AI & SOCIETY (2022) 37:1715–1728
https://doi.org/10.1007/s00146-021-01262-5

OPEN FORUM

AI, big data, and the future of consent

Adam J. Andreotta[1] · Nin Kirkham[2] · Marco Rizzi[3]

Received: 12 March 2021 / Accepted: 18 August 2021 / Published online: 30 August 2021

Abstract
In this paper, we discuss several problems with current Big data practices which, we claim, seriously erode the role of informed consent as it pertains to the use of personal information. To illustrate these problems, we consider how the notion of informed consent has been understood and operationalised in the ethical regulation of biomedical research (and medical practices, more broadly) and compare this with current Big data practices. We do so by first discussing three types of problems that can impede informed consent with respect to Big data use. First, we discuss the transparency (or explanation) problem. Second, we discuss the re-repurposed data problem. Third, we discuss the meaningful alternatives problem. In the final section of the paper, we suggest some solutions to these problems. In particular, we propose that the use of personal data for commercial and administrative objectives could be subject to a 'soft governance' ethical regulation, akin to the way that all projects involving human participants (e.g., social science projects, human medical data and tissue use) are regulated in Australia through the Human Research Ethics Committees (HRECs). We also consider alternatives to the standard consent forms, and privacy policies, that could make use of some of the latest research focussed on the usability of pictorial legal contracts.

Keywords Big data · AI · Privacy · Informed consent · Moral responsibility

1 Introduction

Despite there being clear epistemic and practical benefits from applying AI algorithms to large datasets containing our personal information—commonly referred to as 'Big data'—some recent high profile cases have raised a series of moral and legal concerns, with respect to the behaviour of the companies using these datasets.[1] To name some high profile examples, in 2018, it was revealed that British consulting firm Cambridge Analytica were allowed by Facebook to harvest the personal data of over 80 million of its users (without their permission). Cambridge Analytica then used these data to target American voters in the 2016 presidential election (Isaac and Singer 2019). Google owned YouTube

✉ Adam J. Andreotta
adamandreotta@outlook.com

Nin Kirkham
nin.kirkham@uwa.edu.au

Marco Rizzi
marco.rizzi@uwa.edu.au

[1] School of Management, Curtin University, Kent St, Bentley WA 6102, Australia

[2] Department of Philosophy, The University of Western Australia, 35 Stirling Hwy, Crawley WA 6009, Australia

[3] UWA Law School, The University of Western Australia, 35 Stirling Hwy, Crawley WA 6009, Australia

[1] One problem with the term 'Big Data', as Luciano Floridi (2012) points out, is that it is slightly ambiguous, since the predicate 'big' is vague. In other words, there is no precise point at which a dataset changes from small to big. In this paper we will use the term 'Big Data' in the sense that is most commonly adopted at present—namely, to describe data sets (of ever-increasing sizes) that are too big for humans to analyse for the purpose of identifying new patterns, correlations, and insights. AI algorithms become useful in these domains due to the speed and scale at which they can operate. An ethical issue arises here because AI algorithms have the potential to reveal novel forms of personal information from such data sets. Individuals may have a strong desire for such personal information not to be made public, shared to third parties, or used to modify their behaviour. In short, there is a risk that serious harm can be caused to individuals by the improper use of AI and big data. For a more precise definition of Big Data, see Levin et al. (2015). They characterise Big Data in terms of four key attributes—namely Volume, which refers to the terabytes of new data being added each day; Velocity, which refers to the real time speed at which analyses can now be performed

were fined $US170 million (Singer and Conger 2019) for extracting personal information, without parental consent, from children using the platform, and then using the data to target advertisements towards them. The British parenting club, Bounty, were fined £400,000 in 2019 for sharing data from over 14 million of its users to third parties for marketing purposes (Postelnicu 2019). In 2016, DeepMind Technologies Ltd (a Google subsidiary) initiated a collaboration with the Royal Free London NHS Foundation to train machine-learning algorithms capable of assisting with the management of acute kidney injury. The opacity of the terms of this collaboration have raised a number of pressing questions regarding the protection of privacy, the regulation of data sharing and use, and the structural difficulties to implement a measure of meaningful individual consent and control in the face of increasingly common transfers of population-derived datasets to large private companies (Powles and Hodson 2017).

These cases, and others like them, show that express consent (consent pertaining specifically to the activity to which the consent is given) has not always been sought, let alone received, by companies who use, share, and/or sell, peoples' personal data. This is not only a problem relevant to the big four companies (Amazon, Apple, Facebook, and Google) rather it affects all businesses, and governments, who collect personal information. On the face of it, the best solution to this problem may seem to be the introduction of stricter legislation, with greater penalties for those who fail to gain express consent. However, such an approach, we argue, will provide only at best a partial solution and at worst the mere appearance of a solution. The introduction of ever more detailed terms and conditions forms for users to read, or more 'policy acceptance' boxes for users to tick, *prima facie* may allow companies to secure greater levels of express consent, but it will make the question of whether that express consent amounts to *informed consent* only more complicated, not less. A distinct course of action is that of adopting a stringent rule-based approach to permissible and prohibited conducts. The European General Data Protection Regulation (GDPR) is the prime example of this approach to privacy and data protection (GDPR 2018). The adoption of top-down command-and-control regulatory action in the face of structural power imbalances that can give rise to the large-scale abuses of the sort exemplified above is certainly valid and in all likelihood necessary. There remains, however, scope to argue for the irreplaceable role of individual autonomy, given the centrality of individual subjects in the data production process. This necessary role calls for novel and creative forms of protection attuned to the modern social landscape. This paper focuses squarely on this limb of the Big data rebus.

We do so by focusing on 'informed consent', a concept that takes into account our own psychological constitution, supports rational decision making, preserves autonomy, and respects individuals. For these reasons, it is a significant moral requirement. In this paper, we discuss several problems with current Big data practices which, we claim, seriously erode the significance of informed consent. To illustrate these issues, we consider, in Sect. 2, how the notion of informed consent has been understood and operationalised in the ethical regulation of biomedical research and contemporary clinical practices. We then compare this with current Big data practices. In light of this, we articulate some challenges that current Big data practices face, in Sect. 3; and, finally, in Sect. 4, we sketch some ways in which informed consent may be secured more reliably and meaningfully than is currently the case.[2] In particular, we propose that the use of personal data for some commercial and administrative objectives could be subject to a 'soft governance' ethical regulation, akin to the way that all projects involving human participants (e.g., social science projects, human medical data and tissue use) are regulated in Australia through the Human Research Ethics Committees (HRECs).[3] We also consider alternatives to the standard consent forms, and privacy policies, that could make use of some of the latest research focussed on the usability of pictorial legal contracts.

2 Informed consent

By 'informed consent' we have in mind the sense articulated by Tom Beauchamp, who says—in the context of medicine, clinical practices, and biomedical research—that 'A person gives an informed consent…if and only if the person, with substantial understanding and in substantial absence of control by others, intentionally authorizes a health professional to do something' (2011, p. 517–518). There is a clear ethical dimension to informed consent because, on the standard view, it (ideally) facilitates the transfer of information between two parties, e.g., a doctor and patient (Manson and O'Neill 2007, p. 27). In biomedical contexts, for example,

Footnote 1 (continued)

on these data; Variety, which refers to the different types of data, and variety of sources, that are now being collected; and Veracity, which pertains to the trustworthiness of the data sources (Levin et al. 2015 pp. 1661–1662).

[2] In a literature review on the ethics of Big Data by Mittelstadt and Floridi (2016), it was found that informed consent was one of the biggest concerns of researchers.

[3] See also the 'National Statement on Ethical Conduct in Human Research' (National Health and Medical Research Council 2007 [updated 2018]). This statement provides ethical guidance for Australian researchers whose work involving human subjects.

such information allows a patient to become aware of the potential risks that may arise if they consent to a procedure or determine whether consenting to a procedure may conflict with their values or preferences. As important as information transfer is, there may be dangers in reducing informed consent to this property. Manson and O'Neill, for instance, highlight a communicative element that is important for respecting agency, that may not always be captured in an 'information transfer' (2007, p. 62). After all, it is not only important that patients are given the relevant information, but that they understand and retain it (Kadam 2017).

Informed consent, as Tom Walker (2020) notes, can also provide a kind of symbolic value for patients, because it acknowledges them as decision makers and recognises their personhood. For example, if a person is not consulted before a medical procedure is performed on them, or their personal information is shared without their knowledge, they may feel dehumanised because they have been left out of the decision-making process. A person may wish to opt out of the decision-making process, of course; what matters is that they had the 'opportunity to choose' (Walker 2020, p. 2). O'Neill (2003), further, notes that while patient autonomy is important to focus on, informed consent also matters because it can help to ensure that patients have 'not been deceived or coerced' (2003, p. 5). While scholars have debated the importance each of these different features of informed consent, this brief discussion highlights the reason why securing informed consent is an important ethical requirement.

Over the past 70 years, the development of several ethical principles and codes of ethics has resulted in a greater focus on informed consent in the biomedical context. These include: the Nuremberg code—first created in 1947 (following the Nazi Doctors' Trial) which stated that voluntary consent ought to be sought from patients before undertaking procedures; The Declaration of Helsinki—first adopted in 1964 (revised several times since), which also focused on informed consent, but provided explicit recognition to the vulnerability of certain individuals or groups, who may be incapable of giving consent; and the creation of governing bodies which make recommendations and offer advice to practitioners. In Australia, for example, The Royal Australian College of General Practitioners (RACGP) and The Royal Australian and New Zealand College of Obstetricians and Gynaecologists (RANZCOG) are organisations who are responsible for maintaining ethical standards in their respective fields. One of the RACGP's guidelines, for example, tell practitioners to ask whether 'a "reasonable" person (in the same position) if warned of the risk is likely to attach significance to it' RAGCP 2019).[4]

RACGP and RANZCOG also recommend that practitioners become aware of the power imbalances that may arise between health professionals and patients (RANZCOG 2018). This is because a patient's decision-making abilities can be affected by who it is that informs them (Nimmon and Stenfors-Hayes 2016). Further, practitioners are encouraged to check if their patients really have understood what they have been told. For example, they can ask their patients to explain, in their own words, what the implications of a certain procedure are.

None of this is to say that modern attempts to integrate informed consent into biomedical settings are always successful (Beauchamp 2011). It is only to recognise the extent to which informed consent is now taken seriously by medical professionals. When it comes to Big data use, conversely, we argue that informed consent is not taken as seriously yet. Before laying out three of the main problems with current Big data practices, as we see them, it is worth reflecting on the asymmetry between the seriousness in which informed consent is dealt with in the medical context; and the perfunctory manner in which it is dealt with by big technology companies.

One explanation for this asymmetry is that informed consent in the medical context can be associated with procedures that put a person at risk of death or serious injury, whereas allowing one's personal information to be used, or having one's privacy invaded in some way, might be seen as less of a significant risk inasmuch as it does not usually result in physical harm. However, the kinds of harms that might occur as a result of having one's personal information made public can be very serious. So, it may be that the problem is one of risk perception (Sunstein 2002), where many people have not yet appreciated the relevant risks

[4] It also worth remembering that the concept of informed consent has not always been considered integral to medical ethics. If we look at the Hippocratic physicians of ancient Greece, we not only find a lack of concern for informed consent but also an absence of concern for the truth. The *Corpus Hippocraticum* (the corpus of early medical texts associated with Hippocrates), for example, for all its innovation and focus on the responsibilities of physicians, features instructions to conceal information from the patient where doing so would be useful (Faden and Beauchamp 1986, p. 61).

Footnote 4 (continued)

Interestingly, this formulation echoes the findings of the High Court of Australia in the landmark medical negligence case of *Rogers v Whitaker* (1992) 175 CLR 479. The issue was whether the failure to warn a patient, who was about to undergo eye surgery, of a very unlikely risk constituted negligence on the part of the surgeon. With this decision the court moved past the traditional 'doctor knows best' approach (whereby the decision on whether or not to warn of a certain risk fell within the discretion of the health professional) and embraced a doctrine that upholds the autonomy of the individual patient and their ability to attach significance to particular risks (Sappideen 2010).

enough to expect the kind of informed consent practices that they expect in medical contexts. Given the relative novelty of some of the AI technologies, this is unsurprising even though it is not a justification.

It is important to secure an individual's informed consent with respect to Big data use in our view because, just like in the medical context, failing to do so can cause harm to individuals. For example, recently 5 million facial images were captured by facial-recognition technology in 12 Canadian shopping malls, without shoppers' consent (Bronskill 2020). Upon finding out, some shoppers may have felt dehumanised from having their facial images captured and analysed without first being consulted. Shoppers may be concerned, justifiably, that their information would be sold to third parties or analysed by AI algorithms for marketing purposes (Bronskill 2020). Others could feel their rights as citizens were being violated. Similarly, an underage teenager whose browsing activities are collected and then who is sent targeted alcohol and gambling advertisements to their social media account, with or without consent, may be harmed. The ads may not only be inappropriate given their age, but they may cause direct harm, as they might feel like they are 'being spied on' (Duffy 2021). Data users, intuitively, have a moral responsibility to prevent such harms from occurring (Macnish and Gauttier 2020, p. 52).

As the public has become aware of the economic gains AI and Big data use can bring, opinions about their use have started to form. A recent Eurobarometer survey, for instance, showed that a majority (53%) of participants said that they were 'uncomfortable' with their internet companies using their personal information to tailor ads towards them (Eurobarometer 2015, p. 39). And about a third of those participants said they were 'very uncomfortable' (ibid) with the same act. While the concept of privacy is matter of contention amongst scholars—that is, whether it is one thing, or a series of different things (Solove 2008)—a basic level of privacy remains important to most of us. This is because having control over what we keep private allows us to control what others know about us, it lets us set the levels of intimacy we have with certain people and gives us control about how we shape own our personhood (Solove 2008 Ch. 2).

While Big data use is still a large problem, some companies have responded with practical changes. Apple's recently released operating system IOS14, for example, requires users to first grant apps permission (to opt in) before their personal information is collected. Given the profitability of targeted advertising, some companies like Facebook have pushed back against such measures, claiming that this would hurt small businesses (Thorbecke 2021). Apple's measures would not mean that advertising would cease, of course; however, without users' personal information, the advertisements would be less targeted and thus less profitable (Purtill 2021). Whether or not a company like Apple thinks that consent is a *moral right* that they have a duty to uphold, or whether they have recognised that their customers have a *desire* to be in control of what is done with their personal information, or whether they think a focus on privacy could boost their reputation is hard to say. Pollach (2011, p. 94), has shown that companies claim to be motivated by each of these three reasons. What is clear is that the problem cannot be ignored.

The legislative situation in Australia leaves much to be desired as recently analysed in a report by the Australian Competition and Consumer Commission (ACCC) on Digital Platforms, which made recommendations to adopt several reforms to the loose framework currently provided by the *Privacy Act 1988* (Cth).[5] Other jurisdictions have started to address the seriousness of the issue. The EU, for example, has tried (with mixed results) to put some form of legislative watchdog in place—namely, the GDPR. The GDPR is an immensely ambitious and important piece of legation in this sense, but arguably it cannot fully succeed precisely because of how cumbersome it is (Quelle 2018). Several factors are at play here. As we point out below, there is the issue of individual awareness of the degree of potential harm. This is in part caused by a degree of 'fatigue' (data subjects are not always aware how their data are being used because one cannot follow up on every single instance of potential abuse) (Martin 2019), and also because understanding data use decrees or polices often requires that the data subject possess a high level of technical or legal understanding of how data are legally permitted to be used. Additionally, the more 'click through' agreements we receive, the less likely we are to read them (Lundgren 2020). Again, these reasons help to explain why there are such problems, but they by no means excuse the behaviour of companies, whose actions help to create them.

3 Big data and informed consent

In this section, we will discuss three of the main types of problem that can impede informed consent with respect to Big data. These are: the transparency (or explanation) problem, the re-repurposed data problem, and the meaningful alternatives problem.

3.1 The transparency problem

The transparency (or explanation) problem is an epistemic problem that can arise from several sources. First, it can arise from companies' reluctance to reveal their own internal workings—e.g., they may not want to reveal

[5] See Australian Competition & Consumer Commission (2019), for the details of this report.

trade secrets (Pasquale 2015). The problem can also arise unintentionally, such as when 'black-box' algorithms, often used in deep learning, are implemented (Innerarity 2021). While the inputs and outputs of certain algorithms may be viewable, and in principle explainable, the internal working may not be. A machine learning algorithm may be so complex that not even the creators understand how it works (Burrell 2016; Mittelstadt et al. 2016). If users have a statutory 'right to an explanation' about how their data are being used [as they do, for instance, under the EU's GDPR], and this explanation is a necessary part of their ability to make informed decisions, then the problem of procuring meaningful 'informed consent' becomes extremely challenging in these contexts.[6] This is not to imply that informed consent cannot be sought until a full understanding of an algorithm's inner workings is received. This would be to set the bar unnecessarily high. For example, providing a user with a specific explanation of how a machine learning algorithm works, such as a detailed account of how it uses backpropagation, would not be appropriate in many contexts. More transparency is not necessarily desirable as it may overwhelm the user (Tsamados et al. 2021). What is of central importance is the context, and the potential harm that a decision could cause, that should help decide how much transparency is appropriate (Robbins 2019). To compare this to the medical context, a person does not need to understand how the heart functions before an operation to meaningfully consent to an operation: in such a context they would typically be more concerned about what the health risks of the operation are, or how their life will be affected by the operation. Opacity is typically a problem when the user is unaware of how their data are being used.

O'Neil (2016, p. 4–6), for instance, discusses a case where a fifth-grade schoolteacher received a poor score from an algorithm which assessed teacher performances. And as a result of the low score she received, she lost her job. The algorithm used data about her student grades, amongst other data, to decide that she was performing poorly. Initially the teacher did not know why she received the score, as she thought of herself as a good teacher. She, and others, lacked an understanding of how the algorithm came to such a conclusion. She later learned that the algorithm drew from a limited data set and placed value on results which were seen by some as a non-optimal way of measuring a teacher's worth. What the teacher initially lacked in this situation was an explanation of how the algorithm reached the conclusion it did.

This specific problem about making AI use more transparent has been the focus of much recent attention. Explainable AI (XAI), for instance, is an emerging field of research (Schmelzer 2019; Shaban-Nejad et al. 2021), which seeks to provide more meaningful explanations about how AI algorithms work. XAI is important, like informed consent, because it provides data subjects with a sense of awareness about how their data are being used. Providing a meaningful explanation of how one's data are being used can be humanising because it gives data subjects a greater sense of control, with respect to their information (Colaner 2021). XAI promises data users, 'fairness, trust, and governability' (Colaner 2021).

Another promising idea that may help with the transparency problem comes from research into synthetic data, which refers to data sets that have been generated by AI algorithms, after being trained by real world data. Synthetic data can be useful, especially for research purposes. For example, realistic sets of medical records consisting of synthetic data would not feature any individual's personal information, yet it would still be useful for carrying out medical research (Kearns and Roth 2020, p. 135).

Some of these ideas are promising, and if implemented appropriately, will help with the transparency problem. But as long as we are required to use real world data (whether this is for machine learning training or for actual use), these problems will persist. Even when companies make attempts to be transparent, communication problems persist. Consent forms may be long; and they may be written in a technical language that is hard to parse, making them inaccessible. Companies are not often pressured to change current practices, as so few people read these documents (cf. Cohen 2019, p. 162; Zuboff 2019). According to a recent Consumer Policy Research Centre Report, for example, most Australians (94%) state that they do not read all of the privacy policies that apply to them (Kemp 2018). And in a recent survey of European internet users, it was found that under a fifth (18%) fully read privacy statements (Eurobarometer 2015, p. 84). Hence, so few people know what they are agreeing to, or what they should be objecting to, with respect to these documents.

Such low rates should not be attributed to peoples' lack of care about privacy or their alleged irrationality. A more plausible explanation is that the policies are very long and often written in a highly technical language which is incomprehensible to the average person. The most popular reason (67%) that Europeans in a recent survey gave for why they

[6] This is most evident in Article 15 of the GDPR—'Right of access by the data subject'—where it is stated that data subjects have the right to (i) obtain information about what their data will be used for; (ii) know which parties have access to their data; and (iii) know the length of time their data will be stored for (GDPR 2018). For a recent critique of the GDPR's capacity to ensure that a right to an explanation is secured, see Wachter et al. (2017). They argue that the GDPR does not, in its current form, give data subjects a right to an explanation, due to the fact that the document's language is ambiguous in parts. In their article, they make recommendations about how this issue can be resolved.

do not fully read privacy statements is that they are too long (Eurobarometer 2015, p. 87). The second most popular reason (38%) was that privacy statements are too unclear and hard to read (Eurobarometer 2015, p. 87). These individuals' beliefs are well corroborated. In 2008, two Carnegie Mellon professors (McDonald and Cranor 2008) estimated that the average Internet user, if they read all the privacy policies that they encountered in a year, would require on average 76 work days to get through them. No doubt it would take longer today (cf. Zuboff 2019, p. 50). Further, there is rarely a meaningful alternative to a privacy policy as there is no straightforward opportunity to negotiate the stated terms and conditions.

3.2 The re-purposed data problem

One of the features of Big data analyses is that new AI algorithms can be applied to existing data sets to yield new information. While a human would find it near-impossible to search through tens of thousands of medical records, to discover novel patterns and insights, an AI algorithm can be designed to perform such a task very quickly. As beneficial as this can be in many contexts (e.g., assisting patient care or preventing disease (Arnold 2021)), informed consent may need to be secured again, if the original consent is no longer applicable. For example, someone who consents to sharing their postal code may wish to withdraw consent when they learn such data can be used to determine insurance premiums (see Floridi 2019, p. 110). And this is also true in biomedical contexts, where medical information or tissue samples are often stored and then requested for further research to which the original Participant Information Sheet and Consent Form (PICF) did not refer.

In other cases, it may be much harder for users (and even companies) to predict how certain data could be repurposed in the future. One of the features of AI and Big data, after all, is that surprising or unexpected information, or correlations, can sometimes be revealed from existing data sets (Mittelstadt and Floridi 2012, p. 312). This makes it hard for data subjects to assess the risk of consenting to sharing certain data because it is hard to predict how their data could be used in the future (Cohen and Mello 2019). For example, a person may initially consent to having her Facebook likes publicly viewable by her friends and Facebook. But that person may not consent to having a third party collect and analyse those likes so that they can be used to predict her sexual orientation, religious or political views, intelligence, or happiness (Kosinski et al. 2013), and then use those predictions to target ads towards her. For many users, it is hard to even imagine how such information about a person could be obtained by only looking at Facebook likes.

Cases involving data pings illustrate this problem in an alarming way. According to a recent piece in the New York Times, reporters obtained a leaked file containing over 50 Billon location pings, from over 12 million phones in the United States dating from 2016–2017 (Thompson and Warzel 2019). These data (shown as dots on maps) represent the locations of people's movements in major cities in the US, such as New York and Washington, D.C. For companies (e.g., data brokers) wanting to track where certain individuals go; or which individuals visit a certain place (e.g., the White House or a certain celebrity's house), these data are invaluable.

For most users, this kind of data sharing would likely go beyond what they had in mind, when they consented to having their smart phone access their location data (data about the current phone position in space). For the average user, consenting to giving away their location data is merely the way to see what the weather is like in their current location, or where the closest hospital to them is. Most would not be aware that such data are being shared with data brokers and perhaps fewer would be informed about what data brokers do. Since the location data have been repurposed, the original consent is no longer applicable. Given peoples' concerns about potential threats to their reputation, or opportunity for manipulation, such revelations have caused anxiety for some (cf. Cohen 2019, p. 76).

The general problem of re-purposed data, then, is that data users (e.g., companies and institutions) have not always limited their use of personal data to the purpose for which the subjects' original consent was applicable. This is morally problematic because it ignores the preferences, and potentially the wellbeing, of these subjects (ACCC 2019).[7] The reason informed consent ought to be sought in the first place, we have suggested, is so that subjects have the information required to decide whether consenting to something is in their best interests; or whether consenting to that thing could potentially cause them harm. Such

[7] In Australia, at present, this is primarily a moral problem, as the protections currently afforded by the *Privacy Act 1988* (Cth) are minimal, while in Europe it is also a legal problem (see GDPR 2018). As mentioned above, the recent 'Digital Platforms Inquiry' conducted by the ACCC (Australian Competition and Consumer Commission.

2019) found current Australian legislative and regulatory protections wanting on a number of levels and made recommendations to depart from exclusive reliance on principle-based regulation, and add more rule-based protective requirements, some of which is inspired by the GDPR (2018).

information allows them to perform a kind of risk assessment.[8] It does not follow, of course, that failing to secure consent will always cause harm, as some users may not care what companies or institutions do with their data. The problem is, however, that many people are concerned, or anxious, about their data being used in ways that go beyond what they originally thought would occur when they first consented—if, indeed, their consent ever really amounted to informed consent.

3.3 The meaningful alternatives problem

This issue arises when users are not given alternative choices in cases where they do not wish to consent. If the only way for S to gain access to *P* is to accept a set of conditions *C*, which S is hesitant to consent to, then S's choice is compromised. Currently, users who do not wish to comply with the terms and conditions lack the power to renegotiate. In the medical context, a practitioner may offer their patient reasonable alternatives to a certain procedure, if a patient does not wish to go ahead with a suggested procedure. And when recruiting research participants, medical researchers must ensure that refusal to be involved in a research project will not prevent a patient from receiving the medical treatment that they would standardly receive. Cutting edge medical research has also managed to address both the *re-purposed data problem* and the *meaningful alternatives* problem through the mechanism of 'dynamic consent', which allows participants to constantly recalibrate their initial decision and provide (or refuse) fresh consent to new and emerging uses of their data beyond what they initially consented to (Kaye et al. 2015).

The situation is quite different with respect to online privacy policies, and the offer of reasonable alternatives is far less frequent. For example, if users of a particular app or online service do not wish to accept a set of terms and conditions, in many cases their only choice is to reject them and, in doing so, give up their use of the app or online service. It could be argued that this is coercive, at least to some extent. If a user of an app or online service wants (or needs) to use the service provided enough, and there is no way to do so without subjecting themselves to practices they are uncomfortable with, then they are simply far more likely to accept the terms and conditions. Social media provides a good example of this: many young adults report that they have experienced the fear of missing out, with respect to their social media engagement (Przybylski et al. 2013). The price to pay for opting out of a service, where one is not comfortable with the terms and conditions associated with it, may be social isolation or reduced communication with peers.

4 Big data and the future of informed consent

We have discussed three types of problems that have arisen in the recent use of Big data (though there are others, of course). Given the problems with lengthy consent agreements, and the fact that few people actually read them, we do not think that adding more detailed clauses to already burdensome documents will be enough to overcome these problems and ensure informed consent is secured. In this final section of the paper, we suggest some alternative solutions. While we provide reasons to justify these solutions, it is important to stress that they will need to be empirically tested. The argument we make is therefore normative, not empirical.

4.1 'Soft governance' of personal data for commercial and administrative use

First, we propose that some uses of personal data for commercial and administrative objectives could be subject to a 'soft governance' ethical regulation, akin to the way that any research involving human participants (spanning the collection of human medical data and tissue use, to participation in interviews and focus groups) are regulated in many countries through the Human Research Ethics Committees (HRECs). HRECs review all research proposals that involve human beings (from hard to social science projects), to ensure that they meet accepted ethical standards and guidelines. Using Australia as our example, there are over 200 of these HRECs in operation. Importantly, HRECs perform a key governing role in cases where data can be released for medical research, for purposes beyond those initially consented to by participants, without the need for fresh consent—which can happen where this would be impractical *and* the risk to participants whose data are being re-purposed is minimal (Flack et al. 2019). In our view, at least in the context of

[8] A risk assessment of this kind need not be an elaborate one of course. It is the kind we perform in our everyday lives. For example, when one gets into a car, one (should) know that there is a small chance that they could get into a crash and get seriously injured. Most of us continue to travel by car, however, because it is convenient, and the probability of crashing is typically low. In circumstances where new information is presented to us, however, we may need to revise such probabilities. For example, if one learns that the driver of a car one is about to get into is inebriated, or does not possess a driver's licence, one would typically not consent to allowing them to drive them home. It would be simply too risky for most people—given that the high probability of getting injured significantly outweighs the gains (in this case convenience). Analogously for repurposed data. If repurposing data introduces new risks for data subjects, it is not fair to subject them to such risks unless they first have knowledge of them and have agreed to proceed anyway.

the institutional use of personal data, similar committees could be set up which would be responsible for evaluating policy statements by tech companies, and institutions, to ensure that their proposed policies meet ethical standards and guidelines. Take the example of public universities in Australia. Currently, while every research project involving participants undergoes a stringent review by the competent HREC, collection and use of student and staff data for 'administrative purposes' is entirely deregulated—save for the minimal protections of the *Privacy Act* 1988 (Cth) where applicable. Australia also has Population & Health Services Research Ethics Committees (PHSRECs) who 'grant ethical approval for research proposals' and 'reject research proposals on ethical grounds' (Cancer Institute NSW 2021). One of their jobs is to review applications for accessing data.

An HREC-like soft-governance model, drawing on such existing ethics approval committees, could help prevent the kinds of personal data controversies that have recently become prevalent. Consider the case involving telecommunications company Verizon, who were recently fined US 1.35 million by the Federal Communications Commission (FCC) for using so-called 'supercookies' to track their users' browser data without their consent (Peterson 2016). Supercookies are harder for users to delete (compared to standard cookies) because they do not reside on the users' computers or devices. So, simply clearing one's browser history will not get rid of them. This makes them a valuable tool for capturing browsing data, which can be used by companies for targeted advertising.

Following the FCC fine, Verizon was required to receive consent from users before they could track users' browsing history with supercookies, and also inform its customers about how the system of targeting advertising worked. This is a good result in our view, but it is one that ought to have been implemented at the outset. The use of supercookies to gather browser data without user consent would have been unlikely to pass through a HREC-like review board, on the grounds that users were not informed about the consequences of using the services. Although it is true that fines were issued to Verizon, these can often be written off as business expenses, and may not be enough to deter other companies from finding alternative ways to get such user data.

Consider another example, from Australia, whereby legal action was taken by the ACCC against Google. The ACCC accused Google of misleading its users about the personal data that they were sharing (see Kemp 2019). The ACCC accused Google of violating Australian Consumer Law in two ways: the first was that Google did not tell users that two settings needed to be switched off in order for user data not to be collected and shared; and the second was that Google did not inform its users that location data could be used for other purposes, that went beyond Google's services. It is very unlikely that a HREC-like committee who reviewed these policies would have approved of them on the grounds that Google failed to satisfactorily inform its users about what they were doing with their data, in addition to the risks associated with sharing data. It would be far preferable, in our view, that non-consented personal data are not released in the first place.

Further, there are new advances in algorithm design that could be recommended by HREC-like committees in certain contexts to uphold privacy. Kearns and Roth (2020), for instance, discuss the use of so called 'ethical algorithms.' One example they describe draws upon the idea of randomized polling protocols. This is a technique that adds a percentage of randomness to data sets. Even though this means that some of the data in a set will not be accurate, the data set as a whole will be, if the percentage of randomness is known. This can help protect privacy because if raw data are leaked, information about individuals cannot be easily identified (Kearns and Roth 2020, p. 45). There is always the chance such a record could have been random: not even the data users would be able to tell which data were real and which data were randomised. A HREC-like ethical review process could suggest the use of such algorithms in appropriate contexts.

There is a significant objection to our proposal that is important for us to address here—namely: what would incline a global company like Verizon, Google, or Facebook to subject their practices to this kind of process, even if public institutions like universities could be persuaded to comply? If recent behaviours are a reliable guide to what to expect in the future, would it not simply be too easy for companies or institutions to ignore or dismiss the kinds of ethical recommendations that we have in mind here?

To address this objection, it will be useful to consider what motivates researchers (e.g., psychological or biomedical) to get ethics approval for their experiments or research on human subjects. One prominent reason is that without HREC approval, many academic journals will not publish researchers' work. The consequences of this are quite significant and obvious. If a work is not published in a journal then there is a greater chance that it will fail to reach a wide audience, go unnoticed, or not be taken seriously by other researchers. Academic journals have typically justified such a stance by pointing out that ethics approval preserves the integrity of the research and takes into account the interests of the human beings who are the subjects of such research. Newson and Lipworth (2015), for example, give four reasons why health promotion research (to pick just one type of research) should not be published without ethics approval. They claim that ethics approval helps to (i) ensure there is a legitimacy to the research, which helps to establish public trust; (ii) it ensures that the study is well planned; (iii) it identifies areas in the study which may be risky; and (iv) it gives due respect to humans (2015, pp. 173–174).

We think that the reasons Newson and Lipworth provide for the justification of ethics approvals, with respect to research involving human subjects, can also be used to justify our suggested ethics approval component for certain Big data practices. Although researchers/experimenters and Big data users differ in many important respects, both are dependent on human subjects for their work. As such, both have the potential for improving human lives; and conversely both have the potential to cause great harm. In both contexts, for example, it is possible that human subjects can be used merely as a means to achieve certain ends, such as making profit. Much has been written in moral philosophy about the problems with treating people as mere means; perhaps the most well-known work comes from the eighteenth century German philosopher Immanuel Kant. In Kant's Humanity Formulation of his famous Categorical Imperative, he stresses the importance of not using anyone as a mere means to an end, because doing so fails to recognise them as an agent. Treating someone merely as a means treats them like an object. This is morally wrong, according to Kant, because it does not treat subjects as what they actually are—namely, free agents (1993, pp. 37–38). So, failing to get consent from a subject, or taking advantage of a subject for commercial gain, can in certain situations downplay that subject's own agency, and can thus dehumanise them.

Returning to the problem of motivation, what would motivate large companies like Google, Verizon, or Facebook to subject their practices (e.g., their policies and terms of conditions) to ethical review? What inconveniences would they seek to avoid by complying with ethical (HREC-like) approval? Our response is that at the start, they may be hard to convince. But we have to keep in the forefront of our minds that this is just the start. Big data technology and its applications are still relatively new and as a global culture we are only just beginning to comprehend the ethical issues that are posed by these advances. We cannot expect to have comprehensive ethical regulation of practices that have just arisen—however, that should not incline us to just throw our hands up and let anything go. If we are looking to soft ethical regulation practices in medical research (and human research, more generally) for useful ways of thinking about the ethical regulation challenges posed by Big data, it is worth remembering that some of the earliest uses of human bodies in medical research was utterly unregulated. For instance, in the 17th and 18th Century in Scotland and the UK medical researchers were known to source their cadavers from grave robbers.

Of course, over time the use of human bodies and human participants in medical research has become more and more well regulated. In some respects, the situation of Big data at this time is akin to the very earliest use of human bodies in medical research. We need to recognise that just because the current situation feels a bit like it is beyond proper regulation, given time and continued effort, it is possible to bring even the most complicated situations under good ethical regulation. However, even with all that said, the HREC-like solution to the post-facto regulation of personal data usage, that we have put forward here, might be in the first instance limited to the particular context of public institutions like universities, rather than attempting to bring large corporations under a voluntary regulatory practice that at this stage they have no cultural expectation or economic motive to adhere to.

Universities may be more likely to react positively to the idea of subjecting their own Big data policies and practices to ethical approval because they already exist in an environment where ethics approvals are taken seriously. The data that universities collect, further, are also of great importance, as they collect personal data from their students and staff. Another reason that universities may be more likely to comply with such approvals is that they place such a high premium on their own reputations (academically and ethically). Universities which are seen as failing to comply with ethical review of their Big data practices may be seen as lacking integrity, which may affect the way that the public views them. As the public come to expect such kinds of ethical approvals from their educational initiations, they then may begin to demand them from larger companies, who may be forced to adapt for economic reasons.

The Big data problem does not just apply to the Big four companies (Amazon, Apple, Facebook, and Google). Universities, Governments, and many small and large business also collect personal information. So even though big companies have a lot of influence, we should keep in mind how widespread personal information use is, and how important it is for all data subjects and users to be concerned with its proper use. Accordingly, there is a great opportunity for such organisations to gain reputational benefits from ethical data use practices. As has been shown in the past few years, customer values play an important role in purchasing decisions. More and more customers are deciding to purchase products that align with 'personal values' (Rosmarin 2020). So, companies whose data policies have been subject to ethics approval, may benefit from a competitive advantage over others that do not.[9]

[9] Furthermore, ethical approval of data-use policies by an independent HREC-like body could become the new frontier of the fast-growing field of 'corporate ethics', potentially leading to legislative reform. While it is worth signalling this aspect of the issue, it is beyond the scope of this paper to provide an in-depth analysis of corporate ethics in this space.

If we think about recent changes in attitudes towards informed consent, in the biomedical context, we see that critics were not only able to point out the ethical shortcomings of the practices of their day but were also able to propose and implement alternatives. In doing so, they were able to show that the practices of their day were neither necessary nor unavoidable. Robert Baker (2019), for example, describes the shift that occurred in the 1960s/1970s between the paternalistic 'doctor knows best' paradigm, which treated informed consent as a mechanism to protect researchers; and the 'respect for the autonomy of the patient' paradigm, which takes into account the patient's agency and rights (see Baker 2019, p. 175).[10] What is noteworthy about this change, as Baker notes, is that it not only came about as a result of the criticisms that individual patients and prominent bioethicists of the day made. Alternatives to the doctor knows best paradigm were suggested and implemented. Leading Bioethicists authored influential documents and textbooks such as the *Belmont Report* (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research 1978), which identified and made the case for a greater focus on respect for persons, beneficence and justice; and the *Principles of Biomedical Ethics* by Beauchamp and Childress (1979) which focused on establishing the importance of four ethical principles: autonomy, non-maleficence, beneficence and justice (cf. Page 2012). It was not enough simply to criticize the doctor knows best paradigm; alternatives were proposed, and medical practitioners were convinced to abandon the old paradigm. This is important, as Baker has also observed (2019, p. 173), because people tend to prefer problematic established paradigms to none at all.

Analogously with Big data use. While it is important to criticise current practises, it is just as important to propose and implement alternatives. It is common for people who learn about a certain company's abuse of their personal data to express anguish, and at the same time lament that this is just the way things are. If public institutions like universities would subject their practices to such a soft governance approach we would, as a culture, have glimpse at the new paradigm—we would see that the existing paradigm is not necessary, that alternatives are actually realisable. A successful application in the university setting could then pave the way for businesses, but also governments, to subject their own practices and proposals to ethical review. The idea is applicable to a recent case in Australia: the recent introduction of the 'Data Availability and Transparency Bill 2020'. This bill seeks to simplify data use, for example, by requiring a person only to upload their personal information once, in order for multiple agencies to access it. While the government officials who put forward the bill claim that citizens will benefit from its added simplicity, there is notable lack of focus on informed consent in the bill. This is concerning because the proposal means that it will be easier for organisations, researchers, and companies to access peoples' personal information (Taylor 2021). And since some of the people using government services, and providing their personal information, are vulnerable people, there is the potential for harm to be caused (Taylor 2021). A HREC-like approval mechanism would seek to highlight these issues, and bring to light the ethical shortcomings of the bill. In a recent analysis piece from a set of scholars based at UNSW Sydney, for example, Bennett Moses et al. (2021) point out that the bill does not provide adequate definitions or examples about when it is reasonable to secure consent. The authors raise the concern that what is considered reasonable to one person may not to another. This ambiguity means that there is risk that the kinds of data misuses we have discussed above could be facilitated. A HREC-like review process, on the other hand, would ideally identify these ethical problems before the bill were passed, and make recommendations about how to resolve them.

The importance of engendering public trust in the security of their private information and hence, in the process by which their collection and storage of their data was designed and implemented, was made especially salient with respect to the uptake of COVID-19 contact tracing apps. Convincing members of the public to download and use a COVID-19 contact tracing app was a pivotal element of a successful roadmap out of lockdown (or to avoid lockdown) for many nations. Assurances that the data that people were providing on their whereabouts would only be used for contact tracing purposes and would be deleted immediately at the end of its useful life played a key role in ensuring a broad uptake of the technology. While not necessary in the COVID-19 case, perhaps because of the coercive power of the promise of a return to normal, an independent HREC-like body established to assess such claims and provide assurance that they are being made good, could act to establish and maintain public confidence and trust in a range of justifiable and important data uses.

A controversy has arisen very recently in the state of Western Australia (WA). The state promotes the use of an app, called SafeWA, for WA residents to check in any public location they visit to facilitate contact tracing (the alternative being manual signing in). In June 2021, it became apparent that the WA Police had accessed SafeWA app data for the purpose of facilitating a murder investigation (Manfield 2021). This happened despite the state government's assurances that the data would be used for COVID-19-related contact tracing exclusively (Manfield 2021). While the use of this type of data in the context of criminal investigation does not constitute a Big data-type of repurposing (and

[10] The idea of a paradigm shift is from Thomas Kuhn (1962), who applies it to scientific revolutions.

while it may very well be proportionate and legitimate), it raises an important issue of trust—if these data are used in criminal investigations, then in what other contexts are they used? And who else has access to the data? We do not contend that adopting an HREC-like model could, in and of itself, solve the complexity of this issue. It would however provide a baseline of accountability, whereby the public authorities' potential use of certain data would be subject to an ethical review and approved for a limited range of uses *ex ante*. The case of the SafeWA app is prompting authorities to look into legislative reform. While this may certainly be needed, a soft-governance model would have the benefit of not requiring constant realigning of legislation, while providing broad ranging guarantees across the (vast) spectrum of potential uses and abuses of data.

4.2 Facilitating 'information transfer' through imaging

Even if such soft governance approaches are successful, there still remains the issue of informed consent. In the medical context, recall, subjects still need to be informed about procedures—even if they are deemed legitimate and fair by a HREC. It is still important, thus, that an 'information transfer' occurs, between users and companies, just like in the medical context. And it is still important that individuals feel like they have a real choice and that they are respected enough to be consulted in the first place, as discussed in Sect. 1. There is symbolic value in seeking consent (Walker 2020). After all, some people may wish to give up personal information. What is at issue is whether users understand what they are giving up. After all, a decision still has to be made by a person. This is the case even if the measures have been sought to preserve the subject's own interests, which, as discussed, should not simply take the form of additional consent decrees posted in obscure, hard-to-find sections of a company's website.

An alternative approach to the standard consent forms, and privacy policies, could make use of some of the latest research focussed on the usability of pictorial legal contracts (see Keating and Andersen 2016; Andersen (2018); Brunschwig 2019; McGuire and Andersen 2019). These are contracts that are distinguished from the standard read and sign forms, and are presented in coloured pictures. Consider, for example, the 'Comic Contracts' developed by a company called Creative Contracts, who have sought to address the problem that text agreements cannot be understood by poor, vulnerable and illiterate people (Brunschwig 2019). For example, farm workers from parts of the world where literacy rates are low, may not be able to read standard legal text agreements. Pictorial contracts are beneficial to these farmers not only because they facilitate a proper information transfer, but also because they motivate individuals to comprehend the contract.

Given that only a small percentage of users whose personal data are being sought will be able to fully understand personal data agreements, we think that these 'comic contracts' could prove useful. While most users who give express consent may possess a basic ability to read, many will be in 'vulnerable' positions because they will not be able to decipher the technical language and implications of standard contracts. Given an essential part of informed consent involves a transfer of information, comic contracts may help due to their ability to simplify complex information. There is, after all, evidence to suggest that by simplifying medical consent decrees, a greater comprehension and retention of information ensues (Dresden and Levitt 2001). So, there is reason to believe that the same could be true in the Big data space. Further, the visual aspect of comic contracts may also help with basic comprehension, since there is evidence to suggest that visual explanations can improve learning (Bobek and Tversky 2016). We do not claim that pictorial contracts are a 'magic bullet' that will solve all of the problems outlined here. Neither do we recommend that pictorial contracts should supersede written ones (cf. French 2019). Rather, we think that they may help with the facilitation of information transfer—a key component of informed consent.

One objection to this approach is that it will be ineffective, since there is such a low compliance rate with existing consent decrees (as stated in Sect. 3). Why think that people will engage with consent decrees just because pictures are included? It is difficult to predict future behaviour, of course, but if polls are to be believed there is some evidence to suggest that certain people would be more inclined to engage with consent forms if they were simplified. Recall that the second most popular reason (38%) that participants in a recent European survey gave for why they do not read privacy statements is that they are too unclear and hard to read (Eurobarometer 2015, p. 87). If enough consent decrees could be produced that convey complex information in an understandable way, and new standards could be set, then levels of engagement could very well go up. Andersen (2018), for example, notes of a non-disclosure agreement—one that was legally binding—that was depicted in just three pictures. If applicable on a large scale, this would provide more opportunities than currently exist for people to become informed about how their data are being used. Privacy policies may no longer be seen as obstacles to get past as fast one can possibly can. They may actually inform data subjects.

Our suggestions are by no means the only ways to counter the problems discussed in this paper. Another interesting idea is discussed Lundgren (2020), who notes that consent forms could be bypassed, in certain contexts, with the use of advanced web browsers. He explains that web browsers could be designed to implement 'pre-set responses' that save users' privacy preferences. Such advances could save a data subject time, and the difficulty of comprehending consent forms, since they would not be required to read all the consent forms they encounter. All they would need to do is configure their basic privacy requirements. The browser would then be able to accept or reject the use of any pages or applications on the basis of these preferences. Given the difficulties with current consent form practices that were outlined above, such software advances could be advantageous. Users would, ideally, still need to understand how best to set up their 'pre-set responses', but such an approach may offer improvements upon the current paradigm. Pictorial representations of the kind we have suggested here could even be used to help users configure their 'pre-set responses'.

5 Conclusion

The Big data practices that companies and institutions partake in today are not unavoidable or inevitable. Alternative ways of securing informed consent exist; it is worth exploring the ethical implications of applying them. The introduction of a soft governance approach, based upon the HREC model, could prove useful in ensuring that individuals are not taken advantage of. Since this approach only addresses one half of the problem, we also focused on improving ways in which consent is actually sought. We considered alternatives to the standard consent forms, and privacy policies—namely, the pictorial contracts. This makes the approach we advanced a two-pronged one—one that focuses not only protecting the rights of individuals, but one that respects peoples' own autonomy and decision-making abilities.

These proposals are by no means all that is required to solve the problems of Big data use that we have identified here. And it is important to note that they need to be empirically tested and further developed. As we have suggested throughout, we are still in the early stages of this ethical crisis. Big data technologies and their applications are new and will continue to evolve and bring with them new ethical problems (consider recent developments in AI emotion capturing technology, for example). As a global culture, we have been in this position before, however. Our current day views of informed consent have had a long history and did not spring up overnight. It is worth attempting to understand that history, so that we might minimize the harm that current Big data use is causing. Our aim in the paper has been to make a contribution to that effort.

Funding Not applicable.

Availability of data and material Not applicable.

Code availability Not applicable.

Declarations

Conflict of interest Not applicable.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

References

Andersen CB (2018) Comic contracts and other ways to make the law understandable. The Conversation. Retrieved from: https://theconversation.com/comic-contracts-and-other-ways-to-make-the-law-understandable-90313. Accessed 23 Aug 2021

Arnold MH (2021) Teasing out artificial intelligence in medicine: an ethical critique of artificial intelligence and machine learning in medicine. J Bioeth Inq 18:121–139. https://doi.org/10.1007/s11673-020-10080-1

Australian Competition and Consumer Commission (2019) 'Digital platforms inquiry—final report.' Retrieved from: https://www.accc.gov.au/system/files/Digital%20platforms%20inquiry%20-%20final%20report.pdf. Accessed 23 Aug 2021

Baker R (2019) The structure of moral revolutions: studies of changes in the morality of abortion, death, and the bioethics revolution. MIT Press, Cambridge

Beauchamp TL, Childress JF (1979) Principles of Biomedical Ethics. Oxford University Press, New York.

Beauchamp TL (2011) Informed consent: its history, meaning, and present challenges. Camb Q Healthc Ethics 20:515–523. https://doi.org/10.1017/S0963180111000259

Bennett Moses L, Johns FE, Land LPW, Vaile D, Zalnieriute M, Yastreboff M, Zhao S, Nicholson K, de Sousa T, Whitty M (2021) Inquiry into the data availability and transparency bill 2020 and the data availability and transparency (consequential amendments) bill 2020. UNSW law research paper no. 21–37, Available at SSRN: https://ssrn.com/abstract=3807026 or https://doi.org/10.2139/ssrn.3807026. Accessed 23 Aug 2021

Bobek E, Tversky B (2016) Creating visual explanations improves learning. CRPI 1:27. https://doi.org/10.1186/s41235-016-0031-6

Bronskill J (2020) Malls gathered facial images of five million shoppers without consent: watchdogs. National post. Retrieved from: https://nationalpost.com/pmn/news-pmn/canada-news-pmn/malls-gathered-facial-images-of-five-million-shoppers-without-consent-watchdogs. Accessed 23 Aug 2021

Brunschwig CR (2019) Contract comics and the visualization, audio-visualization, and multisensorization of law. Univ W Aust Law Rev 46 (2):191–217. https://www.law.uwa.edu.au/data/assets/pdf_file/0004/3459415/Brunschwig-FInal.pdf. Accessed 23 Aug 2021

Burrell J (2016) How the machine 'thinks:' understanding opacity in machine learning algorithms. Big Data Soc 3(1):1–12. https://doi.org/10.1177/2053951715622512

Cancer Institute NSW (2021) NSW population & health services research ethics committee. Retrieved from: https://www.cancer.nsw.gov.au/research-and-data/nsw-population-health-services-research-ethics-com. Accessed 23 Aug 2021

Cohen JE (2019) Between truth and power: the legal constructions of informational capitalism. Oxford University Press, Oxford

Cohen IG, Mello MM (2019) Big data, big tech, and protecting patient privacy. JAMA 322(12):1141–1142. https://doi.org/10.1001/jama.2019.11365

Colaner N (2021) Is explainable artificial intelligence intrinsically valuable? AI Soc. https://doi.org/10.1007/s00146-021-01184-2

Dresden GM, Levitt MA (2001) Modifying a standard industry clinical trial consent form improves patient information retention as part of the informed consent process. Acad Emerg Med 8(3):246–252. https://doi.org/10.1111/j.1553-2712.2001.tb01300.x

Duffy C (2021) Facebook approves alcohol, vaping, gambling and dating ads targeting teens, lobby group finds. ABC News. Retrieved from: https://www.abc.net.au/news/2021-04-28/facebook-instagram-teenager-tageted-advertising-alcohol-vaping/100097590. Accessed 23 Aug 2021

Eurobarometer (2015) Data protection. Special Eurobarometer 431. https://ec.europa.eu/commfrontoffice/publicopinion/archives/ebs/ebs_431_en.pdf. Accessed 23 Aug 2021

Faden RR, Beauchamp TL (1986) A history of informed consent. Oxford University Press, New York

Flack F, Adams C, Allen J (2019) authorising the release of data without consent for health research: the role of data custodians and HRECs in Australia. J Law Med 26(3):655–680

Floridi L (2012) Big data and their epistemological challenge. Philos Technol 25:435–437

Floridi L (2019) The logic of information: a theory of philosophy as conceptual design. Oxford University Press, Oxford

French R (2019) Closing address, comic book contracts conference. Univ West Aust Law Rev 46(2):268–271. https://www.law.uwa.edu.au/data/assets/pdf_file/0011/3442655/8.-French-Closing-Address.pdf. Accessed 23 Aug 2021

GDPR (2018) General Data Protection Regulation. https://gdpr-info.eu/n. Accessed 23 Aug 2021

Innerarity D (2021) Making the black box society transparent. AI Soc. https://doi.org/10.1007/s00146-020-01130-8

Isaac M, Singer N (2019) Facebook agrees to extensive new oversight as part of $5 billion settlement. The New York Times. Retrieved from https://www.nytimes.com/2019/07/24/technology/ftc-facebook-privacy-data.html?mBurodule=inline. Accessed 23 Aug 2021

Kadam RA (2017) Informed consent process: a step further towards making it meaningful! Perspect Clin Res 8(3):107–112. https://doi.org/10.4103/picr.PICR_147_16

Kant I (1993) Groundwork for the metaphysics of morals, James W Ellington (trans.). Hackett Publishing Company, Indianapolis

Kaye J, Whitley E, Lund D, Morrison M, Teare H, Melham K (2015) Dynamic consent: a patient interface for twenty-first century research networks. Eur J Hum Genet 23:141–146. https://doi.org/10.1038/ejhg.2014.71

Kearns M, Roth A (2020) The ethical algorithm. Oxford University Press, Oxford

Keating A, Andersen CB (2016) A graphic contract: taking visualisation in contracting a step further. J Strateg Contract Negot 2(1–2):10–18. https://doi.org/10.1177/2055563616672375

Kemp K (2018) 94% of Australians do not read all privacy policies that apply to them—and that's rational behaviour. The Conversation. Retrieved from https://theconversation.com/94-of-australians-do-not-read-all-privacy-policies-that-apply-to-them-and-thats-rational-behaviour-96353. Accessed 23 Aug 2021

Kemp K (2019) The ACCC is suing Google over tracking users. Here's why it matters. The Conversation. Retrieved from: (https://theconversation.com/the-accc-is-suing-google-over-tracking-users-heres-why-it-matters-126020?utm_medium=email). Accessed 23 Aug 2021

Kosinski M, Stillwell D, Graepel T (2013) Digital records of behavior expose personal traits. Proc Natl Acad Sci USA 110(15):5802–5805. https://doi.org/10.1073/pnas.1218772110

Kuhn TS (1962) The structure of scientific revolutions. University of Chicago Press, Chicago

Levin M, Wanderer JP, Ehrenfeld JM (2015) Data, big data, and metadata in anesthesiology. Anesth Analg 121(6):1661–1667. https://doi.org/10.1213/ANE.0000000000000716

Lundgren B (2020) How software developers can fix part of GDPR's problem of click-through consents. AI Soc 35:759–760. https://doi.org/10.1007/s00146-020-00970-8

Macnish K, Gauttier S (2020) A pre-occupation with possession: the (non-) ownership of personal data. In: Macnish K, Galliott J (eds) Big data and democracy. Edinburgh University Press, Edinburgh, pp 42–56

Manfield E (2021) Police access SafeWA app data for murder investigation, prompting urgent law change. ABC News. Retrieved from: https://www.abc.net.au/news/2021-06-15/safewa-app-sparks-urgent-law-change-after-police-access-data/100201340. Accessed 23 Aug 2021

Manson NC, O'Neill O (2007) Rethinking informed consent in bioethics. Cambridge University Press, Cambridge

Martin K (2019) Ethical implications and accountability of algorithms. J Bus Ethics 160:835–850

McDonald AM, Cranor LF (2008) The cost of reading privacy policies. I/S J Law Pol Inf Soc 4(3):543–568

McGuire J, Andersen CB (2019) Improving aurecon's employment contracts through visualisation. Univ W Aust Law Rev 46(2):218–236. http://www.law.uwa.edu.au/data/assets/pdf_file/0007/3442651/4.-AndersenMcGuidre-Future-of-Works.pdf. Accessed 23 Aug 2021

Mittelstadt BD, Floridi L (2016) The ethics of big data: current and foreseeable issues in biomedical contexts. Sci Eng Ethics 22(2):303–341

Mittelstadt BD, Allo P, Taddeo M, Wachter S, Floridi L (2016) The ethics of algorithms: mapping the debate. Big Data Soc. https://doi.org/10.1177/2053951716679679

National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (1978) The Belmont Report. United States Government Printing Office, Washington, DC

National Health and Medical Research Council (2007) The National statement on ethical conduct in human research. Available from: https://www.nhmrc.gov.au/about-us/publications/national-statement-ethical-conduct-human-research-2007-updated-2018. Accessed 23 Aug 2021

Newson A, Lipworth W (2015) Why should ethics approval be required prior to publication of health promotion research? Health Promot J Aust 26(3):170–175. https://doi.org/10.1071/HE15034

Nimmon LS, Stenfors-Hayes T (2016) The "handling" of power in the physician-patient encounter: perceptions from experienced physicians. BMC Med Educ 16:114. https://doi.org/10.1186/s12909-016-0634-0

O'Neil C (2016) Weapons of math destruction: how big data increases inequality and threatens democracy. Crown Publishing Group, New York

O'Neill O (2003) Some limits of informed consent. J Med Ethics 29(1):4–7

Page K (2012) The four principles: can they be measured and do they predict ethical decision making? BMC Med Ethics 13:10. https://doi.org/10.1186/1472-6939-13-10

Pasquale F (2015) The black box society: the secret algorithms that control money and information. Harvard University Press, Cambridge

Peterson A (2016) FCC cracks down on verizon wireless for using 'supercookies'. The Washington post. Retrieved from: https://www.washingtonpost.com/news/the-switch/wp/2016/03/07/fcc-cracks-down-on-verizons-supercookies/. Accessed 23 Aug 2021

Pollach I (2011) Online privacy as a corporate social responsibility: an empirical study. Bus Ethics Eur Rev 20:88–102. https://doi.org/10.1111/j.1467-8608.2010.01611.x

Postelnicu L (2019) Pregnancy club Bounty UK fined £400,000 by data protection regulator. HealthcareITNews. Retrieved from: https://www.healthcareitnews.com/news/pregnancy-club-bounty-uk-fined-400000-data-protection-regulator. Accessed 23 Aug 2021

Powles J, Hodson H (2017) Google DeepMind and healthcare in an age of algorithms. Heal Technol 7:351–367. https://doi.org/10.1007/s12553-017-0179-1

Przybylski AK, Murayama K, DeHaan CR, Gladwell V (2013) Motivational, emotional, and behavioral correlates of fear of missing out. Comput Hum Behav 29(4):1841–1848. https://doi.org/10.1016/j.chb.2013

Purtill J (2021) Apple's iPhone has a new privacy feature that Facebook has tried to stop. ABC News. Retrieved from: https://www.abc.net.au/news/science/2021-04-29/apple-iphone-tracking-operating-system-update-facebook-privacy/100100172. Accessed 23 Aug 2021

Quelle C (2018) Enhancing compliance under the general data protection regulation: the risky upshot of the accountability- and risk-based approach. Eur J Risk Regul 9(3):502–526. https://doi.org/10.1017/err.2018.47

RAGCP (2019) Informed consent: information sheet. Retrieved from: https://www.racgp.org.au/download/Documents/PracticeSupport/informedconsentinfosheet.pdf. Accessed 23 Aug 2021

RANZCOG (2018) RANZCOG medical schools curriculum in obstetrics & gynaecology (AMC Alignment). Retrieved from: https://ranzcog.edu.au/RANZCOG_SITE/media/RANZCOG-MEDIA/About/RANZCOG-Undergraduate-Curriculum-in-Women-s-Health.pdf. Accessed 23 Aug 2021

Robbins S (2019) A misdirected principle with a catch: explicability for AI. Mind Mach 29:495–514. https://doi.org/10.1007/s11023-019-09509-3

Rosmarin R (2020) Sustainability sells: why consumers and clothing brands alike are turning to sustainability as a guiding light. Business Insider. Retrieved from: https://www.businessinsider.com/sustainability-as-a-value-is-changing-how-consumers-shop?r=AU&IR=T. Accessed 23 Aug 2021

Sappideen C (2010) Bolam in Australia: more bark than bite. Univ New South Wales Law J 33(2):386–424

Schmelzer R (2019) Understanding explainable AI'. Forbes. Retrieved from https://www.forbes.com/sites/cognitiveworld/2019/07/23/understanding-explainable-ai/?sh=122b8bc77c9e. Accessed 23 Aug 2021

Shaban-Nejad A, Michalowski M, Buckeridge DL (2021) Explainable AI in healthcare and medicine: building a culture of transparency and accountability. Springer. https://doi.org/10.1007/978-3-030-53352-6

Singer N, Conger K (2019) Google is fined $170 million for violating children's privacy on YouTube. The New York Times. Retrieved from https://www.nytimes.com/2019/09/04/technology/google-youtube-fine-ftc.html. Accessed 23 Aug 2021

Solove DJ (2008) Understanding privacy. Harvard University Press, Cambridge, MA

Stahl BC, Antoniou J, Ryan M, Macnish K, Jiya T (2021) Organisational responses to the ethical issues of artificial intelligence. AI Soc. https://doi.org/10.1007/s00146-021-01148-6

Sunstein C (2002) Risk and reasons: safety, law and the environment. Cambridge University Press, Cambridge

Taylor J (2021) Government agencies could access personal data without consent under new bill. The Guardian. Retrieved from: https://www.theguardian.com/australia-news/2021/may/01/government-agencies-could-access-personal-data-without-consent-under-new-bill. Accessed 23 Aug 2021

Thompson SA, Warzel C (2019) Twelve million phones, one dataset, zero privacy. The New York Times. Retrieved from https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html. Accessed 23 Aug 2021

Thorbecke C (2021) What to know about Apple's new privacy update and why it's riling Facebook. ABC News. Retrieved from https://abcnews.go.com/Business/apples-privacy-update-riling-facebook/story?id=77340719. Accessed 23 Aug 2021

Tsamados A, Aggarwal N, Cowls J, Morley J, Roberts H, Taddeo M, Floridi L (2021) The ethics of algorithms: key problems and solutions. AI Soc. https://doi.org/10.1007/s00146-021-01154-8

Wachter S, Mittelstadt B, Floridi L (2017) Why a right to explanation of automated decision-making does not exist in the general data protection regulation. Int Data Priv Law 7(2):76–99. https://doi.org/10.1093/idpl/ipx005

Walker T (2020) Value of choice. J Med Ethics. https://doi.org/10.1136/medethics-2020-106067

Zuboff S (2019) The age of surveillance capitalism: the fight for a human future at the new frontier of power. Profile Books, London

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