



Legal and Ethical Aspects of Machine Learning: Who Owns the Data?

14

Barbara Prainsack and Elisabeth Steindl

Contents

14.1 Introduction	191
14.2 Opening the “Ethics Bubble”: What Are the Concerns?	192
14.3 Going Beyond FAT: Beyond Medical Ethics	195
14.4 Who Owns Patient Data?	196
14.5 Conclusion	199
References	200

14.1 Introduction

It is no exaggeration to say that we are in the midst of an “AI ethics bubble”. The ethics of artificial intelligence makes headlines in public media and the topic of major international conferences. Technology corporations in particular are channeling funding into the creation of AI ethics institutes and endowed chairs, such as recently seen at universities in Oxford, Munich, and Cambridge, MA (e.g. [1, 2]). While corporations have collaborated with academia for many

decades, if not centuries, what is new here is the strong focus on ethics.

Perhaps this is not surprising, given that AI—used here as an umbrella term for various technologies that mimic human intelligence—has become a symbol for societal concerns about the mastery of machines over people. It is seen as posing various challenges to society, ranging from voter manipulation to other threats to democracy [3], to the technological replacement of human labour [4]. The replacement of human labour is an aspect that is particularly pertinent to medicine as well: Some studies predict that up to half of all the existing jobs in the United States are at risk of automation [5]. Among medical professionals, radiologists and pathologists are seen as particularly vulnerable to technological replacement [6–8]. Against this backdrop, it could be argued, technology companies have a particularly great need to ensure that their devel-

B. Prainsack (✉)
Department of Political Science, University of
Vienna, Vienna, Austria
e-mail: barbara.prainsack@univie.ac.at

E. Steindl
Department of Innovation and Digitalisation in Law,
University of Vienna, Vienna, Austria

opment and use of AI complies with ethical standards.

But there is also a more sinister reason for the current ethics bubble. Corporations that use AI to develop new services, increase market shares, and expand their global reach, are currently pitching “ethics” against “regulation”. Strict regulation of AI, and in particular, machine learning, they argue, puts Europe, North America and other world regions at risk of falling further behind the AI capabilities of China, and is thus problematic. They suggest that rather than putting up “red tape” for technology, societies should ensure the creation of good ethics guidelines that ensure that AI is “trustworthy” ([9], and in reference to [10]). Such playing out of ethics against regulation is, of course, not only politically problematic but also factually flawed: Ethics and regulation take different forms and are issued by different institutions, but they mutually influence and enable each other. Ethical considerations are always part of regulatory processes and guidelines, and regulation, in turn, is necessary to enforce ethical norms and commitments. Also in this chapter, ethical and regulatory and legal aspects are treated as closely intertwined, and not as something that can, or should be, strictly separated.

Before we look at the legal and regulatory aspects of AI in imaging—and zoom into the question of who owns the data that is used for this purpose—let us first look at what the issues are the ethics scholarship has identified in this context.

14.2 Opening the “Ethics Bubble”: What Are the Concerns?

There has recently been a terminological shift in discussions of the ethics of AI. Until about mid-2019, the term “artificial intelligence” was widely used as an umbrella term for all computational processes that mimic human intelligence. More recently, following criticism of the unduly vague and wide use of the term in ethical and regulatory discussions, the terms that are used have become more specific: Policy and academic papers alike increasingly use the term “machine learning” to

denote applications of AI that improve with only very little, or even no, input from humans. Also in this chapter, the term machine learning is used to refer to processes and technologies whereby machines discern patterns in data with only little steering from humans, while “AI” is used to denote instances in which debates refer to even wider areas of machine “intelligence”, or to the attempt to make machines act like humans.

Although AI has a history of many decades (e.g. [11]), there has been an increase in AI technologies in recent years. This is mostly due to increasing computational power and increasing opportunities for automation and digitisation. These, in turn, have been made possible by “datafication”, which means the capturing and storing of information about people’s lives, their bodies, and about their environments, that were previously unrecorded. For example, even a decade ago, the only way to learn about people’s exercise levels was by asking them what type of exercise they had done within a specific period of time, and how much of it. Today, this information is, for many of us, automatically captured by activity trackers built into our smartphones, or measured in other, often remote and unobtrusive ways. The legal scholar Harry Surden called this the end of structural privacy [12], meaning that the domains of our lives and bodies that remain unseen and “unaccounted” are becoming smaller and smaller. There is ever less of us and our lives that is not datafied.

For healthcare, the availability of data about various aspects of the lives and bodies of patients, often over a long period of time, is seen as an unprecedented opportunity. Here, AI is portrayed as an answer to the problem of data interpretation: While the production of data has become relatively cheap, and greater amounts of data are being produced each day, making sense of these data has remained expensive [13]. To bridge this “interpretation gap”, machine learning in particular has been suggested as a solution. Moreover, in many aspects of healthcare, AI is already in use: from telemedicine to supporting communication with patients to billing and insurance. In medical imaging, molecular imaging is expected to benefit significantly from machine learning; and deep

learning based interpretation is hoped to help reduce interobserver variability in nuclear imaging (e.g. [14]; see also [15]).

What are the key ethical challenges related to AI? Over the last years, ethicists and other experts have raised a range of concerns related to AI that can be largely grouped in three clusters: Fairness, accountability, and transparency (FAT). The paradigmatic challenge for fairness is biased training data (see [16, p. 176]): This is the case when a specific population group, such as elderly people, members of minorities, or the uninsured, are underrepresented, or entirely missing, from a data set. It is not always straightforward to know, however, when bias exists, or when it is problematic [17]. For example, in the context of the training of an algorithm to classify pulmonary tuberculosis (e.g. [18]), what constitutes a non-biased dataset: A dataset that is representative of people who typically suffer from TB? One that reflects the demographic composition of the patient population treated in a specific hospital? Or a dataset that represents the demographic composition of the city? Of the entire nation even? Moreover, if it is known, for example, that minority populations have been underrepresented in training data for machine learning for years, would it be mandated for ethical reasons to oversample members of the minority populations in question to make up for previous discrimination? There are no definitive answers to these questions; instead, they illustrate the intricacies of knowing when a bias exists, and when a bias is problematic, that is, when it has a negative impact on equity.¹

¹It is mandated here to clarify the difference between inequality and inequity. The two terms are often conflated in common parlance, but they mean different things. Inequality means that resources or benefits are distributed unequally over different groups. Using the example of health outcomes, if women and men have different life expectancies, that is an inequality. Not all inequalities, however, are also unfair: if the different outcome can be explained by voluntary actions, for example. If Laura and Amir, who are married, and who grew up in similar social strata and in the same town, have different health status because Amir likes to tend to the garden in his spare time while Laura goes paragliding, and due to multiple sport-

While fairness ultimately pertains to questions about equity, the second criterion within the FAT paradigm, accountability, relates to the question of who can be held responsible for outcomes. Here, also legal questions about liability come into play. Very broadly speaking (and without consideration of specific configurations in particular jurisdictions; for more details on these, see [16, 19]), liability for harm caused by machine learning applications can only kick in when someone has been negligent, either a physician or a company. Negligence on the side of physicians or healthcare workers, in turn, requires that there is a duty of care towards patients that was breached. As Schönberger emphasises, not all erroneous predictions by an AI system that caused harm to a patient mean that physicians or healthcare organisations they work for are liable; they can only be held accountable if they used the AI in a way that they should not have [16, p. 197].

The other type of liability besides that of physicians and healthcare providers is product liability. This becomes relevant when patients suffer harm from products that were defective in their design, manufacturing, or warning—in other words, products that did not operate as they should have. The legal concept of liability was developed with the idea in mind that those held liable would be people, not machines. They were written for people who have a sense of responsibility, which machines do not have. Moreover, machines would not be affected by any of the conventional sanctions (e.g. fines) that our law system applies. Algorithms, in contrast to book titles that suggest otherwise (e.g. [20]), do not “want” things—they are not human. This raises a few issues when liability laws are applied to machines: First, if AI works in the form of non-embedded software (meaning that the software is not built into other machines such as phones, cars, or pacemakers) then it is not clear whether it is covered by existing liability legislation such as

ing accidents she now suffers chronic pain, then the difference in health status between them is not an inequity. As a rule of thumb, if we cannot find any factor that justifies different outcomes, then we should treat different outcomes as inequities.

the European Union’s Product Liability Directive, for example. Second, current approaches bypass the problem that the legal concept of liability was designed to apply to humans by holding the people who build or use the machines liable for the actions of the machines. As Schönberger argues [16], the more “autonomous” machines become, that is, the less their actions can be traced back to decisions taken by humans, the more difficult it becomes to hold the humans “behind” the machines accountable. Scholars are discussing a number of ways to address these problems. These include giving some kind of personhood status to intelligent machines (e.g. [21])²; another solution that is discussed is to hold the healthcare professionals that are using AI even more strictly accountable for the “decisions” of the machine than at present. For example, doctors would then be responsible for harm if they did not take adequate measures to evaluate how accurate the algorithm is that they are using [16].

The last notion in the FAT-paradigm is transparency. At times, transparency is a precondition of liability, and at other times, it goes beyond it. While liability refers to the consequences for someone who bears responsibility for something in the case of harm (i.e. in the case of negligence or even intentional wrongdoing), a certain level of transparency is required for the assessment of whether any wrongdoing took place. Especially in the context of unsupervised machine learning, where no function is associated with the input,³ it is often difficult, if not impossible, to know how the software arrived at a specific outcome because the path to achieving the outcome was not designed into the system, and is impossible

²The European Parliament has adopted a resolution in 2017 with recommendations to the Commission on Civil Law Rules on Robotics suggesting to prompt a legal status for robots (https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html?redirect#BKMD-12). The Commission, however, did not follow this recommendation in its recent strategies addressing AI.

³Within supervised machine learning, the machine is told by a human what to look for: e.g. it is shown pictures of dogs and then asked to look for dogs in other images. Within unsupervised machine learning, the machine is not told what to look for, but just commanded to look for patterns.

Table 14.1 A graded scale ethical scrutiny of machine learning in healthcare

Level of ethical sensitivity	Use of AI
Low	AI to support non-medical aspects (e.g. scheduling, video-conferencing)
Intermediate	AI to support diagnosis or treatment choice (“thinking AI”)
High	AI to make decisions (“acting AI”)

for observers to understand. It is because of this lack of transparency that some ethicists have argued that the use of unsupervised machine learning in healthcare is ethically more problematic than supervised machine learning [22]. Such proposals, however, neglect the question of where in healthcare machine learning is put to use. If it is used in core medical contexts, such as for diagnosis and treatment decisions, then the lack of transparency seems much more concerning than if unsupervised machine learning is used within an application to enable video consultations. For this reason, we propose a graded scale ethical scrutiny of machine learning in healthcare (Table 14.1) that distinguishes between three levels of ethical sensitivity: At the lowest level of concern are uses of machine learning (and other AI) for non-medical aspects, such as appointment scheduling or videoconferencing. At the intermediate level are applications of machine learning in key medical activities such as the establishment of a diagnosis or treatment decision, but where machine learning is only aiding human decision making without suggesting a final decision (“thinking AI”). At the highest level of ethical sensitivity is the use of machine learning for key medical activities where the software makes the decision, e.g. if a machine that automatically classified a disease and gave a treatment decision that was binding (“acting AI”), which is so far not part of routine clinical care.

Other factors that are to be considered include whether or not machine learning is supervised (which is less ethically problematic because of higher level of transparency) or unsupervised (more ethically sensitive due to lower levels of

transparency), whether the tool has been validated, and whether the people using the tool are conscious of the possibility and consequences of potential bias (“fairness through awareness”, [23]).

14.3 Going Beyond FAT: Beyond Medical Ethics

Ethics guidelines, ethics codes, as well as papers addressing ethical concerns in connection with AI in healthcare regularly discuss phenomena that map against the FAT paradigm—even if they discuss these issues under different labels. But there are also contributions that raise bigger questions. A statement by the European Group on Ethics on AI, robotics and “autonomous” systems” (2018), for example, draws attention to the need for AI to be put in the service of broader societal and ethical values, including human dignity, responsibility, democracy, justice, equity, solidarity, sustainability, and deliberation. Moreover, scholars such as Karen Yeung use the term “ethics washing” to refer to situations where AI ethics serves mostly as an empty vessel that can be filled with any content that seems suitable, and where ethics lacks the necessary tools to enforce its own claims [9]. Taken together, these points of critique call for an ethics that does not accept current institutional arrangements and configurations of power as they are, and within these, try to make AI “more ethical”. Instead, they call for a political ethics that is concerned also with how new technological practices affect the distribution of entitlements, duties, and resources within and across populations. The FAT paradigm goes some way in that direction, but not far enough.

An important underpinning of such a more political ethics of AI is to leave the specificities of medical ethics behind, and instead treat AI ethics as a form of data ethics. A key argument in favour of the latter is that many ethical issues in connection with machine learning emerge due to the integration and use of large amounts of personal data. But such a move from medical to data ethics may not be as easy to do as it may seem. It would require a fundamental shift in the points of refer-

ence used by ethics frameworks—most prominently the focus only on individual rights. As many scholars have argued, most of the risks in connection with data use are personal and collective, and they cannot be broken down into individual bits (e.g. [24]). Moreover, many of the scholars and approaches that are populating the rapidly growing field of AI ethics were trained in medical ethics or bioethics. It will be difficult to expand (and, in some cases, change) the reference points and institutional structures that these experts are operating with and within.

What is the problem with the categories and focus points of medical ethics—why can they not be transposed to AI ethics? The main reason is that the key reference point of medical ethics is the human body; the early codifications of medical ethics established that people have a right to be informed about, and consent to, what happens to their bodies. This framework emerged partly in response to the horrific human rights infringements of the Nazi period and other instances when harmful or even torturous “experiments” were imposed on people under the guise of science. Data ethics, on the other hand, does not take the physical body as its reference point, but the “data body”—which is of a very different nature. First of all, the data body does not have clear borders and boundaries; the data that represents a person, namely, the data capturing her behaviour, her diseases, etc., is spread over many places and can be accessed by many people at the same time. This means, also, that the frame of an intervention that medical ethics operates with does not work for data ethics. An intervention into a person’s data is not comparable to a body that is operated on to take out a gallbladder, or to test a new drug. There is often no clear beginning and no clear end to an “operation” on a dataset—data is interrogated continuously [25]. In addition, in traditional medical ethics, it is normally clearly apparent who carries out the procedure and who is at risk: The latter is normally the patient. In data ethics, “procedures” can be carried out by many different people in different places at the same time—primary and secondary data users (the latter are researchers, for example, who reuse datasets from other research teams, or

even from the clinic), commercial enterprises, etc. The people at risk from these procedures can be totally unrelated from those who have given their data. In other words, risks in data ethics are not limited to specific individuals, but they are collective.

Understanding AI ethics as a kind of data ethics, and not as a field of application for medical ethics, also affects how we think about data ownership.

14.4 Who Owns Patient Data?

This simple question is not easy to answer. It will concern us for the rest of this chapter. The problem starts with defining ownership. While the related term “property” has clearly definable legal meaning, ownership can relate to legal entitlements, but it can also refer to a moral claim on something. People who say that they own their personal data do not always mean to express a legal opinion. Rather than implying that they have the right to destroy or sell their data, which are some of the key characteristics that distinguish property rights from other entitlements, what they often mean to say is: “I should have a say in who uses my data, what they do with it, and who benefits from it”. In other words, ownership is a very broad concept that includes moral and legal elements.

But let’s start at the beginning. Can we legally own data? In other words, is it possible to own something that is (at least in part) immaterial—as digital data is (see [26])? The law answers this question affirmatively; intellectual property rights protection is an example. It gives people or organisations the right to control intellectual resources that are in part, or even entirely, immaterial.

Within the European Union, the EU General Data Protection Regulation (GDPR) grants special protections to so-called personal data, that is, data that refers to a specific identified or identifiable natural person. Names and addresses are clearly personal data; but IP addresses or genomes are too [27]. Personal data is seen as disclosing things about people and their lives that they may

want to be confidential or even private, and people may suffer harm if this data and information are known or used by others. For these reasons, not only GDPR, but most jurisdictions place restrictions on the collection and use of personal data. But there are crucial differences in how personal data is protected. To put it very generally, in Europe, the predominant view has been to see personal data and information as belonging to people in a moral sense, without being considered property in the legal sense. This means that personal data is not seen as something to be sold, or something that has a market value. The protection of personal data is ensured through privacy rights.

According to European Law, the question of whether data can be owned has multiple layers. One layer refers to the fact that any data has to be categorised as either personal or non-personal data. Personal data is protected by a number of fundamental individual rights, such as the right to be informed, the right of access, the right to rectification, the right to erasure, the right to restrict processing, the right to data portability, the right to object, and rights in relation to automated decision-making and profiling (see Chap. 3 GDPR). These individual rights continue to exist as long as the data has not been anonymised—this means that, taking into account all the means reasonably likely to be used, the data does no longer relate to an identified or identifiable person (i.e. all links to do so have been destroyed). In other words the conception of personal data within the GDPR cannot be aligned with a third party owning somebody else’s personal data.⁴ It also means that, in the European context, the question of ownership only arises regarding non-personal data. And this is where the next layer comes in: as data does not easily fit into either one of the traditional legal categories of material or immaterial, it cannot be subsumed under property that is moveable or intellectual property. The European Commission itself stressed that current

⁴The question of lawfulness of processing of special categories of personal data according to Article 9 GDPR has to be seen apart from any kind of possible ownership and is therefore not discussed here.

intellectual property laws are not a suitable tool for data governance [28].

In the United States, debates about whether personal information should or could be viewed as property have been complex. Some authors see property rights as the best way of protecting personal data [29]. Partially, this notion is rooted in the important role that property rights play in American self-conception. Property rights, understood—in William Blackstone’s deliberately provocative description—as ‘that sole and despotic dominion which one man claims and exercises over the external things of the world, in total exclusion of the right of any other individual in the universe’ [30], are woven into the very foundations of American society and legal culture. Even for those scholars who say that this ideal has never been implemented in actual practice, property rights have nevertheless played a much more important role in the United States than in Europe.

In U.S. discourse, treating personal data as property has served the important purpose of overcoming the shortcomings of U.S. data protection systems [31, pp. 507–508]. In contrast to the European Union, who have a data protection law that applies to the processing of all personal data and expands its territorial scope even beyond European borders, American privacy laws are sector-specific; they are tailored to specific fields such as healthcare or financial services. This has led some scholars to argue that, because American privacy laws are relatively weak, property rights are the best, or even the only, way to ensure people’s control over their data.

Other authors (e.g. [32, p. 1295]) disagree with this stance. They argue that “the *raison d’être* of property is alienability” [32, p. 1295]. The meaning of this statement becomes clear only if we take a closer look at how property rights are organised: It is best conceived as a bundle of entitlements, rather than as one single right. It is the bundle of rights, rather than one specific characteristic, that sets property rights apart from other entitlements to things. Within that bundle, there are some “stand-out” rights that characterise the bundle.

To use an example from the physical world: When someone has borrowed a book from a library, the book is in her possession. She is enti-

tled to do a lot of things: to read the book, to control who else gets to read it, and she can use it for other purposes such as place a laptop on top of it for a videoconference. She can exclude other people from even looking at it. But there are things that this person who has taken a book from the library is not entitled to do: She must not sell or destroy the book. These additional entitlements are reserved to the person or entity that holds property rights. In other words, the bundle of rights granted to a person due to mere possession (e.g. having the book in your house after having taken it from the library) is less “thick” than the bundle of property rights. Property rights include all rights that other forms of possessions include (the right to possession, income, etc., as listed below) plus the right of alienation (selling or destroying).

Another example of the difference between weaker forms of possession on the one hand, and property rights on the other, is renting a flat. As a lawful tenant I am entitled to determine who can enter the flat, how it is decorated, and what is done inside. But only the owner (here: the holder of property rights) holds the additional rights that are also in the bundle, such as selling the flat. (The fact that I am not normally allowed to destroy my flat, even if I hold property rights, illustrates that even property rights are not unlimited—even they can be restricted to protect important other rights and interests. In the interest of public safety and security I am not allowed to burn down my flat, or to neglect it to such an extent that it becomes a public nuisance).

Back to digital data. But how does this difference between property rights and “weaker” forms of possession that apply to tangible goods such as books or flats work with intangible things such as data? As noted, although data has a tangible, material element, including the technical infrastructures that enable its collection, storage, and use, at least a part of them is immaterial.

In order to answer this question it is helpful to unpack the bundle of rights and entitlements that make up property rights. Denise Johnson [33], drawing upon Honore’s famous work in the 1960s [34], names the following entitlements as part of the bundle of property rights:

1. The right to **possess**. Just as the example of the library book, or the rented flat, below, the person who rightfully possesses has exclusive control of a thing. When the thing that is owned is intangible, then, as Honore put it, possession is the right to exclude others from using or benefitting from the thing. Moving to the digital realm, for data in the healthcare domain, such as imaging data and lab results, it is very difficult to conceive what such “exclusive” control would look like. When an imaging department that does a cardiac perfusion scan on a patient owns the imaging data (because the patient may have agreed to this when signing the consent form for the procedure) “exclusive control” means that they can share the data with third parties—they can even sell the data. But does it mean that they can exclude the patient from accessing their own perfusion scan? Wherever GDPR is applicable, this stance would be difficult to argue—because as long as the perfusion scan is seen as personal data—i.e. as data that is linked to an identified or identifiable person (note that this includes pseudonymised data)—then the patient has a right to access—or even initiate the erasure—of her own data even though she does not hold property rights to it [35].
2. The right to **manage** gives people the right to decide who can use the thing that is possessed, and how. It includes the right of lending or contracting out (see also [36]). This right seems relatively unproblematic in connection with digital data, except that it may be difficult to exclude patients from using their own data as long as this data is considered personal data—as explained in point (1). Referring to our example of the perfusion scan explained above, this means that the entity that holds property rights to the perfusion scan data can decide who gets access to it, for what purpose it can be used, and who can commercialise it. They may not, however, be able to refuse patients access as long as the imaging data can be linked to an identified or identifiable person.
3. The right to **income** allows the property rights holder to allow others to use the thing and to pay her for this use. This right is closely related to the previous one, namely the right to manage; the difference between the two is that the right to income focuses on the money that one receives in return—for other people using the thing, for example (see also [36]). This seems no more difficult to enforce in the case of digital data than it is with owning a physical object.
4. The right to **capital**—which is the right that allows a person to alienate the thing, namely to give it away, to consume it, to change it, or to destroy it. The problem here is that it is not so easy to decide what “consuming” or “destroying” data means. Physical things are consumable and rivalrous: They can be ‘used up’, and the use of the good by one person affects the use of the good by others. Many authors argue that the same cannot be said for digital data, as they are considered to be neither consumable nor rivalrous: The perfusion scan data does not disappear, or deteriorate, if lots of people use it; and one research group using it does not detract from the utility of the data for another. Having said this, whereas the data itself is not consumable or rivalrous, their value can be: the value of a dataset can be highest for those who have exclusive use; and it can, of course, be affected by many people using it. Think of proprietary information such as search algorithms, or information on commercial mergers that are likely to affect stock prices, for example. For these reasons, digital data is best described as simultaneous [26]: It can be in more places than one at the same time, it can be copied and used by several people at the same time, independent of what the others are doing, and it leaves traces even when it is deleted. Because the value of data can be rivalrous, it is arguably this multiplicity of data that is the key difference between physical entities and digital data with regard to the right to capital. In situations where those holding property rights to data cannot

control all copy of the dataset (or do not even know where all the different copies are), the right to capital may be difficult to enforce.

5. The right to **security** protects the rights-holder from expropriation. In Quigley's words [36, p. 633], it is "the assurance that a person [...] will not be forced to give it up without adequate recompense." It is not difficult to conceive of this right with respect to digital data.
6. The power of **transmissibility** means that the rights holder can give the thing that s/he owns to somebody else, either before or after his/her death. Also here, it is not difficult to imagine this right to be applied to digital data (for the instrument of post-mortem data donation specifically, see [26, 37]).
7. The **absence of term**: This means that the length of ownership is not time-limited.
8. Now we are moving into the provisions within the bundle of property rights that are duties and liabilities rather than entitlements: The first one is the **prohibition of harmful use**, meaning that even the person who owns a thing is not free to do with it whatever she pleases; the boundaries of her freedom are the rights of others. In the physical world this is best described with a knife: Even if I hold all entitlements of the bundle of property rights to the knife I am not allowed to use it to cut into another person. With regard to data, the prohibition of harmful use raises really interesting questions: Does this only mean that the data owner herself is not allowed to use the data in a harmful way? Or does it include a duty to actively prevent that others can use the data in a harmful way? Does this mean that restrictions of data sharing may be required as a preventive measure? These questions remain open.
9. Those who hold property rights are also **liable to execution**; which means that the thing that is owned can be taken away for the repayment of a debt, for example. It is conceivable that this would apply to digital data: if the data has commercial value, ownership

of a dataset could be taken away to pay for something that the rights holder owns.

10. Last but not least, property rights have a **residuary character**: This means that, even if the property rights holder has given away many entitlements within the bundle (e.g. she has leased her property to someone else), she still holds whatever is left of the bundle. To the extent that the bundle of property rights can be applied to digital data, the residuary character does not pose any additional complications.

In sum, many of the entitlements and duties within the bundle of rights that constitute property rights—which were originally developed for physical things—cannot be neatly transposed to digital data. Because of the multiple nature of digital data (the ability of digital data to be at several places at the same time), it is more useful to speak about the right to control data in the context of medical imaging than about data ownership. Because of the complexities laid out in this chapter, and because of the moral and legal connotations of the term, the notion of ownership tends to confuse more than it clarifies when applied to digital data.

14.5 Conclusion

This chapter started with the diagnosis that we are amidst an "AI ethics bubble", where especially corporate interest in ethics of AI and machine learning is extremely high. Technology corporations and other businesses provide funding for ethics institutes and endowed chairs on AI ethics at leading universities, and co-opt academics into the ethics governance of their own companies. The pitching of "ethics" against "regulation" has been part of this process.

Taking the stance that ethics and regulation, albeit having different emphases, complement and require each other, rather than being clearly separable, this chapter then opened up the "ethics bubble" of AI. Our diagnosis was that most of the ethical concerns identified and discussed in this

context map against the so-called FAT paradigm. It orders concerns in several clusters, including fairness, accountability, and transparency. While this typology is extremely helpful, we proposed to take a step further and go beyond the FAT paradigm. In order to do so, we suggested to go beyond the toolbox of medical ethics and draw more strongly upon the instruments in the growing field of data ethics. This is necessary, we argued, because the reference point of medical ethics is the physical body, which has clear boundaries. The same does not apply to people's data bodies, which are far from clearly bounded: Data is multiple in the sense that it can be in several places at the same time.

What, then, does this mean for the question of data ownership? Who owns the data that medical imaging departments work with? The final section of this chapter seeks to answer this question by discussing how the "bundle of rights" that make property rights can be applied to digital data. We conclude that because of the multiple nature of digital data, some of the entitlements and duties within the bundle of property rights can be applied to digital data only with difficulty.

References

- Moss E, Metcalf J. The ethical dilemma at the heart of big tech companies. *Harvard Business Rev.* 2019. <https://hbr.org/2019/11/the-ethical-dilemma-at-the-heart-of-big-tech-companies>. Accessed 24 Apr 2020.
- Ochigame R. 2019. The invention of "ethical AI". *The Intercept*. <https://theintercept.com/2019/12/20/mit-ethical-ai-artificial-intelligence/?comments=1>. Accessed 24 Apr 2020.
- O'Neill C. *Weapons of math destruction: how big data increases inequality and threatens democracy*. Crown. 2016.
- Ford M. *Rise of the robots: technology and the threat of a jobless future*. New York: Basic Books; 2015.
- Frey CB, Osborne MA. *The future of employment: how susceptible are jobs to computerisation?* *Technol Forecast Soc Chang.* 2017;114:254–80.
- Chockley K, Emanuel E. The end of radiology? Three threats to the future practice of radiology. *J Am Coll Radiol.* 2016;13(12):1415–20.
- Grace K, Salvatier J, Dafoe A, Zhang B, Evans O. When will AI exceed human performance? Evidence from AI experts. *J Artif Intell Res.* 2018;62:729–54.
- Obermeyer Z, Emanuel EJ. Predicting the future—big data, machine learning, and clinical medicine. *N Engl J Med.* 2016;375(13):1216.
- Yeung K, Howes A, Pogrebna G. AI governance by human rights-centred design, deliberation and oversight: an end to ethics washing. *The Oxford handbook of AI ethics*. Oxford: Oxford University Press; 2019.
- European Commission Ethics Guidelines for Trustworthy AI. 2019. <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>. Accessed 16 May 2020.
- Haenlein M, Kaplan A. A brief history of artificial intelligence: on the past, present, and future of artificial intelligence. *Calif Manag Rev.* 2019;61(4):5–14.
- Surden H. Structural rights in privacy. *SMUL Rev.* 2007;60:1605.
- Prainsack B. Precision medicine needs a cure for inequality. *Curr Hist.* 2019;118(804):11–5.
- Choi H. Deep learning in nuclear medicine and molecular imaging: current perspectives and future directions. *Nucl Med Mol Imaging.* 2018;52(2):109–18.
- Choi H, Ha S, Im HJ, Paek SH, Lee DS. Refining diagnosis of Parkinson's disease with deep learning-based interpretation of dopamine transporter imaging. *Neuroimage Clin.* 2017;16:586–94. <https://doi.org/10.1016/j.nicl.2017.09.010>.
- Schönberger D. Artificial intelligence in healthcare: a critical analysis of the legal and ethical implications. *Int J Law Inform Technol.* 2019;27(2):171–203.
- Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classification of skin cancer with deep neural networks. *Nature.* 2017;542(7639):115–8.
- Lakhani P, Sundaram B. Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks. *Radiology.* 2017;284(2):574–82.
- Pesapane F, Volonté C, Codari M, Sardanelli F. Artificial intelligence as a medical device in radiology: ethical and regulatory issues in Europe and the United States. *Insights Imaging.* 2018;9(5):745–53.
- Finn E. *What algorithms want: imagination in the age of computing*. Cambridge, MA: MIT Press; 2017.
- Vladeck DC. Machines without principals: liability rules and artificial intelligence. *Washington Law Rev.* 2014;89:117.
- Jannes M, Friele M, Jannes C, Woopen C. *Algorithms in digital healthcare. An interdisciplinary analysis*. Gütersloh: Bertelsmann Stiftung; 2019.
- Dwork C, Hardt M, Pitassi T, Reingold O, Zemel R. Fairness through awareness. In: *Proceedings of the 3rd innovations in theoretical computer science conference*; 2012. p. 214–26.
- Taylor M. *Genetic data and the law: a critical perspective on privacy protection*. Cambridge: Cambridge University Press; 2012.
- Metcalf J, Crawford K. Where are human subjects in big data research? The emerging ethics divide. *Big Data Soc.* 2016;3(1):1–14. <https://doi.org/10.1093/bds/baw001>.

- [org/10.1177/2053951716650211](https://doi.org/10.1177/2053951716650211). Accessed 16 May 2020.
26. Prainsack B. Data donation: how to resist the iLeviathan. In: *The ethics of medical data donation*. Cham: Springer; 2019. p. 9–22.
 27. Goddard M. The EU General Data Protection Regulation (GDPR): European regulation that has a global impact. *Int J Mark Res*. 2017;59(6):703–5.
 28. European Commission Legal study on ownership and access to data. Final report. 2016. <https://www.op.europa.eu/s/n2Qc>. Accessed 16 May 2020.
 29. Murphy RS. Property rights in personal information: an economic defence of privacy. *Georgetown Law J*. 1996;84:2381–217.
 30. Blackstone W.. Of property in general. Commentaries on the laws of England. 1765–69; Book II: Chapter I. 1979. https://avalon.law.yale.edu/subject_menus/blackstone.asp. Accessed 12 May 2018.
 31. Purtova N. Property rights in personal data: learning from the American discourse. *Comput Law Secur Rev*. 2009;25(6):507–21.
 32. Litman J. Information privacy/information property. *Stanford Law Rev*. 2000;52:1283–313.
 33. Johnson DR. Reflections on the bundle of rights. *Vermont Law Rev*. 2007;32:247. <https://lawreview.vermontlaw.edu/wp-content/uploads/2012/02/johnson2.pdf>
 34. Honoré AM. Ownership. *Making law bind: essays legal and philosophical*. Oxford: Clarendon Press; 1961. p. 161–92 (Originally published in Guest AG, ed. *Oxford essays in jurisprudence*. Oxford: Oxford University Press; 1961. p. 107–47).
 35. Thorogood A, Bobe J, Prainsack B, Middleton A, Scott E, Nelson S, Corpas M, Bonhomme N, Rodriguez LL, Murtagh M, Kleiderman E. APPLaUD: access for patients and participants to individual level uninterpreted genomic data. *Hum Genomics*. 2018;12(1):7.
 36. Quigley M. Property and the body: applying Honoré. *Med Law Rev*. 2007;17:457.
 37. Krutzinna J, Floridi L, editors. *The ethics of medical data donation*. Cham: Springer International Publishing; 2019.