surveillance is under the radar, but I think it's going to a big area of controversy soon.

Effy Vayena

Do you agree with how the debate surrounding contact-tracing apps has been framed?

There is a strong focus on individual privacy, and I think that's very important. But even from aggregated data, inferences can be made about risks that apply to certain groups, which may lead to discriminatory effects among this group. Related but distinct is the issue of whether people will really be free to choose to take part in these digital contact-tracing projects, or whether they are going to be unfairly incentivized or even coerced. Although I think it's reasonable that individual choices and preferences are not in the foreground during an emergency, it worries me that we don't hear much about this. Finally, another issue missing from the debate is how aspects of privacy and voluntariness upstream (for example, choosing to use the app) will affect civil liberties downstream (for example, freedom of movement or freedom to gather).

I am concerned that we are not debating the digital contact-tracing projects in their entirety, from downloading an app to ordering that someone quarantines, but rather only debating as if these steps of such a project are not connected.

What can be done to counteract malicious use?

The technical response to this is to make the systems very robust and secure. But we are told almost constantly by our technical colleagues that you can never have an absolutely bulletproof system. So we also need to put governance of these systems in place that further reduces the risk as much as possible, including deciding how much control individuals have over such systems and how accountable these systems are. We must make every effort to reduce the risk to the absolutely acceptable minimum. In my view, these expectations of good governance are not so different from the ones we have in times of normality.

What do you think of the response from the academic community?

It's interesting that the academic community has been very agile and vocal in ways that I haven't seen before. There have been open letters signed by many scientists and engineers about different technologies, mainly addressed to governments. However, I'm not seeing much active engagement with the public: right now the public are represented only as some polling on what they find acceptable. One has to be pragmatic in an emergency of course, but we have come to realize that we don't have a good system in place for broad public engagement during crises.

What do you think about the role of technology companies during this crisis?

If governments and state actors collaborate with the private sector, the conditions of these collaborations must be negotiated very carefully and transparently. The public must know the exact conditions of any deal, because the stakes are very high in this case, and because an emergency situation gives licence to many actors to do things that they wouldn't have done under a non-emergency state.

Independent oversight bodies are a useful safeguard and can be used not only in drafting the conditions of the collaboration, but also in monitoring each phase of the project. I see these projects as a really huge experiment, with lots of uncertainty, which we have not exactly thought through because the time does not allow it. Because of this we need independent assessment that is ongoing and has powers to stop or to intervene with these upcoming digital public health projects. This will also help very much for public accountability.

How do you hope the WHO guidelines can be practically used and incorporated by states?

As many countries are currently developing digital proximity and contact-tracing systems, I hope the principles we delineated will help them make ethically defensible choices of technologies and implementation strategies. Each context is different and decision makers will have to take their contexts into account; however, meeting ethical requirements is necessary in any context. Given the emergency and the fact that for the first time we are trying out these large-scale digital public health tools, doing it ethically will determine their success.

Are you concerned about other forms of digital surveillance technologies that might be used during the crisis?

There's a risk that a lot of things will creep in that we did not have the chance to think through. So I'm particularly concerned about technologies that we had already flagged as being problematic, or ripe for abuse (for example, face recognition technology) in the days before the crisis, and that we may now resort to as a solution. Any time we introduce some new technology, we need to carefully consider whether this is better than existing alternative ways of solving some of those public health issues.

Frank and Virginia Dignum

Tell us about the agent-based simulations your group are working on.

The purpose of our simulations is to ask whether track-and-trace technologies are effective given what we currently know about the virus, where people are asymptomatic or pre-symptomatic for a long period but are still infectious. A recently published model claimed that approximately 60% of the population will need to use a contact-tracing app for them to be effective. But there are also economic or social factors

that influence the utility of track-and-trace technologies. Our simulations show that contact-tracing apps are even less effective at suppressing the spread than previously thought.

Why is this?

The reason is that we use a more fine-grained behavioural model, modelling many types of people within a population and many types of ways that people interact. From this we get an emergent macro model that seems to be more consistent with what's happening in reality. It also provides full transparency and an understanding of how we get to our conclusions. This is crucial, because validating models with empirical data is very challenging in a novel pandemic in which one can't assume normal conditions.

So, a lot of simulations of epidemic patterns assume relatively normal socio-economic conditions, which currently do not apply.

Especially if we consider machine learning models, in which we learn from data measured in the past. We will learn useful insights if the future is similar to the past. And we are now experiencing a moment in which the present is completely different from the past, and the future is very uncertain. Pretending that we can do anything with these data-driven, machine learning type of models is very risky. We need a much higher level of transparency, explainability and accountability for what we're doing.

What is your view of the debate on centralized versus decentralized approaches for contact-tracing apps?

From the perspective of fundamental human rights, a decentralized approach is of course more suitable. But, independently of centralized or decentralized approaches, the question is: what do we gain with this technology? It might be that one approach is better than the other in providing a suitable app for our mobile phones that can support each one of us to feel safe. But at this moment, we cannot really promise safety at all because we don't know how these things will work. I'm worried that we are giving people a fake sense of security, which may end up as a big problem. So we have to decide what we are doing this for. And if the point is to help public health authorities, then we need to listen to them and see what they really want, and what kind of technology is best for that. Generally, they want a centralized approach with lots of data collection, with more than just the location history of people. Perhaps they don't even need a track-and-trace app, but some technology that helps the contract tracing teams currently doing it manually.

What role — if any — do you see for AI technologies to have a positive impact in this crisis?

I think that there is a lot that can be done with AI during this crisis. But, we have to be extremely careful and transparent about what we are doing and what assumptions are being made. AI could help a lot in developing more powerful modelling tools that combine traditional epidemiology, with complex social and economic factors.

What do you think about the role of technology companies during this crisis?

Infrastructure is very important, so making it available, open and transparent for use by governments and academic institutions is crucial, especially if we need large-scale data processing. Also, technology companies need to be aware of their responsibility in particular regarding the spreading of misinformation. This is something for which we need to demand much more accountability on their part.

David Murakami-Wood

What is your view of the debate on centralized versus decentralized approaches for contact-tracing apps?

Centralization falls at the first challenge: it is simply not necessary for the purpose of controlling, reducing and ultimately eliminating the spread of COVID-19 infections. Beyond this, there are no advantages to centralized contact-tracing solutions in a democratic society that respects human rights, apart from convenience for the state. The kind of countries in which centralized contact-tracing apps have been deployed are generally already 'surveillance states', that is, countries that have a great deal of power over citizens and knowledge about them. These systems tend to be far more than simple contact-tracing, and mean comprehensive surveillance with access to multiple databases and promiscuous data sharing. There might be an epidemiological or public health benefit to this system, but it's almost incidental.

A decentralized architecture, which stores personal information and IDs on the user's mobile device, and which requires that user's permission for sharing through highly protected stages, has exactly the same functionality as a basic centralized architecture, but without the same potential for expansion (although one can't say that every decentralized protocol has zero potential for misuse).

Is anonymization enough to protect us from potential misuse?

Privacy on the whole, with some exceptions, is a typical liberal human right, in that it is premised on the holder of the right being an individual person. But most of the most serious injustices in the world do not relate to individuals (although they experience the outcomes and the daily reality) but to groups of people: classes, races, genders, age groups, sexual orientations and so on. And we already know from research by brilliant scholars like Safiya Noble, Ruha Benjamin and Latanya Sweeny how much even apparently simple algorithms and the software they are built into can embody economic inequalities and social prejudices, and produce biased output that only serves to reinforce the existing inequalities that generated the bias.

Now you might agree with all this and think, 'but what could possibly be wrong with de-identified contact-tracing?' Well, quite a lot. Contact-tracing is about location and movement — tracking, in other words. People move as individuals, but they move in particular patterns that reflect community dynamics. Indeed, Google has a massive effort devoted to turning these kinds of movements into operational urban policy with a system called Replica. These patterns can be mined and compared with lots of other data in our increasingly 'smart' environments — and even with good intentions — this could produce outcomes that are undemocratic and discriminatory. With bad intentions, it could be much worse. 'We know where you live' should be thought of as much with a plural 'you' as with a singular. I'm not saying that this is inevitable, not at all, I'm saying we need to be paying attention, and that these issues need to be discussed before we choose certain paths.

In many ways, what we need to worry about isn't the ordinary person with the cell-phone, or even the often-invoked figure of the malicious hacker. We need to worry about misuse of data within organizations — whether that be for policy or profit. Massive health datasets never sit around for a long time without insurance companies or the police wanting access to them, often for reasons that seem perfectly sensible to the data controllers. But again, this just emphasizes why decentralized solutions are less risky that centralized ones.

So, it seems like decentralized 'privacy-by-design' approaches place a strict limit on the possibility of misuse or mission creep. Is this enough?

There has been a real change in how engineering education and development processes happen, and the pressure to

include ethics and various privacy-by-design ideas has played a huge role in that. But neither of these things are substitutes for dealing with wider political economic issues, as privacy-by-design apps don't in themselves change the overall balance of power between users and companies or citizens and government.

Having said that, decentralized architectures offer far fewer avenues for any kind of mission or function creep than centralized ones. If a government or corporation doesn't have users' data in its database, it simply doesn't have the temptation to find other uses for that data, share it with others or monetize it.

Do you think the academic community have been sufficiently involved in the development of these systems?

Many of the government app developments are being done in secret. I was involved in one such app development right at the start, hoping to change how things were done, but I stepped out of it almost immediately when I realized the secrecy and the pressure to be seen to be constructive would constrain my ability to be a critic, even inside, let alone outside. That has nothing to do with the people who are still involved; they are all good people with the right attitudes. But it highlights, for me, some of the pressures involved in this.

One of the obvious solutions is to 'make your own', and that's exactly what the DP-3T group has done, in producing an open-source, decentralized set of protocols for contact-tracing. That initiative has demonstrated the agility and speed of a genuinely open network of like-minded researchers — obviously, it helps that they are some of the best people out there. And this of course 'inspired' Apple and Google to develop their decentralized system.

Overall, however, there has been very little genuinely public discussion, and certainly very little in the way of accountable, democratic political debate.

What role — if any — do you see for AI technologies to have a positive impact in this crisis?

When it comes down to it, the spread of COVID-19 could have been limited simply by states paying more attention. They should have funded not the development of apps on individual phones but existing systems of global disease surveillance through the World Health Organization. Most of what we're doing now is dealing with the consequences of ignoring our surveillance systems at that scale. If there are long-term implications for AI from this, it's this level we should be focusing on — extending systems of planetary-scale surveillance that benefit humanity collectively, rather than ones that focus on individual human beings.

Interviewed by Yann Sweeney

Published online: 16 June 2020
https://doi.org/10.1038/s42256-020-0194-1