



Mind the App—Considerations on the Ethical Risks of COVID-19 Apps

Floridi Luciano^{1,2}

Published online: 13 June 2020
© Springer Nature B.V. 2020

In the past months, there has been a lively debate about so-called COVID-19 apps¹ developed to deal with the pandemic (Morley et al. 2020b). Some of the best solutions use the Bluetooth connection of mobile phones to determine contacts between people and therefore the probability of contagion, and then suggested related measures. In theory, it may seem simple. In practice, there are several ethical problems (Morley et al. 2020a), not only legal and technical ones.² To understand them, it is useful to distinguish between the *validation* and the *verification* of a system, remembering that these words are used here only in the design/engineering sense, and certainly not philosophically.

The validation of a system answers the following question: “are we building the right system?” The answer is no if the app

- is illegal. For example, the use of an app in the EU must comply with the GDPR; mind that this is necessary but not sufficient to make the app also ethically acceptable (Floridi 2018), see below:
- is unnecessary. For example, there are better solutions;
- is a disproportionate solution to the problem. For example, there are only a few cases in the country;
- goes beyond the purpose for which it was designed. For example, it is used to discriminate people;
- continues to be used even after the end of the emergency.

Assuming the app passes the validation stage—for example, assuming that an app planned in Europe, not just in the EU, is “valid” in the sense just stated—then, it needs to be verified.

¹https://en.wikipedia.org/wiki/COVID-19_apps

²<https://www.lightbluetoothpaper.org/2020/04/12/contact-tracing-in-the-real-world/>

✉ Floridi Luciano
luciano.floridi@oii.ox.ac.uk

¹ Oxford Internet Institute, University of Oxford, 1 St Giles, Oxford OX1 3JS, UK

² The Alan Turing Institute, 96 Euston Road, London NW1 2DB, UK

The verification of a system answers the following question: “are we building the system in the right way?” Here too, the difficulties are considerable.

For once, privacy (or personal data protection to be more precise) is not an unsolvable problem. For example, a Bluetooth-based app can use anonymous data, recorded only in the mobile phone, used exclusively to send alerts in case of contact with people infected, in a non-centralised way. It is not easy, but it is feasible.³ Of course, it is trivially true that there are and there might always be privacy issues. And problems concerning group privacy (Florida 2014) are not being discussed enough (for a valuable exception, see Taddeo 2020). However, the point is that, in this case, privacy risks can be seriously mitigated, and made much less pressing than other issues. Once (or, if the reader is more sceptical than I am, even if) privacy is taken care of, it is other ethical difficulties that need to be resolved. They concern the effectiveness and the fairness of the app.

To be effective, an app must be adopted by a huge number of people. In the UK, for example, the consensus seems to be that the app would be pointless if used by less than 20% of the population. According to comments about the PEPP-PT, very significant effectiveness seems to be reached around the threshold of 60% of the whole population. This means that in the UK, for example, where about 20% of the population does not own or use the right mobile phone (see below), the app would have to be used by 75% of the population to be very successful. In Italy, the app should be consistently and correctly used by something between 11 to 33 million people, out of a population of 55 million. Consider that in 2019, Facebook Messenger was used by 23 million Italians. Even the often-mentioned app TraceTogether in Singapore has been downloaded by an insufficient number of people: 1.4 million users out of a population of 5.6 million residents, that is, 25% of the population, when the “National Development Minister Lawrence Wong, who co-chairs a multi-ministry task force to manage Singapore’s response to the pandemic, has said three-quarters of the population must have the app in order for it to be effective”.⁴ The app in Australia seems to have been equally unsuccessful.⁵

Given that it is unlikely that the app will be adopted so extensively just voluntarily, out of goodwill, social altruism or responsibility, it is clear that it may be necessary to encourage its use to make it effective, but this only shifts the problem. Because the wrong, positive incentives to use a voluntary app—for example, to use the app as a kind of passport, to allow users back to work or have privileged access to some services—have at least two drawbacks:

1) they tempt one to play the system or cheat. For example, one may leave the phone in the drawer, if a “clean” record is required to be able to have access to some services; or one may provide the wrong kind of self-assessment, or even look for “dangerous Bluetooth contacts”, if appearing to be at risk guarantees an earlier medical treatment;

2) they are unfair if only those who have the app have access to some services. This may create first- and second-class citizens. As an example, consider that, according to the latest DESI report of the EU Commission (data are available until 2016⁶), only 44%

³ <https://ico.org.uk/media/about-the-ico/documents/2617653/apple-google-api-opinion-final-april-2020.pdf>

⁴ <https://www.scmp.com/week-asia/people/article/3084903/coronavirus-why-arent-singapore-residents-using-tracetgether>

⁵ <https://www.theguardian.com/world/2020/may/24/how-did-the-covidsafe-app-go-from-being-vital-to-almost-irrelevant>

⁶ <https://ec.europa.eu/digital-single-market/en/scoreboard/italy>

of Italians between 16 and 74 years have “at least basic digital skills”. And not everybody owns or can afford the right model of mobile phone equipped with Bluetooth. In the UK, the NHS COVID-19 App Data Ethics Advisory Board (disclosure: I am a member) has expressed its concern that the app may exclude large numbers of people, given that according to Ofcom 21% of UK adults do not use a smartphone.⁷

Therefore, the positive incentives should be independent of the use of the app; that is, they should not be endogenous (what is offered by the app), but exogenous (what is offered for the app). In the case of the app, an example of an exogenous incentive mechanism may be a lottery. Installing and running the app may be like buying a lottery ticket, and periodical prizes could be given to “app tickets”, in order to gamify the process and incentive people to get as many “tickets” (apps downloaded and activated) as possible. Those tempted to “game the system”—e.g. by having people near them adopting the app to have “more tickets”—would actually be reinforcing the positive trend, ensuring more widespread adoption of the app. And a lottery-like mechanism would avoid unfair discriminations in terms of social advantages or other kinds of facilitation unavailable to those who are on the wrong side of the digital divide. These would all be good outcomes.

If some kind of lottery-based incentive mechanism is not adopted or does not work sufficiently well, the other problem is the potentially unfair nature of a voluntary solution that is used by a limited community (see 2 above). The app works much better the more it is widespread, and it is most widespread where there is more digital literacy and ownership of mobile phones equipped with Bluetooth. Therefore, by making it voluntary and linking it to some benefits or advantages, there is a very concrete risk of privileging the already privileged and their residential areas. The following summary from the EU *Digital Economy and Society Index Report 2019 - Human Capital*⁸ is worth quoting in full:

“In 2017, 43% of the EU population had an insufficient level of digital skills. Seventeen percent had none at all, as they either did not use the internet or barely did so.

According to the digital skills indicator, a composite indicator based on the digital competence framework for citizens*, 17% of the EU population had no digital skills in 2017, the main reason being that they did not use the internet or only seldom did so. This represents an improvement (i.e. decrease) of 2 percentage points compared with 2016. The share of EU citizens without basic digital skills, in turn, went down by 1 percentage point (to 43%). However, these figures imply serious risks of digital exclusion in a context of rapid digitisation. There are proportionally more men than women with at least basic digital skills (respectively, 60% and 55%). In addition, only about 31% of people with low education levels or no education have at least basic digital skills. Forty-nine percent of those living in rural areas have basic digital skills compared with 63% in urban areas. There are still major disparities across member states. The share of people with at least basic digital skills ranges from 29% in Bulgaria and Romania (despite noticeable progress in both these countries in 2017) to 85% in Luxembourg and 79% in the Netherlands.”

⁷ <https://nhsbsa-socialtracking.powerappsportals.com/EAB%20Letter%20to%20NHSx.pdf>

⁸ <https://ec.europa.eu/digital-single-market/en/human-capital>

If the app is not carefully planned, designed, deployed and managed, the digital divide may become a biological divide. It may be objected that a small percentage of people voluntarily adopting the app is better than nothing. I doubt it, given that such a small percentage tends to live together in the same areas. It will be great, but probably mostly for them (benefits and positive incentives included), and largely useless for all the others. Indeed, the unfairness of a voluntary app adopted by a small percentage of the population may be used to support the case for a mandatory solution, that is negative incentives. Where the primary benefit is collective, not individual, it may be contextually ethical to require that those who can afford more do more; that is, those who can adopt a digital solution are asked to adopt it. This “negative” incentive (the mandatory app) could “positively” discriminate the better-off (the digitally educated owners of a Bluetooth-enabled mobile phone) to help, now as a very large group, the worst off. Metaphorically, a compulsory app would work like a tax. Those who can contribute more digitally may be required to do so. Consider that in some cultures, this may be a *de facto* social outcome, not a legal requirement, whenever peer pressure and expectations make a behaviour so strongly expected to be almost compulsory. Let me haste to add that I am not advocating the mandatory nature of the app, just highlighting the counter-intuitive effects of a policy.

In all this, there is a temptation one must resist: a merely political solution, in the following sense. Validation and verification are of course related. Above all, they are related by a fundamental twofold relation:

- lack of validation should stop one from building the app; and
- poor or failed verification should force one to re-consider the validation of the app in the first place, including its necessity.

If it turns out that one cannot build an app rightly, one should not build it at all in the first place (it is not the right solution to build) or change approach completely, because of the positive + negative + opportunity costs incurred by the adopted solution. This relation helps to explain why the app may be criticised in some cases as a mere political solution.

When a solution is labelled as merely political (“it does not matter whether the system works or not, even if it does not, we must show that we are doing something”), then this means that the backward link between verification and validation is severed. It does not matter if (mind, not whether) one cannot build the app rightly—in particular, it does not matter if the app is totally or largely ineffective—one should still build it (validation is ok) because the validation does not concern the system but the signalling process of building the app, and even an ineffective app is still successful in signalling that something has been tried and done.

Sometimes, severing validation and verification is ethically justified. It is often justified in cases of high uncertainty (the “just in case it works” kind of scenarios). In general, politicians love to think that the V/V severance offers a win-win situation: if it works great, if it does not, one has still shown goodwill, care, that one has tried everything, that one cannot be blamed for not trying, etc. Of course, the temptation is to treat all costs—which are the reason against severing validation and verification—as externalities (they affect the future, the next government, etc.). However, I think that this is not one of those cases—that is, I think that it is not a case in which we are

ethically justified in severing validation and verification—because the costs (all kinds of them) are high and cannot be “externalised” (even assuming that “externalisation” was a good practice, which is not): they will hit a population/government pretty quickly pretty soon, potentially making the whole problem worst. Therefore, one should avoid the risk of transforming the production of the app into a signalling process. To do so, the verification should not be severed from, but must feedback on, the validation. This means that if the verification fails, so should the validation, and the whole project ought to be reconsidered. It follows that a clear deadline by when (and by whom) the app and the project behind it may be assessed (validation + verification) and in case be decommissioned, or improved, or even simply renewed as it is, is essential. This level of transparency and accountability should be in place since the very first step in conceiving and designing an app-based policy and a clear, independent and effective form of governance to manage the whole process would be essential.

Technology can help, substantially. I have argued (Floridi 2019) that the human project for the twenty-first century should be based on a new marriage between the green of all our environments (biological and human-made) and the (electric) blue of our digital technologies. It can be done, and the app is a good test. But it must be done successfully and ethically, or the solution will become part of the problem, e.g. in terms of social injustice.

An app by itself will not save us. And the wrong app may be worse than useless, as it will cause ethical problems and potentially exacerbate health-related risks, e.g. by generating a false sense of security, or deepening the digital divide. A good app must be part of a wider strategy, and it needs to be designed to support a fair future. This is possible. But if the opportunity is not seized ethically, better do something else, avoid its positive, negative and opportunity costs, and not play the political game of merely signalling that something (indeed anything) has been tried.

It is clear that we are entering some uncharted areas of digital ethics. The way forward may lie in designing the right policies that incentive the adoption of the app (voluntary, mandatory or a mix of the two), and/or in a different architecture of the app (e.g. decentralised vs. centralised, using GPS data, etc.), and/or the nature of the hardware required (think of a cheap, freely-distributed Bluetooth-based tracker, like those that one can attach to one’s keys to find them at home), and/or how the app is used (think of an app-hub, able to support a whole family through only one mobile phone, in connection with other Bluetooth trackers). But any solution should take care of its ethical implications, and be flexible enough to be improved rapidly, to rectify potential shortcomings, and take advantage of new opportunities, as the pandemic develops. In two recent articles (Morley et al. 2020a, b), Jessica Morley, Josh Cowsli, Mariarosaria Taddeo and I provided a list of 16 questions to be answered to check whether an app is ethically justified and to evaluate the sort of trade-offs that may be required to support its use. I remain moderately optimistic about the opportunity: it is possible to find an ethically good solution, and efforts to devise it are worth making. But I am aware that it will not be easy, especially since a digital solution may not be proposed again to the public if there is an initial failure. Indeed, given the challenges and the risks, by the time this article is published, more countries might have followed Belgium and decided not to adopt any app. Sometimes, “now is not the time”.

References

- Floridi, L. (2014). Open data, data protection, and group privacy. *Philosophy & Technology*, 27(1), 1–3.
- Floridi, L. (2018). Soft ethics, the governance of the digital and the General Data Protection Regulation. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2133), 20180081. <https://doi.org/10.1098/rsta.2018.0081>.
- Floridi, L. (2019). The green and the blue: naïve ideas to improve politics in a mature information society. In C. Öhman & D. Watson (Eds.), *The 2018 yearbook of the digital ethics lab* (pp. 183–221). Cham: Springer International Publishing.
- Morley, J., Cowls, J., Taddeo, M., & Floridi, L. (2020a). Ethical guidelines for COVID-19 tracing apps. *Nature*, 582, 29–31.
- Morley, J., Cowls, J., Taddeo, M., Floridi, L. (2020b). *Ethical guidelines for SARS-CoV-2 digital tracking and tracing systems*. Available at SSRN 3582550.
- Taddeo, M. (2020). The Ethical Governance of the Digital During and After the COVID-19 Pandemic. *Minds & Machines*. <https://doi.org/10.1007/s11023-020-09528-5>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.