Trust and ethical data handling in the healthcare context

Robin Wilton
Technical Outrech Director – Identity and Privacy
Internet Society

June 2017

wilton@isoc.org

Summary

Current practices for regulating the processing of personal data are oriented principally towards notions of governance, risk management, and regulatory compliance - based on data protection laws that have, in some jurisdictions, been in place for decades. Despite this framework, individuals are likely to encounter uses of their personal data which, while legal, may appear to lack fairness or legitimacy, through, for example:

- over collection,
- insufficient care taken with personal data,
- unexpected or unwelcome use,
- excessive sharing.

Such uses may lead users to conclude that third parties are failing to take due account of their wishes and preferences.

On the part of a data controller, a greater focus on ethics (as opposed to just legal compliance), might lead them to act in ways more likely to fulfil the expectations of their users and customers. Ethics has been core to the practice of medicine at least since the formulation of the Hippocratic oath (North 2012), but the digital era introduces new risks which require ethical evaluation and response. The resulting guidance for data controllers should be based on a clear understanding of digital privacy, and associated complexities, so that the abstract notions of trust and ethics can be transformed into applicable principles and practical measures while also reflecting the diverse motivations and interests of different stakeholders.

This article explores the trust-related factors and challenges that arise from the digital and online processing of personal data, particularly in the context of healthcare. The article proposes ethical principles, practical approaches, and resources for putting those principles into practice.

1. The relationship between trust, privacy and data ethics

1.1. Trust and the mediated nature of online services

Everything we do online is mediated through at least one third party: when we browse a website, we do so through the mediating services of our network provider; when we send email, we do so through the mediating services of at least one mail server; social media services mediate our online relationships with friends and family; payments to retail websites pass through third party payment service providers, and so on.

Our reliance on third parties, in our online dealings, exposes us to the risk that they will fall short of our expectations, in their collection and use of information about us. For example:

- Unexpected collection: in 2013, smart TV manufacturer LG acknowledged that its devices had been collecting data about owners' viewing habits without their consent; in 2014 and 2015, a wave of reports and news articles described the ways in which apps downloaded for one purpose (such as to provide a 'flashlight' function or a game) were also collecting unrelated data such as location and contacts, without users being made sufficiently aware of such collection or its purpose. (Albanesius 2013, Komando 2014, Raiche 2015)
- Insufficient care: in 2007, the UK tax authorities notoriously "lost" two CDs containing the personal and banking details of some 25 million citizens. This came as an unwelcome surprise to those affected, and was taken by some as an indication that the authorities had not taken sufficient care of the data entrusted to them. (Hope 2007). Large scale data breaches continue to happen, and continue to affect millions of citizens and consumers, as noted in the Internet Society's 2016 Global Internet Report - (Kende 2016).
- Unexpected/unwelcome use: the 2012 case of Target sending "expectant mother" vouchers to the household of a pregnant teen is frequently cited, albeit sometimes with caveats about whether it is genuine or merely anecdotal. Nevertheless, the cycle of customer profiling and unsolicited advertising gives rise to legitimate concern about such risks. (Hill 2012)
- Unexpected sharing: the growing market for data brokerage services, as part of a data monetisation ecosystem, has been worrying US regulators for several years, and yet is largely invisible to the data subjects concerned. (Singer 2012)

The mediated nature of online services and interactions also introduces a pervasive requirement for trust. To understand why, it may help to start with a candidate definition of trust; for some time now, I have used the following formulation: trust is a belief that someone will act in your interests, even if they have the opportunity and motivation to do otherwise.

Because online services are both remote and mediated, we have to rely on others, whose actions we may neither be able to control, nor inspect. For example, if Bob sends webmail to Alice, it is practically impossible for Bob to inspect the security mechanisms applied to the email by the service provider. If the service provider assures Bob that the email is encrypted in transit, Bob has little option but to take that assurance at face value. The same applies to any assurances they might give as to whether or not the email is encrypted while at rest (e.g. while waiting to be read). In the interests of good custody, an email service provider *could*, on receipt of Bob's emails to Alice, encrypt them for storage until Alice collects them – just as a postal sorting office might keep mail under lock and key until it is sent out for delivery. But in the normal course of events, Bob simply does not know whether his emails are sitting, in clear, on the service provider's servers.

In fact, following Edward Snowden's disclosures about government surveillance programmes, a number of webmail services were revealed to be less secure than their users might have expected (McCullagh - 2013): traffic between the sender's mail server and the recipient's mail server was, in many cases, sent in clear, and therefore open to being read by third parties. In the case of webmail services, some users may have relied on the presence of the "https" padlock symbol as an indication that their emails were confidential, although the https protocol in fact only secures the data while it is in transit, and not while the email contents remain stored at either end of the communications link. Since it is practically impossible for the end user to inspect the security mechanisms applied at the server, users have to trust that the service provider is taking appropriate measures to care for the data at that point.

When it becomes a question of the server-to-server communications provided to the email services, by telecommunications infrastructure providers (such as AT&T, Verizon, NTT, Telia, Telstra and others), it is even harder for end users to determine what security measures are in place to protect the data users entrust to those networks.

1.2. Privacy and contextual integrity

In 2004, Helen Nissenbaum published a paper on the idea of "contextual integrity" as a foundation for our understanding of privacy (Nissenbaum, 2004). In it, she argued that when we disclose information about ourselves, we do so implicitly, and in ways that still involve norms and expectations - about what it is, or is not, appropriate to do with the information that is disclosed.

Where we discover that the information we disclosed in one context - has been processed in another context (and therefore according to a different set of norms and expectations), we are likely to feel that our privacy has been violated.

As Nissenbaum also notes, however: *"One could read these cases simply as public policy disputes in which groups with opposing interests face off against one another, each seeking to promote its own goals, desires, preferences, and interests above those of opponents in the dispute."* This implies that one issue at stake is how to account for those interests appropriately, and resolve the resulting tensions. The issue to which these tensions give rise is, at least in part, an ethical one.

The more pervasive connected devices become in people's lives, the more likely it is that they will have to put their trust in the behaviour of the device, the apps it may run, or the services to which it connects. If users download a 'flashlight' app for their phone, they probably do not do so in the expectation that it will read and share their address book. That is not, after all, an expected thing for a flashlight to do. In a 2014 paper on data ethics by this author, this example was used to illustrate "the principle of no surprises" ([Author] 2014). If what a service provider does with data comes as a surprise to the data subject (especially if the surprise is an unwelcome one), we should ask whether there has been an ethical shortfall in the service provider's behaviour concerning collection, use, sharing, safe custody or disposal of the data.

The example of a flashlight app, harvesting contact data from an individual's phone, may seem like a rather small scale problem. However, the scope and scale of unexpected use of personal data is revealed in work by Sarah Spiekermann and Wolfie Christl, who researched corporate use of data for tracking and commercial use, for their 2016 report "Networks of Control" (Christl et al. 2016). Their findings consider (among other things):

- the use of personal data for behavioural prediction, marketing, financial services and employment decisions
- data brokerage as an industry
- mobile devices, apps, and connected objects.

The report gives a sobering perspective, from which an example is provided below, based on the collection and sharing of personal data by mobile phone apps. Christl and Spiekermann cite multiple case studies and surveys from 2010, 2011, 2012, 2014 and 2015, to show that not only is this a pervasive problem, but current legal safeguards do not appear to be constraining collection and sharing over time.

- In 2014, 26 data protection authorities surveyed 1200 apps across 19 countries, and concluded that 31% of them accessed data not needed for the app's functions, and 59% did not adequately inform the user about what would be done with data collected (Christl et al. 2016, page 49).
- In 2015, a survey of the top 100 free and paid-for apps, across the US, Australia, Brazil and Germany concluded that around 60% of paid apps were connected to tracking mechanisms that collect personal information, as were 85-95% of free apps. About 20% of paid apps were connected to more than 3 trackers (ibid. page 50).

Mobile apps are just one way in which new technology can give rise to the collection of health-related personal data. Other means include:

- "Wearables" such as fitness trackers
- "Smart" devices such as bathroom scales
- Devices for remote diagnosis and patient care
- Passive scanning of travellers' body temperature

Even devices whose primary purpose is not related to health can generate data with equivalent potential significance. For example, the accelerometers in a mobile handset can be used to analyse the gait of the person carrying it, and even detect leg injuries or falls (Lee R, Carlisle A, 2011).

2. Ethical Data Handling as an Organisational Approach

Given this background of data collection and monetization, many of us can probably recall some experience of discovering our personal data being used in a way that struck us as unexpected, unwanted, or unfair. Examples of such undesirable uses are:

- unexpected, unsolicited advertisements for products we have just mentioned on social media;
- unwanted third-party access to the contents of confidential messages;
- unfair pricing, offered on the basis of passively-disclosed information such as the make of laptop one is using.

And then, of course, there are the "unknown unknowns"[1] (Rumsfeld 2002): perhaps personal data held about us results in our being offered higher insurance premiums, or being declined credit, and it is possible we may never know what factors influenced these outcomes. Even if the processing of our data in such cases may be legal, we may find it undesirable nonetheless, and conclude that our interests are being ignored or overridden.

There is a distinction to be drawn between legal processing of personal data, and legitimate processing of personal data, and this article is written from the perspective that it is both valid and useful to draw that distinction. One can contend that there are some practices in the processing of personal data that may be legal, but are not legitimate: for instance, practices that pay lip service to the collection of meaningful user consent, but actually deliberately over-collect data for speculative future use. The goal of promoting ethical data handling is to encourage data controllers[2] to treat legal compliance as the minimum threshold, and adopt an approach to the handling of personal data that remains demonstrably legitimate when assessed against a wider set of values - such as fairness and respect - which are not necessarily codified in law.

An approach based on knowingly and explicitly exceeding what is required by law (in terms of notice and consent, for instance) would help build consumers' and citizens' trust in service providers, enhancing the latter's reputation and reinforcing service users' loyalty. In the commercial context, an improvement to consumer loyalty and brand perception represents a business benefit to the service provider (Hasselbalch et al. 2016). In the public sector, the potential benefit takes the form of a trust dividend which is critical to citizens' engagement with public services. In practice, an explicit ethical approach might lead organisations to give users clearer information about the further use of personal data -- such as its sale to third parties for use in targeted advertising -- or to adopt a default position of opting users *out* of data collection and data sharing, rather than opted-in by default (and putting the onus on the individual to find and exercise the option of opting out).

[1] A reference to Donald Rumsfeld's epistemological taxonomy, in a 2002 briefing.

[2] For the purposes of this paper, by "data controller" I mean an entity which collects and uses (personal) data in ways covered by data protection legislation; by "data subject" I mean an individual to whom personal data relates (I use "subject" in the grammatical sense, rather than in the sense of subordination to the data controller).

Current models of good practice in data protection tend to be based on a so-called "Governance, Risk and Compliance" (GRC) model, through which the data controller aims to minimise the risk associated with collecting and using personal data. This approach is supported by methodologies, professional accreditation bodies, and operational disciplines such as data/records management and audit. However, it falls short of guaranteeing ethical treatment, for a number of reasons.

- First, in some cases the mitigation (taking out insurance against the cost of a data breach, for instance) indemnifies the organisation, but does not protect the data subjects. In other words, there are ways of handling personal data that minimise the risk (to the organisation) of being penalised for lack of compliance, but still violate the privacy of the data subject.
- Second, in practice a GRC approach can result in a "check-list" mentality. Ensuring that 'all the audit boxes are ticked,' can blind the organisation to the real reason why handling personal data in a certain way is important: not just to minimise risk for the organisation, but also to ensure that personal data is handled in ways that genuinely reflect the interests of the data subject.
- Third, the GRC model grew principally in response to the use-cases that arise from "being a data controller" (i.e. having responsibilities under data protection law). It is open to question whether it has kept pace, for instance, with best practice as suggested by the *"Privacy by Design"* principles formulated by the then Ontario Privacy Commissioner, Ann Cavoukian (Cavoukian 2011).

Privacy by Design promotes a preventive approach to data protection: for example, by minimising the collection of personal data in the first place, data controllers can reduce their exposure to privacy-related risk from the outset. The *Privacy by Design* guidelines also explicitly call for "a clear commitment, at the highest levels, to set and enforce high standards of privacy − generally higher than the standards set out by global laws and regulation".

This paper is intended to suggest that there is scope for a collaborative approach, in which organisations develop a culture of ethical data handling, data subjects' interests are better respected, and organisations find that, as a consequence, their risks are lowered and their reputations enhanced. As a corollary: setting a high ethical bar for the organisation can, over time, make it easier for the organisation to achieve regulatory compliance. A culture of respect for personal data can motivate employees to see good data-handling practice as a benefit to the organisation, rather than an onerous compliance obligation or an exercise in box-ticking for its own sake. The example provided earlier (the HMRC data breach of 2007), would suggest that the HMRC staff in question gave more weight to internal processes and spending limits than they did to the ethical implications of disclosing unredacted personal details of 25 million citizens.

However, for a culture of ethical data handling to evolve, the rationale for it needs to be understood by those setting the organisation's strategy and priorities, and then put into practice through stated principles and operational practices. In that respect, I suggest that the ethical use of personal data needs to be given similar weight, in determining organisational behaviour, to that given to other corporate imperatives -- such as profitability in the commercial sector, or societal benefit in the public sector.

2.1. Risk and benefit, from the data subject's perspective

In the healthcare context, personal data can be exceptionally sensitive and intimate. Very few other situations require us to forfeit control (over our bodies, and our data) to such great extents, or with such serious potential consequences.

Certain types of foreseeable harm from the unethical use of patient data are clear, including details of the individual's health being made public without the individual's consent, or the individual suffering in terms of employment prospects, or insurance status. Both these types of harm could also adversely affect the relatives and dependents of the individual.

There are more pernicious forms of harm, too: an individual might be threatened with disclosure of sensitive records (blackmail), whether or not the threat comes from someone with authorised access to the data in question. In fact, if the blackmailer can make a credible *claim* of access to the data, actual access may not even be necessary (CNBC 2016)[3]. Healthcare records, which may touch on matters such as attempted suicide, abortion, mental health, domestic violence, and sexual abuse, can give rise to particularly grave risks, as illustrated by numerous examples in the public domain (Star 2016).[4]

Healthcare data can give rise to other kinds of risk which deserve careful attention. Genetic data, in particular, is revealing not just about the data subject, but about their parents, siblings and children. In this respect, a decision made on the basis of one person's data can, therefore, affect many others. Clinically, of course, that is one of the benefits of genetic data: it can allow accurate inferences to be made about genetically-related individuals, for instance about the likelihood of passing on inherited conditions.

Let us consider two ethical issues raised by this example.

First, under what circumstances is it acceptable to use one person's medical data in order to benefit someone else (or the wider population as a whole)? In the Helsinki Declaration on Ethical Principles for Medical Research Involving Human Subjects (Helsinki 1964), Clauses 6 and 8 acknowledge the beneficial intent of medical research (to understand disease and improve interventions), but also impose strict constraints. The goal of acquiring new knowledge, according to Clause 8, "can never take precedence over the rights and interests of individual research subjects". This has both theoretical and practical implications.

A strict interpretation of Clause 8 could be taken to mean that the privacy rights of an individual cannot be violated even if there is a broader medical benefit in doing so. One might counter that by claiming that a general practice patient is not a research subject in the relevant sense, and that therefore the constraints of the Helsinki Declaration need not apply. However, even though the patient may not be enrolled in a clinical research program, they will have implicit assumptions

[3] An analogous example: Blackmailers threaten to delete organisations' data unless ransom is paid.

[4] An illustrative but far from exhaustive sample of privacy violations and resulting harms, from a relatively privacy-conscious province.

about the purpose and scope of any personal data disclosures they make in order to receive treatment. The patient's primary intent, in visiting his/her general practitioner, is to be treated. On that basis, is it fair to take the details disclosed by that patient for the purposes of being made better, and re-purpose them for broader research? In terms of some countries' data protection laws, almost certainly not - unless that further purposing has been made clear to the patient, and their consent sought.

The Helsinki Declaration is just one point on a line that stretches from the Nuremberg Code (Nuremberg 1947), through the Belmont Report (Belmont 1978), and on to the much more recent Menlo Report (Menlo 2012). Where the Nuremberg, Helsinki and Belmont documents were aimed specifically at medical research, the Menlo Report attempts to take the principles expressed in the earlier reports and apply them to the ethical questions raised by ICT-related research.

Second, what expectations should the patient be entitled to have, regarding the safety of data about them? Particularly when consent is absent, for whatever reason, the data controller may take measures to ensure that any patient records used for broader research purposes cannot be associated with the patient(s) from whom they were derived. One such measure is the anonymisation of patient records. A practical problem with this approach is the increasing likelihood of supposedly anonymised data being re-identified (Narayanan et al. 2006), raising the risk that patient privacy is compromised after the event. A good deal of research, which goes beyond the scope of this article, looks at different models for what constitutes formal anonymity in sets of data (Dwork 2016).

It is worth noting that the reliable, long-term anonymisation of data is a highly technical topic, which one should expect to be beyond the normal scope of a data controller's expertise. If medical data is processed on the basis of anonymisation rather than consent, the question will be to what extent technology vendors implement strong anonymisation technology in data management solutions, and then to what extent these solutions are adopted and deployed. Adopters of such technology will have to find reliable ways to assess the trustworthiness of any claims made about its robustness; that, in turn, should guide their policies regarding retention, sharing, and use of the data.

2.2. Consent, control and agency.

In the healthcare context, patient consent is an important consideration, but cannot always be obtained or guaranteed - for instance, if the patient is incapacitated, or cannot understand what is being asked. As patients, we may, sometimes, have no option but to forego explicit consent, and rely on indirect alternatives, such as legal frameworks or the competence and ‘good will’ of the other party. *Trust* and *control* are inextricably linked, in ways we can explore by referring again to the candidate definition proposed above, in Section 1.1: "trust is the belief that someone will act in your interest, even if they have the means and the motivation to do otherwise".

In other words, the less control you are in a position to exercise, the more you have to rely on being able to trust the other party. Like any other belief, trust may be ill- or well-founded, and based on factors that are often systemic and diverse. For example, you might hear people say:

- "I trust her to keep my data confidential, because I've dealt with her before."
- "I trust her with my medical records, because she's professionally qualified."
- "I trust the clinic to keep my records private, because I've read their privacy policy."
- "I think they will keep my records secure, because they could get sued if they don't."
- "I'm not that worried about medical confidentiality - my data isn't that interesting to anyone."

All of these represent different foundations for the placement of someone's *trust* in the handling of their data. Without examining whether they are justified or not, we should be prepared to accept that when people make trust decisions, they may do so for many varied reasons, which may have little or no connection with the statements or behaviour of the data controller. If there is a wide divergence between the data controller's behaviour, and the patient's reasons for giving consent, we might want to question whether it is ethical for the data controller to rely on that consent. In simple terms: if there are facts about the data controller's behaviour which, if known to the individual, would cause the latter to withhold consent, then it is questionable whether that consent is "informed" in the relevant sense.

The examples, above, of people's reasons for trusting their doctor to handle personal data, are useful because they illustrate a broader principle - that of user agency: the individual's ability to influence the behaviour of systems and devices that have potential impact on their privacy. Where a subject cannot give informed, explicit consent, he/she forfeits agency and has to rely on others to act in his/her interest.

The same principle applies in many other (non-medical) uses of personal data; in fact, individuals' loss of agency is, arguably, one of the factors that characterises our modern, data-intensive, hyper-connected lives. "Smart" devices, and their corresponding web-based services, tend to reduce our involvement in deciding what happens to the data we generate by using them. The LG television sets referred to in Section 1 of this paper exposed a certain set of functions to the owner (the usual abilities to change channels, set volume, select inputs, and so on), but the device also had functions to do with the collection and onward transmission of data about the user, and the user had little or no control over these functions. A consumer might be a perfectly competent user of the device *as a television set,* but lack the means to control the ways in which the television also acts as a device for collecting and forwarding personal data.

This places (or should place) an additional burden on the provider of the service or device, to respect the intentions and preferences of the data subject, and to treat their data accordingly.

The general design question - whether in the healthcare context or elsewhere - is: "what is it, in this system, that gives effect to the user's intentions and preferences?". The answer could be a technical component, a process, or a third party... but if the answer is "nothing", the ethical issue is not being addressed.

2.3. Realism about incentives

When considering how to improve on the way organisations are disposed to use personal data, we should acknowledge that they can have strong - even compelling - reasons to treat personal data as they do, and that it may be hard for them to shift from the economics of the *status quo*. If a business derives most of its revenue from the monetisation of personal data, then a shift to a model that constrains its use of personal data may threaten that revenue stream and, ultimately, the profitability of the business. If we seek to change the behaviour of the business in question, without acknowledging the commercial imperative, we are unlikely to succeed.

This problem is compounded by what economists call "negative externalities". If the benefits of monetizing data fall mainly to the organisation, but the risks and drawbacks fall mainly to the data subject, there may be little apparent incentive for the data controller to behave otherwise. If a supposedly free service is in fact funded by the proceeds of monetizing personal data, that represents a cost to the consumer, and one of which they may be unaware. If the true cost of monetisation is not evident to consumers, it does not influence their behaviour sufficiently to influence, in turn, the bargain offered by the service provider. Where this is the case, the market does not accurately reflect the interests of the buyer. This principle is examined in the Internet Society's Global Internet Report for 2016, which explores the hidden costs and negative externalities associated with data breaches, and makes a number of recommendations about how data controllers can improve their handling of personal data (Kende 2016).

3. The elements of a solution

3.1. "Point" solutions and systemic problems

Like other challenges presented to us by our increasingly data-driven society (such as trust, privacy and security), ethical data handling is a ‘systemic issue,’ for which there is unlikely to be a "point" solution: if the technology is there, but users aren't motivated to adopt it; if the economics of personal data create an overwhelming incentive for unethical use; if users perceive a need but the technology is too hard to use, then the systemic problem will persist.

The reality is that success frequently requires a *set* of point solutions, each of which addresses a particular part of the problem: a point solution that generates user awareness of the problem, a point solution that produces good, usable technology, and a point solution that creates the right regulatory and economic circumstances for the solution to thrive.

Perhaps that is why so many of these systemic problems remain persistently thorny and seemingly unsolved. Based on this author's experience of the Internet Society's approach to systemic problems in Internet-related domains four principal criteria suggest themselves:

- ensure that all the stakeholders in a system are engaged in the search for solutions;
- look at the whole lifecycle of the information in question, including the aspects of economics, regulation, technology, and user motivation;
- recognise that each aspect of the lifecycle will require different kinds of intervention;
- design for sustainable change.

3.2. Abstract problems need practical answers

Apart from their systemic nature, problems of trust, privacy, ethics, and so on can also be hard to address just because they are abstract concepts. "Build me an ethical product," is a worthy request, but one which does not give an engineer, or developer, sufficient practical guidance. However, reference works are emerging which show how the abstract concepts can be broken down into specific, practical features that contribute to trust, privacy, and ethics in the resulting product (Spiekermann 2016) (Dennedy 2014).

For example, one factor in trust, as mentioned earlier, is the ‘consent’ of the data subject. For consent to be ethical, the data subject needs to be adequately informed about the collection and use of the data, and to understand the information they are given; the data controller needs to collect consent through the knowing, competent and freely-given action of the data subject. These are ethical foundations for the seeking of consent. They are related to, but distinct from, the legal and regulatory requirements for gaining consent to the collection and processing of personal data, which may differ significantly from one jurisdiction to another. However, that raises a set of cross-jurisdictional issues, the resolution of which goes well beyond the scope of this paper.

The consent criteria used in the example above can be translated into specific technical measures in, for example, the transaction flow, the user experience, the user interface, and the way in which relevant information was made available to the user. This is not to pre-suppose that designing an ethical application is easy. The level of detail one person might consider to be necessary for an informed decision might appear to another person as excessive, for instance. So who should make those choices, and on what basis? The practice of involving an institutional review board (IRB) or independent ethics committee (IEC) is currently most commonly applied to the supervision of projects involving medical research and clinical trials, but as technical innovation reaches further and further into our lives, it introduces the same levels of intrusiveness and potential harm as the kinds of project for which IRBs are mandated.
As the scope and consequences of personal data collection become ever more significant in our daily lives, I would argue that the IRB/IEC discipline deserves to be applied more widely.

One option might be to mandate a privacy impact assessment (PIA) for projects involving personal data, and for one possible outcome of the PIA to be a recommendation to appoint an IRB. As the former privacy commissioner of Ontario, Ann Cavoukian, put it recently: *"ethics has always formed an essential component of privacy and should be reflected in any privacy by design assessment."* (Cavoukian 2016)

Arriving at the ethical solution might require an adaptive approach from the designer and from the application itself: for instance, an initial, detailed expression of consent might be followed by subsequent periodic confirmatory checks, with the goal of maintaining the legitimacy of the user's consent. In this example, one could consider three possible cases:

1. The user is asked, once and for all, for a blanket expression of consent. This approach is currently common, and does not cater well for possible changes in the user's attitude towards consent over time.
2. The user is reminded, on every access to the service, of the ‘consent’ arrangement/agreement currently in force, and is expected to confirm it before being allowed to proceed. This approach may provide regular evidence that the user has clicked to confirm consent, but may also lead to "click fatigue", where the user comes to consider the consent request as an inconvenience, and simply clicks past it without considering its implications.
3. The user is invited to confirm consent periodically, on the basis of criteria that are clearly explained. For instance, a six-monthly "courtesy check" in case the user's attitude to consent has changed; or an opportunity to turn off discretionary messages from the service provider.

Note that these do not lessen the service provider's obligation to meet legal criteria for consent (such as the need to seek consent if there is a material change in the way personal data is used, relative to the purpose for which consent was originally sought). The third option includes measures for transparency and control, as contributing factors to users' trust in the system.

Explicit consent can also be particularly hard to achieve, when data is being collected through passive means (such as CCTV or security cameras). In such instances, again, we have to ask what it is, in the system, that gives effect to the legitimate expectations of the data subject, bearing in mind that the view of what constitutes "legitimate expectations" may well vary from one culture to another.

But the over-arching principle is that the abstract ideas of *trust* and *consent* can be broken down into elements that can be functionally described and technically implemented.

In addition to consent, the same approach can be applied to other trust-enhancing factors:

- the data subject's ability to review and correct or delete the data held about them;
- user controls over which pieces of data are disclosed, shared or retained;
- management of the specific context in which data is collected and used.

3.3. The role of “context”

It is difficult - and, some would argue, pointless - to produce a single, canonical definition of "privacy". Certainly, privacy can have many different facets or dimensions - socially, personally, technically, legally and so on. For instance, Warren and Brandeis described it as "the right to be let alone" (Warren et al. 1890), a concept with social connotations of self-determination and freedom from outside influence. In doing so, they took a phrase used by Thomas Cooley (Cooley 1888), in a work on tort law, in which he referred to "the right to one's person" as a "right of complete immunity: to be let alone". Warren and Brandeis applied it explicitly to an immunity from the non-physical harm of an invasion of privacy. The European Union's present-day data protection principles express privacy in legal terms that are still couched in a social context: they describe it as one of the fundamental rights and freedoms of natural persons, "in particular with respect to the processing of personal data" (European Commission 1995).

To turn these social and legal principles from rights-based ideals into practical implementation requires a further step, in support of which the Internet Society (Internet Society 2017) formulated the following definition of privacy:

> *"Privacy is about retaining the ability to disclose data consensually, and with expectations regarding the context and scope of sharing."*

This definition intentionally binds privacy to the notion of context, so as to reinforce the idea that with changes in context may come changes in expectation and jurisdictional requirements.

Health data and its use extend across a very broad range of contexts, with scope ranging from the individual to the species. An individual's medical data might identify them uniquely, or as a member of a family, or as someone who stayed in the same hotel as a group of others, or as a member of a particular ethnic group, or as someone with a particular genetic disposition, and so on.

We should therefore perhaps not expect a single ethical formula to work for all cases. In particular, as we see in the domain of privacy, different social, economic and cultural expectations about healthcare lead to different expectations about what is acceptable in terms of ethical and privacy trade-offs for the use of health data. The Internet, as an enabling infrastructure, is notoriously poor at respecting organisational, cultural, jurisdictional or national boundaries, and makes it easy for data - including health data - to travel across all of them. Digital technology makes it incredibly easy to replicate, share and distribute data without regard for the context in which that data originated. When that happens, whether through design, indifference or malice, both privacy and the ethical treatment of data are at risk.

4. Summary and future directions

4.1. Trust

Users' trust in the collection and use of health data will be based on their expectations, which are contextual and may change over time. But users' expectations are influenced by a set of factors including : previous experience, perception of risk/harm, transparency of behaviour, the likelihood of effective remedy, among others - and it one could conjecture that this set of factors is likely to remain consistent. If trust is a belief that others will act in your interest, even if they have the means and the motivation to do otherwise, and if we believe that health data presents many opportunities for your interests to be put at risk, we ought to conclude that trust is a key factor in determining how to process health data appropriately. The next section considers increased transparency as an element of trustworthy processing of personal data.

4.2. Transparency

If we are to raise ethical standards in the use of personal data, there are certain aspects concerning which the data subject should be better informed than tends currently to be the case:

- evidence of fairness in the way personal information is used,
- respect for personal preferences,
- measures to provide enforcement, redress and reconciliation, when personal data is misused.

The more absent or invisible such aspects are, the harder it is for users to determine whether the system is trustworthy. But, as Michelle Dennedy and Sarah Spiekermann explain in their respective books, values such as fairness, respect, and redress can be translated into practical measures that keep the user appropriately informed, make transactions auditable, and provide evidence in support of dispute resolution. The problem of translating these abstract values into technical solutions is not an insoluble one, and there is a growing set of practical resources available to those organisations that are motivated to solve it.

4.3. User agency versus user intervention

All of the above may seem to place a heavier burden on the individual than they currently expect when using online services, since we are implicitly expecting users to engage with, interpret and act on information that is presented to them. We are also expecting them to make rational decisions and be accountable for the consequences. Experience in other areas, such as websites' implementation of the EU cookie directive (Kobie 2015), or browser plug-ins to monitor the replacement of public key certificates, suggests that more interaction with the user is not always better. When applications make too frequent requests for user interaction or confirmation, the net effect can be that users "click through" without thought, or that they disable the supposedly helpful function for the sake of a less interrupted experience. Both of these outcomes defeat the purpose of the helpful function.

But if the right design principles are put into practice (for instance, using the reference works cited previously), the individual's intentions and preferences can be put into practice without their necessarily having to take explicit action in every instance.

A somewhat rudimentary example of this is the way in which an email application might allow the user to specify rules and filters to be applied to some or all incoming mails. Having set the rules, the user can rely on the application to enforce them automatically with no further input. For the sake of this example, we rely on the user knowing how to set filter rules, and we also assume that the technology can enforce this preference without rendering the email application unusable. However, since those assumptions are not always reliable, we should expect to see similar design principles expressed, in future, in far more sophisticated ways, such as the use of machine learning to draw inferences about the user's preferences from an initial set of observations. In this way, it is conceivable that an adaptive "user proxy" could act on the user's behalf to express privacy preferences, perhaps validating those choices directly with the user from time to time (but not on every occasion).

We are surrounded by computing power in the form of intelligent devices, smart objects, increasingly autonomous systems, wearable computers, implants and so on. As a direct consequence, just by going about our daily lives, we generate ever-increasing volumes of data, much of which is harvested and monetised by others. We should expect more of that computing power to be put to work collecting, understanding and applying *our* preferences, and restoring to us some of the agency which current practices have eroded.

For that to happen, the developers of applications, services and devices must embody ethical principles in the design decisions they take. That, in turn, means translating abstract social concepts into manageable, practical functions, implemented in technology, rules and processes. The discipline of doing this may still be in its early stages, but the availability of reference guides such as those by Dennedy (2014), Spiekermann (2016), Hasselbalch (2016) and others give good reason to believe it is possible, and that it can bring benefit to the organisation as well as to the data subject.

5. References

North M, (2012) on nlm.nih.gov "Greek Medicine - The Hippocratic Oath" U.S.
National Library of Medicine, 02 July 2012.
<http://www.nlm.nih.gov/hmd/greek/greek_oath.html>
Accessed 19 June 2017

Hope C, (2007) on telegraph.co.uk
http://www.telegraph.co.uk/news/uknews/1570258/5000-would-have-made-HMRC-discs-safe.html
Accessed 16 Dec 2016

Albanesius C, (2013) on ukpcmag.com :
http://uk.pcmag.com/tv-home-theaters/11794/news/lg-to-fix-unwanted-smart-tv-data-collection
Accessed 16 Dec 2016

Komando K, (2014) on usatoday.com :
http://www.usatoday.com/story/tech/columnist/komando/2014/11/14/free-apps-privacy/18759109/
Accessed 16 Dec 2016

Raiche R, (2015) on wptv.com : http://www.wptv.com/news/science-tech/angry-birds-2-camscanner-apps-stealing-personal-data-raises-questions-about-iphone-security
Accessed 16 Dec 2016

Kende M, (2016) Global Internet Report
https://www.internetsociety.org/globalinternetreport/2016/
Accessed 16 Dec 2016

Hill K, (2012) on forbes.com :
http://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#2f7c6a3934c6
Accessed 16 Dec 2016
Singer S, (2012) in nytimes.com :
http://www.nytimes.com/2012/06/17/technology/acxiom-the-quiet-giant-of-consumer-database-marketing.html
Accessed 16 Dec 2016

McCullagh M, (2013) in cnet.com :
https://www.cnet.com/news/how-web-mail-providers-leave-door-open-for-nsa-surveillance/
Accessed 16 Dec 2016

Nissenbaum H, (2004) Privacy As Contextual Integrity. Washington Law Review

[Author], (2014) Four Ethical Issues In Online Trust. Internet Society
https://www.internetsociety.org/sites/default/files/Ethical%20Data-handling%20-%20v2.0.pdf

Christl W, Spiekermann S, (2016) Networks Of Control
http://crackedlabs.org/en/networksofcontrol

Yoshida T et al., (2006) Gait analysis for detecting a leg accident with an accelerometer
https://www.researchgate.net/publication/4238356

Lee R, Carlisle A, (2011) Detection of falls using accelerometers and mobile phone technology
https://oup.silverchair-cdn.com/oup/backfile/Content_public/Journal/ageing/40/6/10.1093/ageing/afr050/2/afr050.pdf

Rumsfeld D, (2002) US Department of Defense news briefing transcript -
http://archive.defense.gov/Transcripts/Transcript.aspx?TranscriptID=2636
Accessed 19 June 2017

Hasselbalch G, Tranberg P, (2016) Data Ethics, The New Competitive Advantage. Libris, Copenhagen

Cavoukian Dr Ann, (2011) Privacy By Design - Applying The 7 Foundational Principles
https://www.iab.org/wp-content/IAB-uploads/2011/03/fred_carter.pdf
Accessed 19 Dec 2016

Helsinki Declaration (1964) World Medical Association -
http://www.wma.net/en/30publications/10policies/b3/
Accessed 1 Oct 2016

Nuremberg (1947) Mitscherlich A, Mielke F. Doctors of infamy: the story of the Nazi medical crimes. New York: Schuman, 1949: xxiii-xxv.

Belmont (1978) US Dept of Health and Human Services, The Belmont Report https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/index.html

Menlo (2012) US Dept of Homeland Security, The Menlo Report - Ethical Principles Guiding Information and Communication Technology Research

Ramesh R, (2015) NHS disregards patient requests to opt out of sharing medical records
https://www.theguardian.com/society/2015/jan/22/nhs-disregards-patients-requests-sharing-medical-records
Accessed 20 Dec 2016

Narayanan A, Shmatikov V, (2006) Robust De-anonymization of Large Sparse Datasets
https://www.cs.utexas.edu/~shmat/shmat_oak08netflix.pdf
Accessed 1 Oct 2016

CNBC (2016) Hackers blackmail US Police Departments
http://www.cnbc.com/2016/04/26/ransomware-hackers-blackmail-us-police-departments.html
Accessed 1 Oct 2016

Ontario Star (2015) Hospital privacy violations go unreported
https://www.thestar.com/life/health_wellness/2015/01/13/hundreds_of_hospital_privacy_violations_go_unreported.html
Accessed 1 Oct 2016

Dwork C, Roth A (2014) The Algorithmic Foundations of Differential Privacy
Foundations and Trends in Theoretical Computer Science Vol. 9, Nos. 3–4 (2014) 211–407

Spiekermann S (2016) Ethical IT Innovation. Auerbach

Dennedy M (2014) The Privacy Engineer's Manifesto. Springer

Cavoukian, Dr. Ann (2016) Twitter:
https://twitter.com/anncavoukian/status/803566281190424577

Warren S, Brandeis L (1890) Harvard Law Review December 1890
http://faculty.uml.edu/sgallagher/Brandeisprivacy.htm
Accessed 1 Oct 2016

Cooley T (1888) A Treatise On The Law Of Torts. Callahan & Co., Chicago

European Commission (1995) Directive 95/46 Article 1, and Preamble para. 2

Internet Society (2017) Why Privacy Matters
http://www.internetsociety.org/what-we-do/internet-technology-matters/privacy-identity

Kobie N, (2015) Why the cookies law wasn't fully baked

https://www.theguardian.com/technology/2015/mar/19/cookies-how-to-avoid-being-tracked-online
Accessed 20th Dec 2016

Doliner M, (2016) Safari invalid certificate handling sucks
https://kingant.net/2016/08/safari-invalid-certificate-handling-sucks/
Accessed 20th Dec 2016