On health data architecture design

Prashant Agrawal[*] Subodh Sharma[*] Ambuj Sagar[†] Subhashis Banerjee[‡*]

October 22, 2021

The National Digital Health Mission (NDHM) – announced by the Prime Minister on the Independence Day, 2020 – aims to develop the backbone needed for the integrated digital health infrastructure of India. Developing countries like India, with significant health challenges, perhaps need such an infrastructure the most. This can help not only with diagnostics and management of individual health, but also with broader public health monitoring, socio-economic studies, epidemiology, research, prioritising resource allocation and policy interventions. Digitisation cannot be a substitute for the fundamentals – for example, investment in nutrition and welfare, primary healthcare services and healthcare professionals – but it can potentially make healthcare more organised, efficient, and effective.

However, most attempts to build such large-scale, nation-wide digital systems for health (Temperton, 2016; Charette, 2018; Lovell, 2019; Shrikanth & Parkin, 2019), and national digital identity systems crucial for supporting such infrastructures (The London School of Economics and Political Science, 2005; GOV.UK Press Release, 2011; Khera, 2019), have been mired in controversies. They have often been questioned on privacy and fairness grounds and have been difficult to operationalise. Some have had to be abandoned altogether. In India too, the recent momentum and concerns around informational privacy guarantees have occurred in the context of the creation of new government databases and digital infrastructures for welfare delivery (Unique Identification Authority of India, 2020; Government of India, 2020; Khera, 2019; Banerjee & Sharma, 2019). The concerns are manifold, and rushing into a design and implementation without adequate due diligence is fraught with risks.

In this note we investigate the considerations necessary for building such an infrastructure. We argue that to be able to meet the broad social objectives not only do the crucial privacy and fairness concerns have to be comprehensively addressed, but also that the theory of public good based on such an infrastructure needs to be carefully developed and the operational requirements and risks need to be clearly understood. In particular, an effective proportionality analysis by balancing the utility versus the risks becomes untenable when either the utility or the risks are inadequately or imprecisely modelled. We examine the necessary elements of a conceptual architecture required to enable such a proportionality analysis.

1 Uncertain theory of public good

That a well-functioning personal health data infrastructure can potentially lead to public good and welfare is undeniable. Individual health records – accessible across healthcare centres and hospitals

[*]Computer Science and Engineering, IIT Delhi, New Delhi 110016. Email: {prashant,svs,suban}@cse.iitd.ac.in
[†]School of Public Policy, IIT Delhi, New Delhi 110016. Email: asagar@sopp.iitd.ac.in
[‡]Computer Science, Ashoka University, Plot #2, Rajiv Gandhi Education City, P. O. Rai, Sonepat, Haryana 131029 (on leave from IIT Delhi). Email: suban@ashoka.edu.in

by treating physicians – can certainly facilitate better management of health episodes. Additionally, reliable history of illnesses, reports and medication may ensure that patients do not have to be treated blind.

As the COVID-19 pandemic has shown, such an infrastructure possibly could also have played a crucial role in public health management. Mandatory recording of test reports and reliable death records would have helped epidemiologists accurately estimate the extent of the disease impact in various locales and geographies as the disease progressed, and plan for eventualities better. This data was sorely missed, and, from several accounts, the deaths were probably underreported by a significant factor (Balakrishnan, 2021; Banaji, 2021) leading to complacency. Also, real-time availability of test reports with fine-grained individual-level details like location, occupation, work-place location and environment, commute patterns, etc. could have immediately led to better contact-tracing and containment strategies. Moreover, it is obviously helpful to be able to correlate test reports with vaccination records of each individual, and longitudinally track individuals in the population to accurately estimate vaccine efficacy and breakthrough infections. Such data may be crucial both for monitoring variants and future vaccine design. In the absence of a health record infrastructure, we had to make do with tracking small cohorts by independent research groups, and there would always be doubt as to whether such sampled data is representative.

Also, periodic cross-sectional surveys such as the National Family Health Survey (NFHS) (International Institute for Population Sciences, 2021), Annual Health Survey (AHS) (Office of the Registrar General & Census Commissioner, India, 2021), Rapid Survey On Children (RSOC) (Ministry of Women and Child Development, GOI, 2021), Comprehensive National Nutrition Survey (CNNS) Ministry of Health & Family Welfare, GOI (2021), Surveys of the National Nutrition Monitoring Bureau (NNMB) (National Institute of Nutrition, ICMR, GOI, 2021) etc. – though tremendously useful for macro-level understanding of national health indicators – are limited in many ways. For example, definite understanding of the causal factors leading to the alarmingly high rate of stunting and wasting in India (India State-Level Disease Burden Initiative CGF Collaborators, 2020; Deshpande & Ramachandran, 2021) has remained elusive. Regression studies using aggregated cross-sectional data can only do so much, and real-time longitudinal tracking of nutrition and health episodes of individual mothers and children through home check-ups, primary care centre visits, pregnancy episodes, childbirth and immunisation – for the whole population – may certainly be of immense value. Such high frequency clinical data – in other situations – may also be used for disease forecasting, epidemiology, managing infectious diseases and monitoring public health in general.

However, just highlighting the potentials, as we have done above and as done in the NDHM blueprint (Ministry of Health and Family Welfare, Government of India, 2019), does not develop the full theory of public good using digital health data. That would require clearly identifying the exact data analysis and inference techniques that would be required and the various outcomes that may be expected, for both diagnostics and public health; analysing the exact requirements for future research and innovation; identifying the correlations with other socioeconomic data that may be necessary for research and public health monitoring; and carefully analysing the exact frequency and nature of clinical data that should be recorded at various interfaces, such as at point of care measurements, personal equipment, imaging devices, nutrition records, and the observations made by primary healthcare professionals and specialists etc. It would also require analysis of the feasibility and cost of such recordings, the error models for such data, the error control strategies and assessment of the impact of errors and missing data on inference and data analysis.

Developing such a theory of public good will inevitably require a multidisciplinary approach – including practitioners and researchers – and should not be left primarily to IT professionals. Otherwise, it may lead to aberrations like Cowin (Ministry of Health and Family Welfare, Government

of India, 2020a) and Aarogya Setu (National Informatics Centre, Ministry of Electronics and Information Technology, Government of India, 2020). On the one hand, instead of concentrating on the backend, and facilitating the supply chain and the correlation of vaccination records with tests and infection breakthroughs, Cowin focussed on centralisation of vaccination scheduling which should ideally have been left to local administrations whose needs and service populations were different and diverse (Gupta & Jain, 2021). There is a reason for which health is a state subject in the federal structure of our constitution. On the other hand, Aarogya Setu rushed into contact tracing using Bluetooth and GPS without pausing to evaluate what theory – if any – may facilitate computing infection risk estimates from Bluetooth proximity estimates and GPS locations, and what may be the error rates of such estimates (Banerjee *et al.*, 2020a,b). These perhaps exemplify how not to do digitisation for health.

Even in the current NDHM policy documents (Ministry of Health and Family Welfare, Government of India, 2019, 2020b) there is a serious lack of such analysis, and, as a result, questions regarding the preparedness of the health infrastructure for contributing accurate, standardised and meaningful data remain unanswered (Malhotra *et al.*, 2021).

2 Operationalisation and use cases

Even if a theory of public good is well established, it is the translation of the theory into the operational elements that becomes pivotal for the 'utility vs pain' question, and careful design of the use cases – the processes that define the interfaces between the digital and the human elements – is perhaps the most crucial aspect of digitisation.

Poorly thought out use cases that fail to account for the diverse cultural background and social realities can easily lead to unforeseen behavioural adaptations resulting in unacceptable quality of service delivery. The sheer variety of the actors – ranging from poor rural populace – including children, over-worked and under-staffed healthcare professionals in primary healthcare centres, midwives and other home care workers, Anganwadi workers in childcare centres, Accredited Social Health Activist (ASHA) workers, patients and health professionals in government and private secondary and tertiary care hospitals, imaging and diagnostic centres, pathology labs, specialists, desk workers, insurance personnel, epidemiologists, researchers, administrators, bureaucrats, policymakers and several others – with widely different levels of digital proficiency – adds to the challenge. Resource crunch, with doubtful internet connectivity in several areas, compound the problems. Denial of services and exclusions due to authentication failures, internet and server failures or digitisation errors, and delays and other increased cost of transactions due to poor scheduling and processing are the typical poor outcomes of digitisation use cases that the designers need to explicitly handle. This requires careful modelling of not only the technical and administrative processes themselves, but also the risks of technological and administrative failures as part of a threat model, and ensuring that some well-articulated property of non-exclusion is never violated.

Most importantly, proper use case design and analysis will require participation of not only IT professionals but practitioners with ground level knowledge and users with direct stake in the system. Hence it is crucial to ensure that they buy-in to the proposal, and their participation is eager and voluntary.

The operational requirements also require identifying and understanding the diverse data sources and their complexity. This may involve understanding the constraints of personnel, resources and equipment at various data generation and consumption points, understanding their primary functions and ensuring that they are not hampered in any way. It also requires an understanding of the frequency of data generation, error models, access rights, interoperability, sharing, data analysis,

dissemination and other usage requirements, and designing the data organisation and application programming interfaces appropriately.

Laying out the use case and operational design blueprints for all to reflect upon is imperative before any implementation is even considered. The NDHM policy documents (Ministry of Health and Family Welfare, Government of India, 2019, 2020b) do not seem to address these considerations adequately.

3 Privacy and denial of rights

In a digitisation attempt as complex and sensitive as this, privacy is a crucial concern, and the potential tensions between public good and individual rights are bound to generate disquiet. These concerns must be examined threadbare, as must the suitable ways to navigate them. Any data infrastructure endeavour that fails to effectively address privacy concerns is bound to get mired in controversies and endless litigations. In fact, most attempts at building health data infrastructures worldwide – including in the UK, Sweden, Australia, USA and several other countries – have led to serious privacy-related controversies and have not yet been completely successful.

In order to develop a suitable operational standard for privacy protection, we must first clearly understand the nature of privacy harms. Often, privacy is conflated with security, that too only against external breaches, and this flawed understanding leads to problematic solutions.

3.1 The privacy concerns

The most common fear of digitisation, especially when enforced by governments, is that of Orwellian mass surveillance and misuse of data. All digitisation requires unique digital identities, and the use of a universal health identity across all health-related transactions can create an infrastructure for totalitarian observation of citizen's health data. Also, linking of health identities with general-purpose identities like Aadhaar, as the National Digital Health Blueprint (NDHB) (Ministry of Health and Family Welfare, Government of India, 2019) suggests, exacerbates the problem.

The other common fear of digitisation is the secrecy aspect of privacy, i.e., the fear of exposing one's private world to the public space which may potentially cause embarrassment or public judgment. The secrecy aspect is obviously important in the healthcare context, and it is primarily the considerations of maintaining secrecy that lead to the emphasis on the prevention of data leaks and identification of specific individuals from statistical database queries.

However, surveillance and secrecy are not the only privacy concerns with digitisation. A far more common and subtle manner of erosion of privacy is by the way of losing control of information about oneself to insensitive, uncaring and opaque bureaucracies, who may use it for their own interests with little regard to the direct or indirect harms caused to the individual. In the healthcare context, leaking health data to unauthorised entities may result in direct harms through social prejudice and discrimination. But also unpredictable use of health information by unpredictable entities may result in indirect harms such as mis-profiling, profiling for commercial interests and predatory targeting, incorrect or unfair scoring using out-of-context data, exposing vulnerabilities to malicious actors, etc. This may lead to fallouts like illegal denial of jobs or jacked up costs of insurance based on individualised health data, direct drug marketing, predatory advertisements targeted to the vulnerable, etc. Solove (2004) argues that *Kafkaesque* is a more appropriate metaphor for describing such a situation and the helplessness of individuals in fixing it.

Big-data analytics and machine learning algorithms, which are important reasons for building a national digital health infrastructure in the first place, contribute greatly towards Kafkaesque threats

to privacy and liberty. It is forcefully argued by O'Neil (2016) that big-data analytics systems, by the very fact that they are designed by the privileged and often for profit, magnify inequality and historical biases. Often such systems use poor proxies to make decisions about human life, lack transparency, accountability or flexibility in their decisions, have a tendency to become ubiquitous because of their perceived efficiency gains, and are generally greatly damaging if left unchecked. Thatcher *et al.* (2016) argue that "As algorithms select, link, and analyse ever larger sets of data, they seek to transform previously private, unquantified moments of everyday life into sources of profit." Similar concerns have also been raised by Zuboff (2018) and Eubanks (2018).

Other common Kafkaesque dangers, as also highlighted in Section 2, are poorly thought out use cases, incomplete case analyses and incompetent programming (Khera, 2019). As common fallouts, one may suddenly find herself being deregistered from services due to no fault of hers, or having to unnecessarily run around to get things corrected when she was not the one responsible for the mistakes in the first place. Being denied hospital treatment, pension or welfare because perhaps a name is misspelt or because fingerprints do not match will be cases in point. Such callous omissions are obvious threats to rights to privacy, liberty and life (Banerjee, 2017).

The Orwellian big brother and the Kafkaesque arguments certainly raise crucial concerns, but they do not necessarily imply that privacy protection is impossible with digitisation. We first need to carefully model the direct and indirect privacy risks and then evaluate – through the lens of proportionality – how these risks may be balanced against a well-developed theory of the public good. Because of the enormous benefits that digitisation and analytics promise in healthcare, we must earnestly look for solutions to mitigate the risks.

3.2 Limitations of standard approaches for privacy protection

The data protection principles for privacy in digital databases have mainly been based on the tenets of informed consent and notice; collection, purpose and storage limitation; participation of individuals; transparency; regulations, enforcement and accountability (The Planning Commission: Government of India, 2011; Srikrishna *et al.*, 2017). However, the enforcement mechanisms are typically weak. We try to understand why.

3.2.1 Informed consent

Although there is considerable focus on user consent and notice in much of privacy jurisprudence, consent holds little meaning in case of a nationwide health infrastructure roll out, presenting a false sense of choice to individuals where actually there is none. Also, it is unreasonable to expect individuals to be able to give *informed* consent because it is unrealistic to assume that they can predict the possibilities and scale of Kafkaesque dangers from an opaque bureaucracy. In general, as pointed out by Solove (2004), and also by Matthan (2017) and Srikrishna *et al.* (2017), notice and consent are generally ineffective because of information overload, limited choice and consent fatigue. This is often reflected in the customary negligent clicking of 'I Agree'. In view of this, NDHM's (Ministry of Health and Family Welfare, Government of India, 2019, 2020b) overreliance on a consent-based architecture appears to be problematic. There is a strong case for a rights-based approach that shifts a significant part of the responsibility of privacy protection and accountability from the individual to the data controller, irrespective of the level of consent.

3.2.2 Ex-post accountability vs ex-ante protection

Typical methods of ensuring accountability are based on post-facto auditing and punitive measures in case of breaches or user complaints. However, despite causing grave distress, Kafkaesque privacy

harms are very hard to detect. Even when they are detectable, it is hard to show a causal relationship of a tangible harm with a given breach or insider malfeasance at a data controller. For example, it may turn out to be impossible to know for sure whether a person has lost her job because of the officially put out reason or because her personal medical data was accessed without authorisation and used to discriminate against her. Thus, privacy violations should be prevented from happening in the first place, i.e., privacy protection must be *ex-ante*.

3.2.3 Conflating privacy with security

On the technology side, privacy has often been conflated with security and the focus has been on keeping the data secure from *external* threats, more or less ignoring the possibility of insider attacks either due to rogue system administrators or due to conflict of interest of the entire data controller organisation. This has led to an excessive emphasis on data encryption and other safeguards that are controlled – and are hence potentially overridable – by privileged insiders at the data controller. Moreover, the indirect and subtle nature of Kafkaesque harms mean that often well-intending projects end up having unintended privacy and fairness side-effects unless they are carefully controlled. Hence, security, though necessary, is not sufficient for privacy. What we require is ensuring purpose limitation through independent regulatory oversight on the data controller's data processing activities, careful risk modelling and analysis, and, above all, public participation and transparency.

3.2.4 Anonymisation

The other oft-touted solution to protecting privacy while releasing personal data for unrestricted use, as has also been proposed in the NDHB (Ministry of Health and Family Welfare, Government of India, 2019), is *anonymisation*. Anonymisation is the process of removing personal identifiers from a database by suppressing information, coarsening data or adding noise, with the goal of making it impossible to identify any individual from the released data. However, almost a decade of research in the field of de-anonymisation has shown that anonymisation is often unreliable. A small number of data points about individuals coming from various sources, none uniquely identifying, can completely identify them when combined together (Narayanan & Shmatikov, 2008). Reports in the literature have shown that anonymised census data (Rocher *et al.*, 2019; Ullman, 2021), social-network data (Narayanan & Shmatikov, 2009; Narayanan *et al.*, 2011), genetic data (Gymrek *et al.*, 2013; Erlich *et al.*, 2018), location data (de Montjoye *et al.*, 2013), credit card data (de Montjoye *et al.*, 2015), writing style data (Narayanan *et al.*, 2012), web-browsing data (Su *et al.*, 2017), etc., can be robustly de-anonymised to re-identify individuals (Narayanan & Shmatikov, 2019). This is backed by theoretical results (Datta *et al.*, 2012; Aggarwal, 2005) which show that for high-dimensional data, anonymisation is not possible unless the amount of noise introduced is so large that it renders the database useless. Thus, release of even anonymised personal data for unrestricted use must be a strict no-no.

3.2.5 Differential and inferential privacy

An absolute notion of informational privacy might be that no information about an individual could be inferred with access to a statistical database that could not be inferred without any such access. In her celebrated result Dwork (2006) not only proved that such absolute *inferential privacy* is impossible to achieve, but also observed that if an adversary has access to arbitrary auxiliary information, an individual's inferential privacy may be violated even when she does not participate in the database, because information about her can be leaked by correlated information of other individuals. This

led to the development of the notion of *differential privacy*, which attempts to limit the information gained by an adversary when an individual's data is collected versus when it is not collected, thus limiting the *additional* privacy risk an individual incurs by participating in a database. Note that differential privacy is a considerably weaker notion because even though differential privacy guarantees that individuals cannot be identified, de-anonymisation and other correlation attacks can still infer a lot of information about them from differentially private databases (Ullman, 2021).

3.2.6 Algorithmic fairness

Although differential privacy addresses the secrecy aspect of privacy – by preventing reidentification of any individual – the Kafkaesque issues of potential misuse and discrimination due to biased applications of machine learning and big-data analytics remain. The European GDPR has proposed 'right to explanation' as a countermeasure (The European Parliament and the Council of European Union, 2016). However, predictive analytics rarely support causal reasoning, and, without expert manual audit of algorithmic and data biases, the algorithmic explanations will most likely turn out to be inane. Moreover, the adverse outcomes of perverse machine learning applications are Kafkaesque, and the consequent damages are not immediately obvious. So timely explanations may never even be sought or examined.

It is evident from the above discussion that an effective privacy protection architecture must be based on independent regulatory oversight and focus on ex-ante prevention of violations by ensuring purpose limitation rather than on fixing ex-post accountability. The NDHM's blueprint document (Ministry of Health and Family Welfare, Government of India, 2019) does briefly mention a Security and Privacy Operations Center that appears to be envisaged to play such a regulatory role. However, details regarding not only its legal structure and the precise regulatory obligations, but also how – or on what basis – the regulator may discharge its obligations are completely unclear. Also, there is no well-articulated threat model for privacy protection. As such, the approach appears to be highly non-standard.

Below we outline what may be the necessary elements of an effective regulatory architecture.

4 Elements of a privacy architecture

The lack of a standard grammar for articulating the operational requirements for privacy in large public service applications often results in the proponents and the opponents talking past each other in privacy debates, with one side forcefully proclaiming 'privacy-by-design', and the other side throwing the proverbial 'kitchen sink' of privacy concerns. For example, claims such as 'NDHM has built-in privacy-by-design because it prevents data aggregation in its federated architecture where the data remains in its original location' or 'storing identity information isolated in separate hard-to-access confidential stores mitigates privacy leaks' arise out of imprecise articulation of privacy threat models and do not stand up to scrutiny. It is not at all clear what precise aspects of privacy do these measures mitigate and how, because data aggregation depends more on data access patterns and post-access purpose limitation than on where the data is located, and without an unambiguous specification of their properties it is hard to ascertain what the confidential stores of digital identities actually achieve, especially with possibilities of insider leaks and threats of re-identification with a myriad of correlation techniques. In the absence of adversary threat models debates on proportionality based on such claims are not particularly useful. Indeed, we have seen in the past that not only did the outcomes in the majority and minority opinions of Aadhaar judgement (The Aadhaar judgment, 2018) turn out to be diametrically opposite, but the privacy debates in several other recent applications have also been repetitive in nature without making much headway.

Whereas modelling of utility depends on the context of the application and must vary from case to case, specifying the operational requirements for privacy requires a standard tool. In what follows we present a conceptual architecture for privacy analysis based on identifying a precise threat model and defining an ideal functionality of use cases.

4.1 Privacy analysis of use cases and ideal functionality

Any use case of a large public service application will necessarily leak some information at the peripheral interfaces, where it interacts with either humans or other digital systems. There are also other unavoidable risks, such as leakage of control-flow information over a network (for example, which entities have communicated with each other and when), and possible authentication failures and system outages. In an otherwise perfectly secure system, these should be the only risks in an application. Hence the first obligation of the regulator must be to accurately model these inevitable risks and carry out a privacy risk assessment of whether they are acceptable.

This requires first identifying the regulatory boundary of the application, i.e., the scope of control of the regulator, analysing the interfaces at the regulatory boundary, and capturing the potential technical and administrative failures as well as malicious actors as parts of an adversary *threat model*. The regulator should then develop a model of the *ideal functionality* of the application that captures the intended ideal-world execution of the application and explicitly models the communications with the adversary under the considered threat model. For example, the regulator must model a health professional to whom a patient's data may be revealed as a potential adversary who may leak or misuse the information, an administrative process as a potential adversary whose errors may cause denial of service, and a researcher with whom certain anonymised statistical data may have to be shared as a potential adversary who may orchestrate a de-anonymisation attack using other auxiliary information. It must be the responsibility of the regulator to analyse how or to what extent the use cases may be hardened and evaluate the unavoidable risks.

The ideal functionality thus acts as an abstract specification that sets a standard for real-world implementations. It also identifies and models the unavoidable privacy risks associated with the application. Making the ideal functionality public allows any analyst to clearly understand the scope of regulatory control, the threats considered, and whether the inherent risks with a proposed use case are acceptable and meet the standards of proportionality. A data protection framework is incomplete till it lays out the operational standards for such an analysis.

The technological obligation for any subsequent real implementation would be to guarantee that the ideal functionality is never violated, or, if unavoidable, clearly document the gaps if any. In fact, according to computer science security principles, it must be (mathematically) proved that a real implementation is *indistinguishable* from a virtual simulation of the ideal functionality and there are no additional privacy, fairness or exclusion risks except those already modelled in the ideal functionality.

4.2 Purpose limitation through audited, tamper-proof programs

Since unpredictable use of personal data can lead to arbitrary Kafkaesque dangers, one critical component of the ideal functionality must be that all application programs are pre-approved by the regulator and run exactly as specified, i.e., are *tamper-proof*, even in the presence of malicious privileged insiders. The ideal functionality must also specify that the application programs do not leak any information, other than those pre-defined at the interfaces, and that all inter-program data exchanges and all data storages are secure and the accessed data is used only for authorised purposes by authenticating only the approved tamper-proof programs.

This would require the regulators to perform rigorous privacy, security and fairness audit of all application programs. This is likely going to be a predominantly manual review process since, above all, it is hard to encode fairness in an automated algorithm (Kleinberg *et al.*, 2016). Once the programs are audited and approved, they should be published on a public bulletin board, along with the regulators' digital signatures, to invite public scrutiny and debate. A guarantee must also be provided that indeed it is only the published programs that can be executed.

4.3 Virtual identities

All public service applications need a digital identity infrastructure to identify and authenticate individuals. One obvious way to minimise privacy loss across applications is to use different *virtual identities* with different applications, such that the identifiers are impossible to correlate. For example, one may use different identifiers with different care or diagnostic centres for different medical episodes, and they should only be able to correlate or access information that is allowed by the regulator.

Such virtual identities also allow controlled interoperability via anonymous credentials (Chaum, 1985; J. Camenisch and A. Lysyanskaya, 2004). Anonymous credentials allow one to transform a credential issued against a virtual identity they own to an identical credential against another virtual identity they own, such that the issuing authority does not know the purpose for which the transformed credential is later used, and a service provider obtaining the transformed credential does not know about the individual's information stored with the issuing authority. The unlinkability and untraceability can also be designed so that they may be overridden on case-specific situations, for example by a regulatory approval, or to permit legitimate analytics by linking of silos.

Virtual identities ensure that ideal functionalities for each application can be predominantly analysed in isolation. It is to be noted, however, that it may still be possible to de-anonymise using other auxiliary information (see Section 3.2), and for this reason, unrestricted release of data should always be viewed as risky and, as much as possible, data should be released outside the regulatory boundary only for human consumption and not for copying or forwarding to other unknown data processing elements.

4.4 A dynamic authorisation architecture

Another crucial function of the regulator ought to be to determine and clearly define who can access what data and for what purposes, based on legal sanction or on authorisations, in conformance with approved regulations. Purpose limitation needs to be built into such authorisations, and all purpose extensions and authorisation renewals should be explicitly considered.

This, in turn, would require specification of a dynamic authorisation architecture in the ideal functionality. Some of the data access authorisations may be role based and static; for example defining what parts of a patient's data a treating physician may be able to access depending on the nature of the complaint. However, granting access to the treating physician in the first place must be a dynamic authorisation process, with support for 'grant' or 'revoke' updates. Such authorisations to use personal data should not only be defined for the purposes of operations, investigation and review by human agents, but also for granting programmatic access for data mining to pre-approved tamper-proof programs.

4.5 Implementation notes

The ideal tamper-proof programmes that we mention above may not be practically realisable and only approximations may be possible. Nevertheless, they are useful theoretical concepts for privacy modelling and defining ideal functionalities. Near tamper-proof programs may be realised using a variety of emerging technologies such as hardware-based trusted execution environments (TEEs), remote attestation, fully homomorphic encryption, secure multiparty computation, etc. Practicality and the security guarantees offered by the chosen techniques would be important deciding factors to consider. And, it will be of utmost importance to document the vulnerabilities and the risks associated with these techniques and factor them in to any proportionality analysis.

Virtual identities can be realised mostly using anonymous credential techniques, along with tamper-proof programs to link them under regulatory supervision wherever required.

The dynamic authorisation architecture can be specified as a privacy policy using standard programming language techniques and it can be enforced automatically at runtime using suitable parsers. Standard techniques exist.

To prove that a real implementation is indistinguishable from the ideal functionality, existing proof techniques in computer security can be leveraged.

Finally, it will be necessary to build the required state and regulatory capacity for such analyses. There is really no alternative. Without these capabilities the privacy risk analysis and their mitigation will always be ad hoc and suspect.

5 Conclusion

Building a health data infrastructure at a national scale is a problem of unprecedented complexity, and it requires the highest standards of due diligence and guarantees. In this note we have tried to highlight some of the considerations that should go into such a design effort, and have suggested an analysis framework through defining an ideal functionality of an application as precisely as possible. It is to be noted that the ideal functionality specification can never be a static one-time affair, but has to continuously evolve with understanding, analysis and feedback. However, it is an indispensable conceptual tool for analysing the fairness and proportionality of an application.

Also, most importantly, such an endeavour must follow a due process of adequate public consultation and ensure that it is backed by a law. There also has to be the will to build the required regulatory capacity and an effective, rights-based specialised data protection framework for health data infrastructure.

References

Aggarwal, Charu C. 2005. On k-anonymity and the Curse of Dimensionality. *Pages 901–909 of: Proceedings of the 31st international conference on very large data bases*. VLDB '05. VLDB Endowment.

Balakrishnan, Paran. 2021. *Study estimates 1.21 million Indians have died from Covid-19*. https://www.telegraphindia.com/india/study-estimates-1-21-million-indians-have-died-from-covid-19/cid/1817293. [Online May 31, 2021].

Banaji, Murad. 2021. *Estimating Covid-19 Fatalities in India*. https://www.

theindiaforum.in/article/estimating-covid-19-fatalities-india. [Online May 10, 2021].

Banerjee, Subhashis. 2017. *A Welfare Test for Aadhaar*. http://indianexpress.com/article/opinion/columns/a-welfare-test-for-aadhaar-upa-nda-aadhaar-card-4921582/. [Online; posted 4-November-2017].

Banerjee, Subhashis, & Sharma, Subodh. 2019. Privacy Concerns with Aadhaar. *Commun. ACM*, **62**(11), 80.

Banerjee, Subhashis, Raman, Bhaskaran, & Sharma, Subodh. 2020a. *How Reliable and Effective Are the Mobile Apps Being Used to Fight COVID-19?* https://thewire.in/tech/covid-19-mobile-apps-india. [Online April 16, 2020].

Banerjee, Subhashis, Raman, Bhaskaran, & Sharma, Subodh. 2020b. *On the proportionality of Aarogya Setu*. https://www.livemint.com/opinion/online-views/covid-19-tracking-app-on-the-proportionality-of-aarogya-setu-11594183812518.html. [Online July 8, 2020].

Charette, Robert N. 2018. *Australians Say No Thanks to Electronic Health Records*. https://spectrum.ieee.org/riskfactor/computing/it/australians-choosing-to-optout-of-controversial-my-health-record-system. [Online July 27, 2018].

Chaum, David. 1985. Security Without Identification: Transaction Systems to Make Big Brother Obsolete. *Commun. ACM*, **28**(10), 1030–1044.

Datta, Anupam, Sharma, Divya, & Sinha, Arunesh. 2012. Provable De-anonymization of Large Datasets with Sparse Dimensions. *Pages 229–248 of:* Degano, Pierpaolo, & Guttman, Joshua D. (eds), *Principles of security and trust*. Berlin, Heidelberg: Springer Berlin Heidelberg.

de Montjoye, Yves-Alexandre, Hidalgo, César A, Verleysen, Michel, & Blondel, Vincent D. 2013. Unique in the Crowd: The privacy bounds of human mobility. *Scientific reports*, **3**.

de Montjoye, Yves-Alexandre, Radaelli, Laura, Singh, Vivek Kumar, & Pentland, Alex "Sandy". 2015. Unique in the shopping mall: On the reidentifiability of credit card metadata. *Science*, **347**(6221), 536–539.

Deshpande, Ashwini, & Ramachandran, Rajesh. 2021. *Picture This: How caste increases stunting in Dalit kids*. https://ceda.ashoka.edu.in/picture-this-how-caste-increases-stunting-in-dalit-kids/. [Online July 29, 2021].

Dwork, Cynthia. 2006. Differential Privacy. *Pages 1–12 of: Proceedings of the 33rd International Conference on Automata, Languages and Programming - Volume Part II*. ICALP'06. Berlin, Heidelberg: Springer-Verlag.

Erlich, Yaniv, Shor, Tal, Pe'er, Itsik, & Carmi, Shai. 2018. Identity inference of genomic data using long-range familial searches. *Science*, **362**(6415), 690–694.

Eubanks, Virginia. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. USA: St. Martin's Press, Inc.

Government of India. 2020. *Direct Benefit Transfer, Government of India*. https://dbtbharat.gov.in. [Accessed May 31, 2020].

GOV.UK Press Release. 2011. *National identity register destroyed as government consigns ID card scheme to history*. https://www.gov.uk/government/news/national-identity-register-destroyed-as-government-consigns-id-card-scheme-to-history. [Online posted 10-February-2011].

Gupta, Apar, & Jain, Anuksha. 2021. *India's technocratic approach to vaccination is excluding the digitally-deprived*. https://indianexpress.com/article/opinion/columns/indias-technocratic-approach-to-vaccination-is-excluding-the-digitally-deprived-7315442/. [Online May 15, 2021].

Gymrek, Melissa, McGuire, Amy L., Golan, David, Halperin, Eran, & Erlich, Yaniv. 2013. Identifying personal genomes by surname inference. *Science*, **339**(6117), 321–324.

India State-Level Disease Burden Initiative CGF Collaborators. 2020. Mapping of variations in child stunting, wasting and underweight within the states of india: the global burden of disease study 2000-2017. *Eclinicalmedicine*, **22**(100317).

International Institute for Population Sciences. 2021. *National Family Health Survey, India*. http://rchiips.org/nfhs/. [Accessed August 7, 2021].

J. Camenisch and A. Lysyanskaya. 2004. Signature Schemes and Anonymous Credentials from Bilinear Maps. *Pages 56–72 of:* Franklin, Matt (ed), *Advances in cryptology – crypto 2004*. Berlin, Heidelberg: Springer Berlin Heidelberg.

Khera, Reetika. 2019. *Dissent on Aadhaar: Big Data Meets Big Brother*. Orient BlackSwan. Edited volume.

Kleinberg, Jon M., Mullainathan, Sendhil, & Raghavan, Manish. 2016. Inherent trade-offs in the fair determination of risk scores. *Corr*, **abs/1609.05807**.

Lovell, Tammy. 2019 (February). *Swedish healthcare advice line stored 2.7 million patient phone calls on unprotected web server*. https://www.healthcareitnews.com/news/emea/swedish-healthcare-advice-line-stored-27-million-patient-phone-calls-unprotected-web. [Online; posted 20-February-2019].

Malhotra, Shefali, Garg, Rohin, & Rai, Shivangi. 2021. *Analysis of the NDHM Health Data Management Policy*. https://drive.google.com/file/d/1sEBg-syzsbe159x4PGkAHzcZilct0cQq/view. [Accessed September 21, 2021].

Matthan, Rahul. 2017 (July). *Beyond Consent: A New Paradigm for Data Protection - Discussion Document 2017-03*. http://takshashila.org.in/wp-content/uploads/2017/07/TDD-Beyond-Consent-Data-Protection-RM-2017-03.pdf.

Ministry of Health & Family Welfare, GOI. 2021. *Comprehensive National Nutrition Survey (CNNS)*. http://www.nhm.gov.in/index1.php?lang=1&level=2&sublinkid=1332&lid=713. [Accessed August 7, 2021].

Ministry of Health and Family Welfare, Government of India. 2019. *National Digital Health Blueprint*. https://ndhm.gov.in/home/ndhb. [Accessed September 21, 2021].

Ministry of Health and Family Welfare, Government of India. 2020a. *Cowin*. https://https://www.cowin.gov.in. [Accessed July 31, 2021].

Ministry of Health and Family Welfare, Government of India. 2020b. *National Digital Health Mission - Health Data Management Policy*. https://ndhm.gov.in/health_management_policy. [Accessed September 21, 2021].

Ministry of Women and Child Development, GOI. 2021. *Rapid Survey On Children (RSOC)*. https://wcd.nic.in/acts/rapid-survey-children-rsoc-2013-14. [Accessed August 7, 2021].

Narayanan, A., Paskov, H., Gong, N. Z., Bethencourt, J., Stefanov, E., Shin, E. C. R., & Song, D. 2012 (May). On the Feasibility of Internet-Scale Author Identification. *Pages 300–314 of: 2012 ieee symposium on security and privacy*.

Narayanan, Arvind, & Shmatikov, Vitaly. 2008. Robust De-anonymization of Large Sparse Datasets. *Pages 111–125 of: Proceedings of the 2008 ieee symposium on security and privacy*. SP '08. Washington, DC, USA: IEEE Computer Society.

Narayanan, Arvind, & Shmatikov, Vitaly. 2009. De-anonymizing Social Networks. *Pages 173–187 of: Proceedings of the 2009 30th ieee symposium on security and privacy*. SP '09. Washington, DC, USA: IEEE Computer Society.

Narayanan, Arvind, & Shmatikov, Vitaly. 2019. *Robust de-anonymization of large sparse datasets : a decade later*. http://randomwalker.info/publications/de-anonymization-retrospective.pdf.

Narayanan, Arvind, Shi, Elaine, & Rubinstein, Benjamin I. P. 2011. *Link Prediction by De-anonymization: How We Won the Kaggle Social Network Challenge*. https://arxiv.org/pdf/1102.4374.pdf.

National Informatics Centre, Ministry of Electronics and Information Technology, Government of India. 2020. *Aarogya Setu Mobile App*. https://www.mygov.in/aarogya-setu-app/. [Accessed May 31, 2020].

National Institute of Nutrition, ICMR, GOI. 2021. *National Nutrition Monitoring Bureau (NNMB)*. https://www.nin.res.in/researchdivision/publichealth.html. [Accessed August 7, 2021].

Office of the Registrar General & Census Commissioner, India. 2021. *Annual Health Survey (AHS)*. https://censusindia.gov.in/vital_statistics/AHSBulletins/ahs.html. [Accessed August 7, 2021].

O'Neil, Cathy. 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*. New York, NY, USA: Crown Publishing Group.

Rocher, Luc, Hendrickx, Julien M, & de Montjoye, Yves-Alexandre. 2019. Estimating the success of re-identifications in incomplete datasets using generative models. *Nature communications*, **10**(1), 3069.

Shrikanth, Siddarth, & Parkin, Benjamin. 2019 (July). *India plan to merge ID with health records raises privacy worries*. https://www.ft.com/content/4fbb2334-a864-11e9-984c-fac8325aaa04. [Online; posted 17-July-2019].

Solove, Daniel J. 2004. *The digital person: Technology and privacy in the information age*. New York, NY, USA: New York University Press.

Srikrishna, B. N., Sundararajan, Aruna, Pandey, Ajay Bhushan, Kumar, Ajay, Moona, Rajat, Rai, Gulshan, Krishnan, Rishikesha, Sengupta, Arghya, & Vedashree, Rama. 2017. *White Paper of the Committee of Experts on a Data Protection Framework for India*. [Online; Accessed January 9, 2018].

Su, Jessica, Shukla, Ansh, Goel, Sharad, & Narayanan, Arvind. 2017. De-anonymizing Web Browsing Data with Social Networks. *Pages 1261–1269 of: Proceedings of the 26th international conference on world wide web*. WWW 2017. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee.

Temperton, James. 2016. *NHS care.data scheme closed after years of controversy*. https://www.wired.co.uk/article/care-data-nhs-england-closed. [Online July 6, 2016].

Thatcher, Jim, O'Sullivan, David, & Mahmoudi, Dillon. 2016. Data Colonialism through Accumulation by Dispossession: New Metaphors for Daily Data. *Environment and planning d: Society and space*, **34**(6), 990–1006.

The Aadhaar judgment, 2018. 2018. *K S Puttaswamy and Another v Union of India (2018): Writ Petition (C ivil) No 494 of 2012, Supreme Court judgment dated 26 September*. https://www.supremecourtofindia.nic.in/supremecourt/2012/35071/35071_2012_Judgement_26-Sep-2018.pdf. [Accessed March 29, 2019].

The European Parliament and the Council of European Union. 2016. *Regulation (EU) no 2016/679*. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32016R0679.

The London School of Economics and Political Science. 2005 (June). *The Identity Project: An assessment of the UK Identity Cards Bill and its implications*. http://www.lse.ac.uk/management/research/identityproject/identityreport.pdf.

The Planning Commission: Government of India. 2011 (December). *Report of the group of experts on privacy chaired by Justice A P Shah*. http://planningcommission.nic.in/reports/genrep/rep_privacy.pdf.

Ullman, Jonathan. 2021. *Statistical Inference is Not a Privacy Violation*. https://differentialprivacy.org/inference-is-not-a-privacy-violation/. [Online June 3, 2021].

Unique Identification Authority of India. 2020. *Aadhaar*. https://uidai.gov.in. [Accessed May 31, 2020].

Zuboff, Shoshana. 2018. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. 1st edn. Profile Books.