



Ethical and Legal Challenges of Artificial Intelligence-Driven Health Care

Gerke, Sara; Minssen, Timo; Cohen, Glenn

Published in:
Artificial Intelligence in Healthcare

Publication date:
2020

Document license:
[Unspecified](#)

Citation for published version (APA):
Gerke, S., Minssen, T., & Cohen, G. (2020). Ethical and Legal Challenges of Artificial Intelligence-Driven Health Care. In A. Bohr, & K. Memarzadeh (Eds.), *Artificial Intelligence in Healthcare* (1 ed.). Elsevier.

Ethical and Legal Challenges of Artificial Intelligence-Driven Health Care

This draft is current as of April 23, 2020, for the definitive version check out *Artificial Intelligence in Healthcare*, Adam Bohr, Kaveh Memarzadeh (eds.), ISBN: 9780128184387, Elsevier when published.

Authors:

Sara Gerke^{1*}

Petrie-Flom Center, PMAIL, Harvard Law School

Timo Minssen²

CeBIL, University of Copenhagen

I. Glenn Cohen³

Harvard Law School

¹ Address correspondence to Sara Gerke, The Petrie-Flom Center for Health Law Policy, Biotechnology, and Bioethics at Harvard Law School; The Project on Precision Medicine, Artificial Intelligence, and the Law (PMAIL); Harvard University; 23 Everett Street, MA 02138, Cambridge, USA. E-mail: sgerke@law.harvard.edu.

² Timo Minssen, Centre for Advanced Studies in Biomedical Innovation Law (CeBIL), Karen Blixens Plads 16, 2300 Copenhagen, DK. E-mail: timo.minssen@jur.ku.dk.

³ I. Glenn Cohen, Harvard Law School, Griswold 503, Cambridge, MA 02138, USA. E-mail: igcohen@law.harvard.edu.

Abstract

This chapter will map the ethical and legal challenges posed by artificial intelligence (AI) in health care and suggest directions for resolving them. Section 1 will briefly clarify what AI is and Section 2 will give an idea of the trends and strategies in the United States (U.S.) and Europe, thereby tailoring the discussion to the ethical and legal debate of AI-driven health care. This will be followed in Section 3 by a discussion of four primary *ethical* challenges, namely (1) informed consent to use, (2) safety and transparency, (3) algorithmic fairness and biases, and (4) data privacy. Section 4 will then analyze five *legal* challenges in the U.S. and Europe: (1) safety and effectiveness, (2) liability, (3) data protection and privacy, (4) cybersecurity, and (5) intellectual property law. Finally, Section 5 will summarize the major conclusions and especially emphasize the importance of building an AI-driven health care system that is successful and promotes trust and the motto “*Health AIs for All of Us*”.

Keywords: Artificial Intelligence (AI), ethical challenges, U.S. and EU law, safety and effectiveness, data protection and privacy

From clinical applications in areas such as imaging and diagnostics to workflow optimization in hospitals to the use of health apps to assess an individual's symptoms, many believe that artificial intelligence (AI) is going to revolutionize health care. Economic forecasters have predicted explosive growth in the AI health market in the coming years; according to one analysis, the market size will increase more than tenfold between 2014 and 2021 (Accenture, 2017). With this growth comes many challenges, and it is crucial that AI is implemented in the health care system ethically and legally. This chapter will map the ethical and legal challenges posed by AI in health care and suggest directions for resolving them.

We will begin by briefly clarifying what AI is and giving an overview of the trends and strategies concerning ethics and law of AI in health care in the United States (U.S.) and Europe. This will be followed by an analysis of the ethical challenges of AI in health care. We will discuss four primary challenges: (1) informed consent to use, (2) safety and transparency, (3) algorithmic fairness and biases, and (4) data privacy. We then shift to five *legal* challenges in the U.S. and Europe, namely (1) safety and effectiveness, (2) liability, (3) data protection and privacy, (4) cybersecurity, and (5) intellectual property law. To realize the tremendous potential of AI to transform health care for the better, stakeholders in the AI field, including AI makers, clinicians, patients, ethicists, and legislators, must be engaged in the ethical and legal debate on how AI is successfully implemented in practice.

1 Understanding AI	
2 Trends and Strategies	
Challenges	
3 Ethical	4 Legal
3.1 Informed Consent to Use	4.1 Safety and Effectiveness
3.2 Safety and Transparency	4.2 Liability
3.3. Algorithmic Fairness and Biases	4.3 Data Protection and Privacy
3.4 Data Privacy	4.4 Cybersecurity
	4.5 Intellectual Property Law

Table. 1: Overview of this Chapter

1 Understanding “Artificial Intelligence” (AI)

The term “artificial intelligence”, or in abbreviated form “AI”, is widely used in society but its precise meaning is contested both in scholarly work and legal documents and we will not insist on a single definition here but instead pick out a few sub-types: Machine Learning (ML), a subset of AI, has been the most popular approach of current AI health care applications in recent times since it allows computational systems to learn from data and improve their performance without being explicitly programmed (Mehta and Devarakonda, 2018, p. 2020). Deep learning, a subset of ML, employs artificial neural networks with multiple layers to identify patterns in very large data sets (Yu, Beam, and Kohane, 2018, p. 720; Mehta and Devarakonda, 2018, p. 2020). Most notably, as we will see below, there are additional ethical and legal challenges in cases

where ML algorithms are closer to “black boxes” (i.e., the results are very difficult for clinicians to interpret fully) (Yu et al., 2018, p. 727; Mehta and Devarakonda, 2018, pp. 2019–2021).

2 Trends and Strategies

In this section we discuss the U.S. and Europe’s strategies for AI and how they strive to compete against their biggest competitor China, thereby tailoring the discussion to the ethical and legal debate of AI in health care and research. We will also look at AI trends and discuss some examples of AI products that are already in clinical use in the U.S. and Europe.

2.1 U.S.

During Barack Obama’s presidency, the U.S. Government’s reports on AI emphasized, among other things, the applications of AI for the public good as well as aspects of fairness, safety, and governance (U.S. Government, 2016a, pp. 13, 14 and 30–34, 2016b, and 2016c). One of the reports also stressed the need to improve fairness, transparency, and accountability-by-design as well as building ethical AI (U.S. Government, 2016b, pp. 26, 27).

Since Donald Trump’s presidency, the U.S. AI strategy has shifted to a more free market-oriented approach (Dutton, 2018). The White House, for instance, hosted the Artificial Intelligence for American Industry Summit in May 2018. One of the key takeaways from the summit breakout discussions was that the Trump Administration aims to remove regulatory barriers to AI innovations (White House, 2018, pp. 3, 5). In July 2018, the Executive Office of the President announced that American leadership in AI is one of the top Administration R&D budget priority areas for 2020 (pp. 1, 2). In February 2019, Trump signed the “Executive Order on Maintaining American Leadership in Artificial Intelligence” to address the criticism that the U.S. has taken a hands-off approach to AI in contrast to other countries such as China (White House, 2019; Knight, 2019). With this executive order, Trump launched a coordinated Federal

Government strategy, namely the American AI Initiative, guided by five key areas of emphasis: (1) investing in AI R&D, (2) unleashing AI resources, (3) setting AI governance standards, (4) building the AI workforce, and (5) international engagement and protecting the advantage of the U.S. in AI (White House, 2019; U.S. Government, 2019).

Only recently, in January 2020, the White House published draft guidance for the regulation of AI applications. It contains ten principles that agencies should consider when formulating approaches to AI applications: (1) public trust in AI, (2) public participation, (3) scientific integrity and information quality, (4) risk assessment and management, (5) benefits and costs, (6) flexibility, (7) fairness and non-discrimination, (8) disclosure and transparency, (9) safety and security, and (10) interagency coordination (White House, 2020a). In February 2020, the Whitehouse also published an annual report on the American AI Initiative, summarizing the progress made since Trump signed the executive order. This report, for example, highlights that the U.S. led historic efforts on the development of the Organization for Economic Co-operation and Development (OECD) Principles of AI that were signed by over 40 countries in May 2019 to promote innovative and trustworthy AI and respect democratic values and human rights (White House, 2020b; OECD, 2019). In June 2019, the G20 also released AI Principles drawn from the OECD Principles of AI (White House, 2020b, p. 22; G20, 2019).

The White House has also launched a new website (“AI.gov”) that focuses on AI for the American people and aims to provide a platform for those who wish to learn more about AI and its opportunities.

There are also numerous AI-related bills that have been introduced in the U.S. Congress since Trump’s inauguration on 20 January 2017, such as the *SELF DRIVE Act* (H.R.3388), the *FUTURE of Artificial Intelligence Act of 2017* (H.R.4625 and S.2217), and the *AI JOBS Act of*

2019 (H.R.827). The SELF DRIVE Act is the only bill that has passed one chamber (i.e., the U.S. House of Representatives), and none of these bills are directly related to the ethical and legal aspects of AI in health care. However, the two bills of the FUTURE of Artificial Intelligence Act of 2017, for example, stipulates the Secretary of Commerce to set up a Federal advisory committee that shall provide advice to the Secretary (Sec.4(a) and (b)(1)). This committee shall also study and assess, inter alia, how to incorporate ethical standards in the development and implementation of AI (Sec.4(b)(2)(E)) or how the development of AI can affect cost savings in health care (Sec.4(b)(2)(L)). There are also legal developments related to AI at state and local levels (FLI Team, 2019). For instance, the State of California unanimously adopted legislation in August 2018 (ACR-215) endorsing the 23 Asilomar AI principles (FLI Team, 2018 and 2019).

AI's are already in clinical use in the U.S. In particular, AI shows great promise in the areas of diagnostics and imaging. In total, the FDA has already cleared or approved around 40 AI-based medical devices (Topol, 2019; The Medical Futurist, 2019). For example, in January 2017, *Arterys* received clearance from the U.S. Food and Drug Administration (FDA) for its medical imaging platform as the first ML application to be used in clinical practice (FDA, 2017a; Marr, 2017). It was initially cleared for cardiac magnetic resonance (MR) image analysis, but *Arterys* has meanwhile also received clearance from the FDA for other substantially equivalent devices (Arterys, 2019).

IDx-DR is the first FDA authorized AI diagnostic system that provides an autonomous screening decision without the need for a human being to interpret the image or results additionally (FDA, 2018a; IDx Technologies Inc., 2018). In April 2018, the FDA permitted marketing of this AI-based device to detect more than a mild level of the eye condition diabetic

retinopathy in adult patients (ages 22 and older) diagnosed with diabetes (FDA, 2018a and 2018b). The physician uploads the images of the patient's retinas to a cloud server, and the IDx-DR software then provides the physician with the recommendation either to rescreen in 12 months or to refer the patient to an eye specialist when more than mild diabetic retinopathy is detected (FDA, 2018a).

In May 2018, the FDA also granted marketing authorization for Imagen's software *OsteoDetect* for helping clinicians in detecting a common type of wrist fracture, called distal radius fracture, in adult patients (FDA, 2018c and 2018d). *OsteoDetect* uses ML techniques to analyze two-dimensional X-ray images to identify and highlight this type of fracture (FDA, 2018c and 2018d).

2.2 Europe

The European Commission adopted its AI strategy for Europe in April 2018. In this Communication, the Commission (2018a, pp. 3, 13–16) launched a European initiative on AI that aims to, inter alia, ensure an appropriate ethical and legal framework, for example, by creating a European AI Alliance and developing AI ethics guidelines. The Commission (2018a, p. 6) also stresses in this Communication that the entire European Union (EU) should strive to increase the (public and private) investment in AI to at least € 20 billion by the end of 2020.

The European Commission's High-Level Expert Group on AI (AI HLEG) – which was appointed by the European Commission in June 2018 and is also the steering group for the European AI Alliance – published the Ethics Guidelines in April 2019. The Guidelines promote the slogan “Trustworthy AI” and contain seven key requirements that AI systems need to fulfill in order to be trustworthy: “(1) human agency and oversight, (2) technical robustness and safety, (3) privacy and data governance, (4) transparency, (5) diversity, non-discrimination and fairness,

(6) environmental and societal well-being and (7) accountability” (AI HLEG, 2019a, p. 2). For the purpose of its deliverables, the AI HLEG also published a document on the definition of AI (AI HLEG, 2019b). Further, in June 2019, the AI HLEG published another deliverable that provides “Policy and Investment Recommendations for Trustworthy AI” (AI HLEG, 2019c).

The European Commission (2018a, p. 17) encourages all EU Member States to develop a national AI strategy, and several States have already released one such as the United Kingdom (UK) (House of Lords, 2018; UK Government, 2018; UK Department of Health and Social Care, 2019) and Germany (German Federal Government, 2018). The European Commission (2018b) also agreed upon a coordinated plan on AI with EU Member States, Norway, and Switzerland in December 2018 to promote the development and use of AI in Europe. The overall goal of working together is to ensure that Europe becomes the world-leading region for the development and application of “cutting-edge, ethical and secure AI” (European Commission, 2018b).

Only recently, in February 2020, the European Commission released a White Paper on AI that contains a European approach to excellence and trust. At the same time, the Commission also published a Communication on a European strategy (2020b) for data and a Report on the liability implications and safety of AI, the Internet of Things (IoT), and robotics (2020c). The Commission’s White Paper, in particular, emphasizes that “Europe can combine its technological and industrial strengths with a high-quality digital infrastructure and a regulatory framework based on its fundamental values to become a global leader in innovation in the data economy and its applications” (European Commission, 2020a).

There are also already AI health applications in Europe, and more are in the pipeline. For example, *Ada* (2020) is an AI health app that assesses an individual’s symptoms and gives guidance (e.g., suggest to the user a visit to a doctor or to seek emergency care). *Ada* (2020) has

been CE-marked (Class I) in Europe – a basic requirement to putting a medical device on the market within Europe – and complies with the EU General Data Protection Regulation 2016/679 (GDPR).

In August 2018, researchers at DeepMind and Moorfields Eye Hospital in London, UK, published in *Nature Medicine* the study results of an AI system that can read eye scans and make referral recommendation, comprising more than 50 common diagnoses; the system was trained on 14,884 scans and showed a success rate of 94% (De Fauw et al., 2018). DeepMind’s health team has meanwhile transitioned to Google Health, and Moorfields Eye Hospital is “excited to work with Google Health on the next phase to further develop this AI system so it can be used by patients all around the world” (Moorfields Eye Hospital, 2019).

Another example is *Ultromics* (2019). The team at the University of Oxford “is dedicated to reducing misdiagnosis and enabling earlier prevention of cardiovascular disease” (Ultromics, 2019). Ultromics’s *EchoGo Pro*, for example, is an outcome-based AI system with CE-marking in Europe that predicts coronary artery disease at an early stage (Ultromics, 2019).

Corti (2019) is a software developed by a Danish company that leverages ML to help emergency dispatchers make decisions. Corti can detect out-of-hospital cardiac arrests (i.e., those that occur in the public or home) during emergency calls faster and more accurately than humans by listening in to calls and analyzing symptoms, the tone of voice, breathing patterns, and other metadata in real time (Corti, 2019; Maack, 2018; Vincent, 2018).

3 Ethical Challenges

As the prior section suggests, the use of AI in the clinical practice of health care has huge potential to transform it for the better, but it also raises ethical challenges we now address.

3.1 Informed Consent to Use

Health AI applications, such as imaging, diagnostics, and surgery, will transform the patient-clinician relationship. But how will the use of AI to assist with the care of patients interface with the principles of informed consent? This is a pressing question that has not received enough attention in the ethical debate, even though informed consent will be one of the most immediate challenges in integrating AI into clinical practice (there is a separate question about informed consent to train AI we will not focus on here; Cohen, Amarasingham, Shah, Xie, and Lo, 2014). There is a need to examine under what circumstances (if at all) the principles of informed consent should be deployed in the clinical AI space. To what extent do clinicians have a responsibility to educate the patient around the complexities of AI, including the form(s) of ML used by the system, the kind of data inputs, and the possibility of biases or other shortcomings in the data that is being used? Under what circumstances must a clinician notify the patient that AI is being used *at all*?

These questions are especially challenging to answer in cases where the AI operates using “black-box” algorithms, which may result from non-interpretable machine-learning techniques that are very difficult for clinicians to understand fully (Cohen, 2018a; Yu et al., p. 727). For instance, Corti’s algorithms are “black box” because even Corti’s inventor does not know how the software reaches its decisions to alert emergency dispatchers that someone has a cardiac arrest. This lack of knowledge might be worrisome for medical professionals (Vincent, 2018). To what extent, for example, does a clinician need to disclose that they cannot fully interpret the diagnosis/treatment recommendations by the AI? How much transparency is needed? How does this interface with the so-called “right to explanation” under the EU’s GDPR (discussed further below in Section 4.3.2)? What about cases where the patient may be reluctant

to allow the use of certain categories of data (e.g., genetic data and family history)? How can we properly balance the privacy of patients with the safety and effectiveness of AI?

AI health apps and chatbots are also increasingly being used, ranging from diet guidance to health assessments to the help to improve medication adherence and analysis of data collected by wearable sensors (UK Nuffield Council on Bioethics, 2018, pp. 3, 4). Such apps raise questions for bioethicists about user agreements and their relationship to informed consent. In contrast to the traditional informed consent process, a user agreement is a contract that an individual agrees to without a face-to-face dialog (Klugman, Dunn, Schwartz, and Cohen, 2018, p. 40). Most people do not take the time to understand user agreements, routinely ignoring them (Klugman et al., 2018, p. 40; Cohen and Pearlman, 2018). Moreover, frequent updates of the software make it even more difficult for individuals to follow what terms of service they have agreed to (Gerke, Minssen, Yu, and Cohen, 2019). What information should be given to individuals using such apps and chatbots? Do consumers sufficiently understand that the future use of the AI health app or chatbot may be conditional on accepting changes to the terms of use? How closely should user agreements resemble informed consent documents? What would an ethically responsible user agreement look like in this context? Tackling these questions are tricky, and they become even more difficult to answer when information from patient-facing AI health apps or chatbots is fed back into clinical decision-making.

3.2 Safety and Transparency

Safety is one of the biggest challenges for AI in health care. To use one well-publicized example, IBM Watson for Oncology (2020) uses AI algorithms to assess information from patients' medical records and help physicians explore cancer treatment options for their patients. However, it has recently come under criticism by reportedly giving "unsafe and incorrect"

recommendations for cancer treatments (Ross and Swetlitz, 2018; Brown, 2018). The problem seems to be in the training of Watson for Oncology: instead of using real patient data, the software was only trained with a few „synthetic” cancer cases, meaning they were devised by doctors at the Memorial Sloan Kettering (MSK) Cancer Center (Ross and Swetlitz, 2018). MSK has stated that errors only occurred as part of the system testing and thus no incorrect treatment recommendation has been given to a real patient (Ross and Swetlitz, 2018).

This real-life example has put the field in a negative light. It also shows that it is of uttermost importance that AIs are safe and effective. But how do we ensure that AIs keep their promises? To realize the potential of AI, stakeholders, particularly AI developers, need to make sure two key things: (1) the reliability and validity of the data sets, and (2) transparency.

First, the used data sets need to be reliable and valid. The slogan “garbage in, garbage out” applies to AI in this area. The better the training data (labeled data) is, the better the AI will perform (Figure Eight, 2020). In addition, the algorithms often need further refinement to generate accurate results. Another big issue is data sharing: In cases where the AI needs to be extremely confident (e.g., self-driving cars), vast amounts of data and thus more data sharing will be necessary (Figure Eight, 2020). However, there are also cases (e.g., a narrow sentiment AI based off text) where less data will be required (Figure Eight, 2020). In general, it always depends on the particular AI and its tasks how much data will be required.

Second, in the service of safety and patient confidence some amount of transparency must be ensured. While in an ideal world all data and the algorithms would be open for the public to examine, there may be some legitimate issues relating to protecting investment/intellectual property and also not increasing cybersecurity risk (discussed below in Sections 4.4 and 4.5). Third party or governmental auditing may represent a possible solution.

Moreover, AI developers should be sufficiently transparent, for example, about the kind of data used and any shortcomings of the software (e.g., data bias). We should learn our lessons from examples such as Watson for Oncology, where IBM kept Watson's unsafe and incorrect treatment recommendations secret for over a year. Finally, transparency creates trust among stakeholders, particularly clinicians and patients, which is the key to a successful implementation of AI in clinical practice.

The recommendations of more "black-box" systems raise particular concerns. It will be a challenge to determine how transparency can be achieved in this context. Even if one could streamline the model into a simpler mathematical relationship linking symptoms and diagnosis, that process might still have sophisticated transformations beyond the skills of clinicians (and especially patients) to understand. However, perhaps there is no need to open the "black box": It may be that at least in some cases positive results from randomized trials or other forms of testing will serve as a sufficient demonstration of the safety and effectiveness of AIs.

3.3 Algorithmic Fairness and Biases

AI has the capability to improve health care not only in high-income settings, but to democratize expertise, "globalize" health care and bring it to even remote areas (Wahl, Cossy-Gantner, Germann, and Schwalbe, 2018). However, any ML system or human-trained algorithm will only be as trustworthy, effective, and fair as the data that it is trained with. AI also bears a risk for biases and thus discrimination. It is therefore vital that AI makers are aware of this risk and minimize potential biases at every stage in the process of product development. In particular, they should consider the risk for biases when deciding (1) which ML technologies/procedures they want to use to train the algorithms, and (2) what data sets (including considering their quality and diversity) they want to use for the programming.

Several real-world examples have demonstrated that algorithms can exhibit biases that can result in injustice with regard to ethnic origins and skin color or gender (Cossins, 2018; Fefegha, 2018; Sharkey, 2018; Short, 2018; Obermeyer, Powers, Vogeli, and Mullainathan, 2019). Biases can also occur regarding other features such as age or disabilities. The explanations for such biases differ and may be multifaceted. They can, for example, result from the data sets themselves (which are not representative), from how data scientists and ML systems choose and analyze the data, from the context in which the AI is used (Price II, 2019), etc. In the health sector, where phenotype- and sometimes genotype-related information are involved, biased AI could, for instance, lead to false diagnoses and render treatments ineffective for some subpopulations and thus jeopardize their safety. For example, imagine an AI-based clinical decision support software that helps clinicians to find the best treatment for patients with skin cancer. However, the algorithm was predominantly trained on Caucasian patients. Thus, the AI software will likely give less accurate or even inaccurate recommendations for subpopulations for which the training data was underinclusive such as African American.

Some of these biases may be resolved due to increased data availability and attempts to better collect data from minority populations and better specify for which populations the algorithm is or is not appropriately used. However, a remaining problem is that a variety of algorithms are sophisticated and non-transparent. In addition, as we have seen in the policing context, some companies developing software will resist disclosure and claim trade secrecy in their work (Sharkey, 2018; Wexler, 2018). It may therefore likely be left to non-governmental organizations to collect the data and show the biases (Sharkey, 2018).

In cases of “black-box” algorithms, many scholars have argued that explainability is necessary when an AI makes health recommendations, especially also to detect biases (London,

2019). However, does this view really hold true? Some argue that what matters is *not* how the AI reaches its decision but that it is accurate, at least in terms of diagnosis (London, 2019). The safety and effectiveness of health AI applications that are “black boxes” could, for example, be demonstrated – similar to the handling of drugs – by positive results of randomized clinical trials.

A related problem has to do with where AI will be deployed. AI developed for top-notch experts in resource-rich settings will not necessarily recommend treatments that are accurate, safe, and fair in low-resource settings (Price II, 2019; Minssen, Gerke, Aboy, Price, and Cohen, 2020). One solution would be *not* to deploy the technology in such settings. But such a “solution” only exacerbates preexisting inequalities. More thought must be given to regulatory obligations and resource support to make sure that this technology does improve not only the lives of the people living in high-income countries but also of those people living in low- and middle-income countries.

3.4 Data Privacy

In July 2017, the UK Information Commissioner’s Office (ICO) ruled that the Royal Free NHS Foundation Trust was in breach of the UK Data Protection Act 1998 when it provided personal data of circa 1.6 million patients to Google DeepMind (ICO, 2017a and 2017b). The data sharing happened for the clinical safety testing of “Streams”, an app that aims to help with the diagnosis and detection for acute kidney injury (ICO, 2017a and 2017b). However, patients were not properly informed about the processing of their data as part of the test (ICO, 2017a and 2017b). Information Commissioner’s Elizabeth Denham correctly pointed out that “the price of innovation does not need to be the erosion of fundamental privacy rights” (ICO, 2017b).

Although the Streams app does not use AI, this real-life example has highlighted the potential for harm to privacy rights when developing technological solutions (House of Lords,

2018, p. 90). If patients and clinicians do not trust AIs, their successful integration into clinical practice will ultimately fail. It is fundamentally important to adequately inform patients about the processing of their data and foster an open dialog to promote trust. The lawsuit *Dinerstein v. Google* (2019) and *Project Nightingale* by Google and Ascension (Copeland, 2019) are recent case studies showing patient privacy concerns in the context of data sharing and the use of AI.

But what about the ownership of the data? The value of health data can reach up to billions of dollars, and some evidence suggests that the public is uncomfortable with companies or the government selling patient data for profit (House of Lords, 2018, pp. 88, 89). But there may be ways for patients to feel valued that do not involve ownership per se. For example, the Royal Free NHS Foundation Trust had made a deal with Google DeepMind to provide patient data for the testing of Streams in exchange for the Trust's free use of the app for five years (House of Lords, 2018, p. 89). Reciprocity does not necessarily require ownership, but those seeking to use patient data must show that they are adding value to the health of the very same patients whose data is being used (Cohen, 2018b).

Beyond the question of what is collected, it is imperative to protect patients against uses outside the doctor-patient relationship that might deleteriously affect patients, such as impacts on health or other insurance premiums, job opportunities, or even personal relationships (Gerke et al., 2019). Some of this will require strong antidiscrimination law – similar to regimes in place for genetic privacy (Roberts, Cohen, Deubert, and Lynch, 2017); but some AI health apps also raise new issues, such as those that share patient data not only with the doctor but also with family members and friends (Gerke et al., 2019). In contrast to the doctor who is subject to duties of confidentiality set out by governing statutes or case law, family members or friends will probably *not* have legally enforceable obligations of such kind (Gerke et al., 2019). Does this

need to be changed? Another sensitive issue is whether and, if so, under what circumstances patients have a right to withdraw their data. Can patients request the deletion of data that has already been analyzed in aggregate form (Gerke et al., 2019)?

4 Legal Challenges

Many of the ethical issues discussed above have legal solutions or ramifications; while there is nothing sacrosanct in our division between the two, we now shift to challenges we associate more directly with the legal system.

4.1 Safety and Effectiveness

As we discussed previously (Section 3.2), it is of uttermost importance that AIs are safe and effective. Stakeholders can contribute to a successful implementation of AI in clinical practice by making sure that the data sets are reliable and valid, perform software updates at regular intervals, and being transparent about their product, including shortcomings such as data biases. In addition, an adequate level of oversight is needed to ensure the safety and effectiveness of AI. How this plays out varies between the U.S. and Europe. So, how is AI regulated in the U.S. and Europe? How can AI makers bring their products to the U.S. and European markets? The initial step of the analysis as to whether AI products need to undergo review is whether such products are medical devices.

4.1.1 U.S.

Let us start with the legal regulation in the U.S.

4.1.1.1 Medical Devices

The FDA regulates medical devices in the U.S. A medical device is defined in Section 201(h) Sentence 1 of the U.S. Federal Food, Drug, and Cosmetic Act (FDCA) as “an instrument,

apparatus, implement, machine, contrivance, implant, in vitro reagent, or other similar or related article, including any component, part, or accessory, which is—

- (1) recognized in the official National Formulary, or the United States Pharmacopeia, or any supplement to them,
- (2) intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease, in man or other animals, or
- (3) intended to affect the structure or any function of the body of man or other animals, and

which does not achieve its primary intended purposes through chemical action within or on the body of man or other animals and which is not dependent upon being metabolized for the achievement of its primary intended purposes.”

For example, medical devices include simple tongue depressors, pacemakers with micro-chip technology and in vitro diagnostic products such as reagents and test kits (FDA, 2018e).

4.1.1.2 Medical and Certain Decision Support Software

The 21st Century Cures Act (Pub. L. No. 114–255) was signed into law by the former President, Barack Obama, on 13 December 2016. Initially, it was hoped by some that the FDA would start to regulate medical advisory tools such as Watson for Oncology fully (Ross and Swetlitz, 2017). Ross and Swetlitz (2017) reported, however, that IBM had a large team of lobbyists pushing for proposals to prevent regulatory hurdles facing health software. Indeed, on 29 November 2016 – a day before the U.S. House of Representatives passed the 21st Century Cures Act – the company expressed its strong support for the Act in a press release, emphasizing that it “will support health innovation and advance precision medicine in the United States” (IBM, 2016; Ross and Swetlitz, 2017). The 21st Century Cures Act (Sec. 3060) introduced an exemption in Section 520(o) of the FDCA for medical and certain decisions support software

that does not fulfill the device definition. Section 201(h) of the FDCA was also amended by adding a second sentence which explicitly states that software functions under Section 520(o) FDCA do not fall under the term “device.”

Software Functions under Section 520(o)(1)(A)–(D) of the FDCA

Section 520(o)(1)(A)–(D) of the FDCA contains the following four categories of software functions that shall generally *not* fall under the device definition in Section 201(h) of the FDCA:

- (1) The software function is intended “for administrative support of a health care facility” (including e.g., business analytics, appointment schedules, and laboratory workflows);
- (2) The software function is intended “for maintaining or encouraging a healthy lifestyle and is unrelated to the diagnosis, cure, mitigation, prevention, or treatment of a disease or condition”;
- (3) The software function is intended “to serve as electronic patient records” (and “is not intended to interpret or analyze patient records”); or
- (4) The software function is intended “for transferring, storing, converting formats, or displaying clinical laboratory test or other device data and results.”

The FDA has also published nonbinding *Guidance for Industry and Food and Drug Administration Staff on Changes to Existing Medical Software Policies Resulting from Section 3060 of the 21st Century Cures Act* (2019a) to provide clarification of the interpretation of Section 520(o)(1)(A)–(D) of the FDCA.

Software Functions under Section 520(o)(1)(E) of the FDCA

Section 520(o)(1)(E) of the FDCA exempts specific clinical decision support software from the device definition in Section 201(h) of the FDCA. In order to be generally exempt from the device definition, a software function must meet the following four criteria simultaneously:

- (1) The software function is *not* “intended to acquire, process, or analyze a medical image or a signal from an in vitro diagnostic device or a pattern or signal from a signal acquisition system”;
- (2) The software function is intended “for the purpose of (...) displaying, analyzing, or printing medical information about a patient or other medical information (such as peer-reviewed clinical studies and clinical practice guidelines);”
- (3) The software function is intended “for the purpose of (...) supporting or providing recommendations to a health care professional about prevention, diagnosis, or treatment of a disease or condition;”
- (4) The software function is intended “for the purpose of (...) enabling such health care professional to independently review the basis for such recommendations that such software presents so that it is not the intent that such health care professional rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient”

(Sec. 520(o)(1)(E) of the FDCA; FDA, 2019b, pp. 6, 7).

In September 2019, the FDA (2019b) issued *Draft Guidance for Industry and Food and Drug Administration Staff on Clinical Decision Support Software* that contains nonbinding recommendations on the interpretation of the criteria in Section 520(o)(1)(E) of the FDCA.

In particular, the FDA clarifies that the term “clinical decision support” (CDS) is defined broadly and means software functions that meet the first two criteria and part of the third

criterion (i.e., intended “for the purpose of (...) supporting or providing recommendations”) (FDA, 2019b, p. 8). A CDS function is only exempt from the device definition when the fourth criterion is additionally fulfilled (FDA, 2019b, pp. 8, 9). Thus, it is decisive to determine whether the software function enable the “health care professional to independently review the basis for such recommendations that such software presents”. The FDA clarifies in its draft guidance that “the software developer should describe the underlying data used to develop the algorithm and should include plain language descriptions of the logic or rationale used by an algorithm to render a recommendation. The sources supporting the recommendation or the sources underlying the basis for the recommendation should be identified and available to the intended user (e.g., clinical practice guidelines with the date or version, published literature, or information that has been communicated by the CDS developer to the intended user) and understandable by the intended user (e.g., data points whose meaning is well understood by the intended user)” (FDA, 2019b, p. 12). The FDA also states that health care professionals rely primarily on software recommendations – and thus are unable “to independently review the basis for such recommendations” – if they cannot be expected to independently understand the meaning of the information on which the recommendations are based (FDA, 2019b, p. 12). An example includes when inputs that are used to generate a recommendation are not described (FDA, 2019b, p. 12).

The FDA also makes clear that it does not intend at this time to enforce compliance with applicable regulatory requirements with respect to certain software functions that are intended for health care professionals, caregivers or patients and may meet the device definition but are low risk (FDA, 2019b, pp. 9, 16-18). For example, even if the fourth criterion is *not* fulfilled and health care professionals rely primarily on software recommendations, the FDA does not intend at this time to enforce compliance with the relevant device requirements as long as the device

CDS functions inform clinical management for non-serious health care situations or conditions (FDA, 2019b, p. 16). The agency thus focuses its oversight especially on device CDS software functions that inform clinical management for serious or critical health care conditions or situations (FDA, 2019b, p. 17). The FDA also clarifies in its draft guidance that it also intends to focus its regulatory oversight on software functions that are devices but are not classified as CDS (FDA, 2019b, pp. 24-27).

4.1.1.3 Other FDA Initiatives

There are many other important FDA initiatives we cannot do justice here, including its *Guidance on Software as a Medical Device (SAMd): Clinical Evaluation* (2017b) and the launch of the so-called “*Software Pre-Cert Pilot Program*.” The latter will enable some digital health developers to become pre-certified based on excellence in identified criteria (e.g., patient safety, clinical responsibility, and product quality) and bring their lower-risk software-based medical devices with more streamlined FDA review to market or no review at all (FDA, 2017c, p. 5–7; FDA, 2019c). The FDA also published a Working Model that contains suggestions for the main components of the Pre-Cert Pilot Program (FDA, 2019c and 2019d). Although there are still many open questions, the program is an innovative regulatory experiment that may hold lessons for peer countries and should be closely followed.

In particular, the FDA (2019e) has only recently, in April 2019, proposed a regulatory framework for public comment for *Modifications to Artificial Intelligence/Machine Learning (AI/ML) - Based Software as a Medical Device (SaMD)*. SaMD is “software intended to be used for one or more medical purposes that perform these purposes without being part of a hardware medical device” (IMDRF, 2013, p. 6). The FDA’s discussion paper proposes a new, total product lifecycle (TPLC) regulatory approach for AI/ML-based SaMD that are *medical devices* to allow

those devices to adapt and optimize their performance in real-time to continuously improve while ensuring their safety and effectiveness (FDA, 2019e). While we praise the FDA's efforts in the field, it will be essential for regulators to focus especially on the development of a process to continuously monitor, manage, and identify risks due to features that are closely tied to AI/ML systems' reliability (e.g., concept drift, instability, and covariate shift) (Babic, Gerke, Evgeniou, and Cohen, 2019). Moreover, when there is substantial human involvement in decision-making, it becomes even more challenging for regulators to determine the effects of the update of such devices (Gerke, Babic, Evgeniou, and Cohen, 2020a).

4.1.2 Europe

Let us now shift to the legal particularities in Europe.

4.1.2.1 Medical Devices and New Legal Developments

There are also new legal developments in the EU: Two new EU Regulation entered into force on 25 May 2017, namely the *Medical Device Regulation* (2017/745 – MDR; see Art. 123 (1) of the MDR) and the *Regulation on in vitro diagnostic medical devices* (2017/746 – IVDR; see Art. 113(1) of the IVDR). With some exceptions, the MDR was supposed to become effective on 26 May 2020 (Art. 123 (2) and (3) of the MDR). However, due to the need for medical devices to combat COVID-19, the European Parliament (2020) postponed the MDR's application by one year (i.e., 26 May 2021). The MDR will repeal the *Medical Device Directive* (93/42/EEC – MDD) and the *Directive on active implantable medical devices* (90/385/EEC – AIMD) (Art. 122 of the MDR). The IVDR will become effective as planned on 26 May 2022 (Art. 113(2) and (3) of the IVDR), thereby especially repealing the *Directive on in vitro diagnostic medical devices* (98/79/EC – IVDD) (Art. 112 of the IVDR).

4.1.2.2 MDR

The new MDR will bring some changes in the classification process of medical devices. Software that does not fall under the medical device definition of the MDD may soon be classified as a medical device under the MDR. In particular, the new medical device definition in Art. 2(1) of the MDR also considers software that is used for human beings for the *medical purpose of prediction or prognosis* of disease as a medical device. However, the MDR also explicitly clarifies in Recital 19 that “software for general purposes, even when used in a healthcare setting, or software intended for life-style and well-being purposes is not a medical device.”

Similar to the MDD, medical devices under the MDR will be classified into four categories, namely Classes I, IIa, IIb, and III, based on the intended purpose of the medical devices and their inherent risks (Art. 51(1) of the MDR).

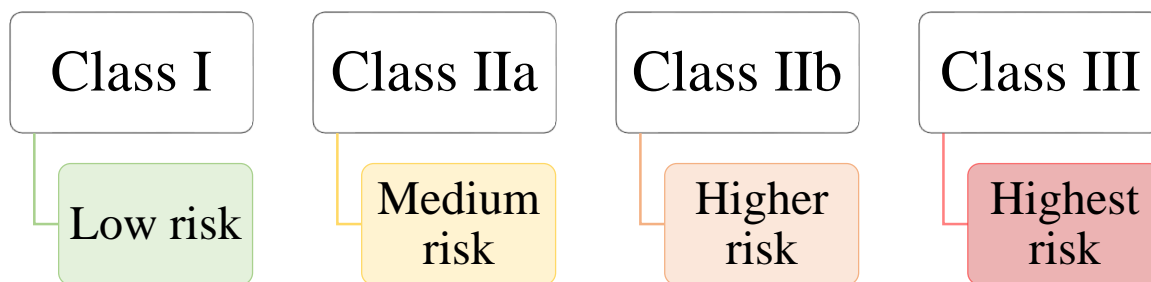


Fig. 1: Classification of Medical Devices

The MDR also introduces new implementing and classification Rules for software in Chapters II and III of Annex VIII. In particular, the MDR contains a new classification Rule that focuses explicitly on software. According to this Rule, “software intended to provide information which is used to take decisions with diagnosis or therapeutic purposes is classified as class IIa, except if such decisions have an impact that may cause:

- death or an irreversible deterioration of a person’s state of health, in which case it is in class III; or
- a serious deterioration of a person’s state of health or a surgical intervention, in which case it is classified as class IIb.

Software intended to monitor physiological processes is classified as class IIa, except if it is intended for monitoring of vital physiological parameters, where the nature of variations of those parameters is such that it could result in immediate danger to the patient, in which case it is classified as class IIb.

All other software is classified as class I”

(Rule 11 in Chapter III of Annex VIII of the MDR).

This new Rule will also lead to reclassifications, meaning software that was originally classified as a medical device under the MDD may be classified in another class category under the MDR. For example, CDS software such as Watson for Oncology will probably be at least classified as a class IIa medical device under the MDR since it “provide(s) information which is used to take decisions with diagnosis or therapeutic purposes.” Depending on the decision’s impact, the AI-based CDS software could even be classified as a class III (if it may cause “death or an irreversible deterioration of a person’s state of health”) or class IIb device (if it may cause “a serious deterioration of a person’s state of health or a surgical intervention”) (Rule 11 in Chapter III of Annex VIII of the MDR). In October 2019, the Medical Device Coordination Group (MDCG) also released non-binding guidance on qualification and classification of software under the MDR and IVDR (MDCG, 2019).

A CE marking (similar to the current MDD) will especially indicate the conformity with the applicable requirements set out in the MDR so that a medical device can move freely within

the EU and be put into service in accordance with its intended purpose (Recital 40 and Art. 2(43) of the MDR). In particular, manufacturers of medical devices shall undertake an assessment of the conformity of their devices prior to placing them on the market (Art. 52 and Annexes IX to XI of the MDR). The applicable conformity assessment procedure is based on the classification (Class I, IIa, IIb, or III) and type (e.g., implantable) of the particular device (Art. 52 of the MDR). For example, Class I devices have a low level of vulnerability and thus the conformity assessment procedure can generally be carried out under the sole responsibility of the manufacturers (Recital 60 and Art. 52(7) of the MDR). In contrast, Class IIa, IIb and III devices that have a higher risk than Class I devices entail the involvement of a notified body, a conformity assessment body designated in accordance with the MDR (Recital 60 and Art. 2(42) of the MDR).

4.2 Liability

New AI-based technologies also raise challenges for current liability regimes. It will be crucial to creating an optimal liability design that figures out responsibilities.

4.2.1 U.S.

Imagine the following case: An AI-based CDS software (see Sec. 4.1.1.2) gives an incorrect treatment recommendation (in the sense that it is not one a non-AI clinician would have arrived at) that the clinician adopts resulting in harm to the patient. In this situation, the clinician would likely be liable for *medical malpractice*. Clinicians must treat patients with due expertise and care; they need to provide a standard of care that is expected of relevant members of the profession. At present, it appears that clinicians could thus be held liable even though they engaged in good faith reliance on a “black-box” ML algorithm because AI-based CDS software is considered a tool under the control of the health professional who makes the ultimate decision;

she remains the captain of the ship and thus responsible for its course. But should that result obtain in a case where the software function does not enable the “health care professional to independently review the basis for such recommendations that such software presents” (see Sec. 520(o)(1)(E)(iii) of the FDCA)? In the other direction, could we imagine a future where the use of AI-based technology becomes the standard of care, and thus the choice *not* to use such technology would subject the clinician to liability (Froomkin, Kerr, and Pineau, 2019)? At the moment, however, using advanced AI does not (yet) appear to be part of the standard of care. Thus, to avoid medical malpractice liability, physicians *can* use it as a confirmatory tool to assist with existing decision-making processes as opposed to *needing* to follow its recommendations out of fear of liability. (Price, Gerke, and Cohen, 2019).

Setting the optimal liability regime depends heavily on what one thinks the “problem” is. If one is concerned that the deployment of AI-based technology in the clinical space is associated with a high risk for patients to get hurt, one might want to keep the current medical malpractice regime that attempts to meet both of tort law’s two functions: (1) deterrence and (2) compensation of the victims. By contrast, if one believes that over the run of cases, reliance on AI promotes patient health, then it may be a problem if physicians prove reluctant to rely on these algorithms, especially the more opaque ones, when they remain on the hook for resulting liability (see also Price II, 2017, p. 12). This might drive the policymaker to a different model.

Some have proposed *product liability* against the makers of AI, a tort that generally entails a strict liability of the manufacturer for defects. However, there are considerable challenges to win such a claim in practice. Courts have hesitated to apply or extend product liability theories to health care software developers since such software is currently primarily considered as a tool to support clinicians make the final decision (Price II, 2017, pp. 11, 12).

A different approach would be to focus on *compensation, even without deterrence*. An example of such a system in the U.S. is vaccine *compensation*. Vaccine manufacturers pay into a fund, and the system collectivizes the risk by paying out to those that are harmed by vaccines. AI manufacturers could do the same, which would compensate patients and spread the risks across the industry, but may give individual makers of AI less incentives to ensure the product's safety.

Beyond clinicians and AI makers, one must also consider the liability of the hospitals that purchase and implement the AI systems. Lawsuits might be brought against them under the theories of *corporate negligence and vicarious liability*. One interesting theory for hospital liability is “negligent credentialing”—just as hospitals may be liable if they do not adequately review the credential and practice of physicians and other staff they employ (Bezaire, Felton, Greve, and Allen, 2017), they may have similar duties when they “hire” an AI.

Still another possibility would be to pair a liability shield with a more rigorous *pre-approval scheme* that would immunize health care professionals and manufacturers from some forms of liability because of the approval process. Whether this is desirable depends in part on one's view of litigation versus administrative law regimes: is ex ante approval by a regulator preferable to ex post liability at the hands of a judge or jury?

4.2.2 Europe

Europe is also not (yet) ready for the new liability challenges that AI-based technology will bring along with it. There is currently no fully harmonized EU regulatory framework for liability on AI and robotics such as care and medical robots in place. However, Europe has taken several steps to address the issue of liability.

One first step in the right direction was the publication of a resolution by the European Parliament called *Civil Law Rules on Robotics: European Parliament resolution of 16 February*

2017 with recommendations to the Commission on Civil Law Rules on Robotics

(2015/2103(INL)). This resolution, among other things, questions whether the current liability rules are sufficient and whether new rules are required “to provide clarity on the legal liability of various actors concerning responsibility for the acts and omissions of robots” (Sec. AB). It also points out that the current scope of *Council Directive concerning liability for defective products (85/374/EEC – Product Liability Directive)* may not adequately cover the new developments in robotics (Sec. AH). The resolution emphasizes “that the civil liability for damage caused by robots is a crucial issue which also needs to be analysed and addressed at Union level in order to ensure the same degree of efficiency, transparency and consistency in the implementation of legal certainty throughout the European Union for the benefit of citizens, consumers and businesses alike” (Sec. 49). It thus asks the European Commission for “a proposal for a legislative instrument on legal questions related to the development and use of robotics and AI foreseeable in the next 10 to 15 years, combined with non-legislative instruments such as guidelines and codes of conduct” (Sec. 51). The resolution recommends that the European Commission should define in this legislative instrument which of the two approaches should be applied: either *strict liability* (i.e., which “requires only proof that damage has occurred and the establishment of a causal link between the harmful functioning of the robot and the damage suffered by the injured party”) or the *risk management approach* (i.e., which “does not focus on the person ‘who acted negligently’ as individually liable but on the person who is able, under certain circumstances, to minimise risks and deal with negative impacts”) (Sec. 53–55 and Annex to the resolution). It also recommends an obligatory insurance scheme and an additional compensation fund to ensure that damages will be paid out in situations where no insurance cover exists (Annex to the resolution).

As a second step, in April 2018, the European Commission adopted its AI strategy (see Section 2.2). A first mapping of liability challenges for emerging digital technologies, such as AI, advanced robotics and the IoT, was provided in a *Commission Staff Working Document on Liability for Emerging Digital Technologies* also published in April 2018 together with the AI strategy (European Commission, 2018c).

Further, in November 2019, the independent Expert Group on Liability and New Technologies – New Technologies Formation (NTF) that was set up by the European Commission released a report on liability for AI and other emerging digital technologies such as IoT (NTF, 2019). The NTF's findings include that liability regimes are mainly regulated by the EU Member States except for strict liability of producers for defective products that is regulated by the Product Liability Directive at EU level (NTF, 2019, p. 3). The NTF's opinion is that the Member States' liability regimes are a good starting point for new technologies and provide at least basic protection of victims (NTF, 2019, p. 3). However, the NTF also identifies several points in its report that need to be changed at national and EU levels (NTF, 2019, p. 3). For example, the NTF emphasizes that “a person operating a permissible technology that nevertheless carries an increased risk of harm to others, for example AI-driven robots in public spaces, should be subject to strict liability for damage resulting from its operation” (NTF, 2019, p. 3). It also states, for instance, that “a person using a technology which has a certain degree of autonomy should not be less accountable for ensuing harm than if said harm had been caused by a human auxiliary” (NTF, 2019, p. 3).

Only recently, in February 2020, the European Commission also published a report on the safety and liability implications of AI, the IoT and robotics (European Commission, 2020c). The Commission understands the importance of these technologies and aims to make “Europe a

world-leader in AI, IoT and robotics” (European Commission, 2020c, p. 1). To achieve this aim, the Commission states that “a clear and predictable legal framework addressing the technological challenges is required” (European Commission, 2020c, p. 1). The Commission, in accordance with the NTF, argues that “in principle the existing Union and national liability laws are able to cope with emerging technologies” (European Commission, 2020c, p. 17). However, it also identifies some challenges raised by new digital technologies such as AI that need to be addressed by adjustments in the current national and EU regulatory frameworks such as the Product Liability Directive (European Commission, 2020c, pp. 16, 17).

We welcome the European Commission efforts to identify and address the liability issues raised by AI and other emerging digital technologies. As a next consequent step, changes need to be made at national and EU levels to implement the NTF’s and European Commission’s findings. Such updates of the liability frameworks should be carried out as soon as possible to have provisions in place that adequately deal with these new technological developments. Updated frameworks are needed to create clarity, transparency, and public trust.

4.3 Data Protection and Privacy

In the world of big data, it is of pivotal importance that there are data protection laws in place that adequately protects the privacy of individuals, especially patients. In the following, we will give an overview of relevant provisions and legal developments on data protection and privacy in the U.S. and Europe.

4.3.1 U.S.

The Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule (45 C.F.R. Part 160 as well as subparts A and E of Part 164) is the key federal law to protect health data privacy (Price II and Cohen, 2019, p. 38). However, HIPAA has significant gaps when it

comes to today's health care environment since it only covers specific health information generated by "covered entities" or their "business associates." HIPAA does not apply to non-health information that supports inferences about health such as a purchase of a pregnancy test on Amazon (Cohen and Mello, 2018, p. 232; Price II and Cohen, 2019, p. 39). Moreover, the definition of "covered entities" also limits its scope; it generally includes insurance companies, insurance services, insurance organizations, health care clearinghouses, and health care providers (45 C.F.R. §§ 160.102, 160.103), but not much beyond that (Cohen and Mello, 2018, p. 231; Price II and Cohen, 2019, p. 39). In particular, much of the health information collected by technology giants such as Amazon, Google, IBM, Facebook, and Apple that are all investing heavily in the field of AI in health care, and are not "covered entities," will fall outside of HIPAA (Price II and Cohen, 2019, p. 39). HIPAA also does not apply in cases of user-generated health information (Cohen and Mello, 2018, p. 232; Price II and Cohen, 2019, p. 39). For example, a Facebook post about a disease falls outside of HIPAA's regime (Cohen and Mello, 2018, p. 232).

A different problem with HIPAA is its reliance on de-identification as a privacy strategy. Under HIPAA de-identified health information can be shared freely for research and commercial purposes (Cohen and Mello, 2018, p. 231; 45 C.F.R. § 164.502(d)(2)). It provides two options for de-identification: (1) a determination by someone with appropriate knowledge of and experience with usually accepted scientific and statistical methods and principles; or (2) the removal of 18 identifiers (e.g., names, social security numbers, and biometric identifiers) of the individual or of relatives, household members, or employers of the individual, and no actual knowledge of the covered entity that the information could be used to identify an individual (45 C.F.R. § 164.514(b)). But this may not adequately protect patients because of the possibility of

data triangulation – to re-identify data thought to be de-identified under the statute through the combination of multiple data sets (Price II and Cohen, 2019, pp. 39, 40; Gerke, Yeung, and Cohen, 2020b). The problem of data triangulation has also recently been featured in a lawsuit, *Dinerstein v. Google* (2019), in which the plaintiffs alleged that the defendants shared medical records with Google containing enough information that enabled Google to potentially re-identify patients given all of its other data at hand.

For all these reasons, HIPAA is not adequate to protect the health privacy of patients. It is time for federal law to take seriously the protection of health-relevant data that is not covered by HIPAA (Cohen and Mello, 2018, p. 232; Arney et al., 2019, pp. 9, 16). Such a federal law should facilitate both innovations, including health AI applications, and adequate protection of health privacy of individuals.

While HIPAA preempts less protective state law, it does not preempt states whose laws are more protective. Inspired by the EU GDPR, California recently has taken action at the state level: The California Consumer Privacy Act of 2018 (CCPA) became effective on 1 January 2020 (Cal. Civ. Code § 1798.198). The CCPA grants various rights to California residents with regard to personal information that is held by businesses. The *term* “*business*” is defined in Section 1798.140(c) of the California Civil Code and applies to “a sole proprietorship, partnership, limited liability company, corporation, association, or other legal entity that is organized or operated for the profit or financial benefit of its shareholders or other owners that collects consumers’ personal information or on the behalf of which that information is collected and that alone, or jointly with others, determines the purposes and means of the processing of consumers’ personal information, that does business in the State of California, and that satisfies one or more of the following thresholds:

- (A) Has annual gross revenues in excess of twenty-five million dollars (...).
- (B) Alone or in combination, annually buys, receives for the business's commercial purposes, sells, or shares for commercial purposes, alone or in combination, the personal information of 50,000 or more consumers, households, or devices.
- (C) Derives 50 percent or more of its annual revenues from selling consumers' personal information."

The CCPA defines the term "*personal information*" broadly as "information that identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household", including a real name, alias, postal address, social security number, and biometric information (Cal. Civ. Code § 1798.140(o)(1)). In particular, personal information is *not* "publicly available information" – "information that is lawfully made available from federal, state, or local government records" (Cal. Civ. Code § 1798.140(o)(2)).

The CCPA does not apply to protected health information that is collected by HIPAA covered entities or their business associates (Cal. Civ. Code § 1798.145(c)(1)). However, it applies to a great deal of information in so-called "shadow health records" – health data that is collected outside of the health system (Price II, Kaminski, Minssen, Spector-Bagdady, 2019, p. 449). Thus, the CCPA is a welcome attempt to at least partially fill in legal gaps and improve the data protection of individuals.

4.3.2 Europe

The General Data Protection Regulation (GDPR – 2016/679) has been applied since 25 May 2018 (Art. 99(2) of the GDPR) in all EU Member States and introduced a new era of data protection law in the EU.

The GDPR particularly aims to protect the right of natural persons to the protection of personal data (Art. 1(2) of the GDPR). It applies to the “processing of personal data in the context of the activities of an establishment of a controller or a processor” *in the EU*, notwithstanding of whether the processing takes place in an EU or non-EU country, such as in the U.S. (Arts. 2, 3(1) of the GDPR). In addition, the GDPR may also have implications for U.S. companies. For example, the Regulation applies in cases where the processor or controller is *established in a non-EU country* and processes “personal data of data subjects who are in the Union” for “the offering of goods or services” (e.g., newspapers and affiliated websites for free or for a fee) to such data subject in the EU or for the “monitoring” of the data subjects’ behavior (Art. 3(2) of the GDPR; Gerke, 2018). The GDPR also applies where a controller processes personal data and is *established in a non-EU country*, but “in a place where Member State law applies by virtue of public international law” (Art. 3(3) of the GDPR).

Art. 3(1)	Art. 3(2)	Art. 3(3)
Processing of personal data	Processing of personal data of data subjects who are in the EU	Processing of personal data
In the context of the activities of an EU establishment of a controller or a processor	Non-EU establishment of a controller or a processor	Non-EU establishment of a controller
Processing takes place within or outside the EU	The processing activities are related to: (a) the offering of goods or services (paid or for free) to such data subjects in the EU; or (b) the monitoring of the data subjects’ behavior as far as their behavior takes place within the EU	But in a place where Member State law applies by virtue of public international law

Table 2: GDPR’s Territorial Scope

The term “personal data” is defined as “any information relating to an identified or identifiable natural person (‘data subject’)” (Art. 4(1) of the GDPR). The GDPR defines “processing” as “any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means”, including collection, structuring, storage, or use (Art. 4(2) of the GDPR). Whereas a “controller” is “the natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data”, a “processor” means “a natural or legal person, public authority, agency or other body which processes personal data on behalf of the controller” (Art. 4(7), (8) of the GDPR).

In the health care context, the definition of “data concerning health” under Art. 4(15) of the GDPR is, in particular, relevant: “personal data related to the physical or mental health of a natural person, including the provision of health care services, which reveal information about his or her health status.” The EU’s GDPR is thus a lot broader in its scope compared to US’ HIPAA, which only covers specific health information generated by “covered entities” or their “business associates” (discussed in Section 4.3.1).

According to Article 9(1) of the GDPR, the processing of special categories of personal data such as genetic data (Art. 4(13) of the GDPR), biometric data (Art. 4(14) of the GDPR), and data concerning health is prohibited. But Article 9(2) of the GDPR contains a list of exceptions to paragraph 1 (Gerke, 2018). For example, the prohibition in Article 9(1) of the GDPR shall usually not apply in cases where “the data subject has given explicit consent (...) for one or more specified purposes” or where the “processing is necessary for reasons of public interest in the area of public health” or “for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes” (Art. 9(2)(a), (i), and (j) of the GDPR; Gerke, 2018).

The EU Member States can also decide to introduce or maintain further requirements, including limitations, but only “with regard to the processing of genetic data, biometric data or data concerning health” (Art. 9(4) of the GDPR).

Non-compliance with these GDPR’s conditions shall result in administrative fines up to 20 million EUR or – if higher – up to four percent of an undertaking’s annual global turnover of the previous year (Art. 83(5) of the GDPR). The first fines in the health care context have already been imposed under the GDPR. For example, a hospital in Portugal was charged 400 thousand EUR for two breaches of the GDPR: First, 300 thousand EUR for the permit of “indiscriminate access to a set of data by professionals, who should only be able to access them in specific cases”; and second, 100 thousand EUR for the incapacity to “ensure the confidentiality, integrity, availability and permanent resilience of treatment systems and services” (Lusa, 2018).

The GDPR also contains provisions that are especially relevant to AI-infused medicine. For example, where personal data are collected, the controllers must generally provide data subjects with *information* about “the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject” (Arts. 13(2)(f), 14(2)(g) of the GDPR). In addition, data subjects have the *right of access* to the personal data concerning them that are being processed and the information about “the existence of automated decision-making, including profiling, (...) and (...) meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject” (Art. 15(1)(h) of the GDPR).

“Automated decision-making” means a decision that is made – without any human involvement – solely by automated means (Working Party, 2018, p. 20). The term “profiling” is defined in Article 4(4) of the GDPR as “any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person’s performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements.” Thus, the term “profiling” is a subset of the term “processing” with two additional requirements, namely the processing must be (1) automated, and (2) for evaluation purposes (Goodman and Flaxman, 2017, p. 52).

Under Article 22(1) of the GDPR, data subjects shall also “have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.” Article 22(2) of the GDPR lists some exceptions to Article 22(1) of the GDPR, but these exceptions do generally not apply where decisions are based on genetic and biometric data as well data concerning health (Art. 22(4) of the GDPR).

It is highly controversial, however, whether the GDPR actually grants a “right to explanation” and what such a right means (Burt, 2017; Goodman and Flaxman, 2017; Kaminski, 2018; Wachter, Mittelstadt, and Floridi, 2017). Recital 71 of the GDPR explicitly mentions “the right (...) to obtain an explanation of the decision reached after such assessment.” Some scholars doubt the legal existence and the feasibility of such a right to explanation of *specific automated decisions*, inter alia, because Recital 71 of the GDPR is not legally binding, and a right to explanation is not mandated by the legally binding requirements set out in Article 22(3) of the GDPR. (Wachter et al., 2017). Thus, according to this view, there is from the outset *no* legally

binding right of the data subject to receive insight into the internal decision-making process of algorithms (Wachter, Mittelstadt, and Russell, 2018), and thus to open the “black boxes” of health AI applications. However, if a legally binding right to explanation of specific automated decisions does *not* exist, Articles 13(2)(f), 14(2)(g), 15(1)(h) of the GDPR at least entitle data subjects to obtain “meaningful information about the logic involved, as well as the significance and the envisaged consequences” of automated decision-making systems (Wachter et al., 2017). This information includes the purpose of an automated decision-making system, how the system works in general, the predicted impact as well as *other system functionality* such as decision trees and classification structures (Wachter et al., 2017).

It is also likely that companies that are controllers under the GDPR must carry out a data protection impact assessment for new AI-based technologies that shall be deployed in the clinical space. In general, Article 35(1) of the GDPR requires such an assessment, prior to the processing, for “new technologies” where the processing “is likely to result in a high risk to the rights and freedoms of natural persons.” Article 35(3) of the GDPR explicitly states when a data protection impact assessment shall especially be required such as in cases of “a systematic and extensive evaluation of personal aspects relating to natural persons which is based on automated processing, including profiling, and on which decisions are based that produce legal effects concerning the natural person or similarly significantly affect the natural person” or “processing on a large scale of special categories of data” (e.g., genetic data and data concerning health). Recital 91 of the GDPR clarifies that personal data should *not* be considered “on a large scale if the processing concerns personal data from patients (...) by an individual physician.” Article 35(7) of the GDPR contains a list of what the assessment shall at least include, such as a

description of the envisaged processing operations, an assessment of the risks to the freedoms and rights of data subjects, and the measures envisaged to address the risks.

As complementation to the GDPR, the Regulation (EU) 2018/1807 entered into force in December 2018 and has been directly applicable since 28 May 2019 (Art. 9 of Regulation 2018/1807). This Regulation contains a framework for the free flow of *non-personal data* in the EU by laying down rules to the availability of data to competent authorities, data localization requirements, and the porting of data for professional users (Art. 1 of Regulation 2018/1807). It applies to the *processing of electronic data* (other than personal data as defined in Art. 4(1) of the GDPR) in the EU, which is either “provided as a service to users residing or having an establishment in the Union”, irrespective of whether the service provider is established in an EU or non-EU country or “carried out by a natural or legal person residing or having an establishment in the Union for its own needs” (Arts. 2(1), 3(1) and (2) of Regulation 2018/1807). In cases of data sets composed of personal and non-personal data, the Regulation (EU) 2018/1807 does also apply to the non-personal data part of such data sets (Art. 2(2) of Regulation 2018/1807). However, the GDPR applies in cases where the personal and non-personal data in data sets are inextricably linked (Art. 3(2) of Regulation 2018/1807).

4.4 Cybersecurity

Cybersecurity is another important issue we need to consider when addressing legal challenges to the use of AI in health care. In the future, much of the health care related services, processes, and products will operate within the IoT. Unfortunately, much of the underlying infrastructure is vulnerable to both cyber and physical threats and hazards (U.S. Department of Homeland Security, 2019). For example, sophisticated cyber actors, criminals, and nation-states can exploit vulnerabilities to steal or influence the flow of money or essential (health care)

information (U.S. Department of Homeland Security, 2019). Such actors are increasingly developing skills to threaten, harm, or disrupt the delivery of vital (medical) services (U.S. Department of Homeland Security, 2019). Targets in the health sector may include hospital servers, diagnostic tools, wearables, wireless smart pills, and medical devices (Pinsent Masons, 2017). All can be infected with software viruses, Trojan horses, or worms that risk patients' privacy and health (Gerke et al., 2019). Moreover, corrupted data or infected algorithms can lead to incorrect and unsafe treatment recommendations (Gerke et al., 2019). Hostile actors could get access to sensitive data such as health information on patients or could threaten patients' safety by misrepresenting their health. AIs are, in particular, vulnerable to manipulation (Gerke, Kramer, Cohen, 2019). For example, Finlayson et al. (2019) have shown in a recent publication that the system's output can completely be changed so that it classifies a mole as malignant with 100% confidence by making a small change in how inputs are presented to the system (Gerke, Kramer, Cohen, 2019).

The need for increased cybersecurity was shown in the "WannaCry" ransomware attack, a global cyber attack using sophisticated hacking tools that crippled the National Health Service (NHS) in the UK, hit the international courier delivery services company FedEx and infected more than 300,000 computers in 150 countries (Graham, 2017). Events like these not only resulted in reactions at national level such as in the UK (Smart, 2018) but also prompted a *new Cybersecurity Act* (Regulation (EU) 2019/881) that came into force on 28 June 2019.

The new Cybersecurity Act's goals are to achieve a high level of cyber resilience, cybersecurity, and trust in the EU while ensuring the internal market's proper functioning (Art. 1(1)). In particular, it lays down a *European cybersecurity certification framework* to ensure that certified information and communications technology (ICT) products, ICT services,

and ICT processes in the EU fulfill an adequate level of cybersecurity (Art. 1(1)(b)). The Act also lays down the tasks, objectives, and organizational matter relating to the European Union Agency for Cybersecurity (ENISA) (Art. (1)(a)).

There is also new progress in the U.S.: The Cybersecurity and Infrastructure Security Act of 2018 (H.R.3359) was signed into law by President Donald Trump on 16 November 2018 (U.S. Department of Homeland Security, 2019). This Act (Sec. 2) amended the Homeland Security Act of 2002, and, in particular, redesignated the National Protection and Programs Directorate of the Department of Homeland Security as the Cybersecurity and Infrastructure Security Agency (CISA) (Sec. 2202; 6 U.S.C. 652; Congress, 2018). CISA augments the U.S. national capacity to defend against cyber attacks and will help the federal government provide cybersecurity tools, assessment skills, and incident response services to safeguard sensitive networks (U.S. Department of Homeland Security, 2019).

While the latest legal developments in the U.S. and Europe will hopefully promote the safety of AI-driven products, services, and processes in the health care sector, cyber-attacks are often a global issue; data sharing and breaches frequently do not stop at the U.S. or European borders but occur around the world (Gerke et al., 2019). Thus, there is the need for an internationally enforceable, large-scale regulatory framework on cybersecurity that ensures a high level of cybersecurity and resilience across borders (Gerke et al., 2019). It will not be easy to set up such a framework since it will require to properly balance the different interests of all stakeholders involved. (Gerke et al, 2019).

4.5 Intellectual Property Law

Translating AI and big data into safe and effective “real world” products, services, and processes is an expensive and risky venture. As a result, the commercial protection of AI and

data-driven health care/life science technologies have become an exceedingly important topic (The Economist, 2017; Carrier, 2011, Katz and Shapiro, 1992; Lemley and Shafir, 2011). At the same time, there are continuing discussions about open science and innovation and the primary objective of more data sharing as well as increasing debates over access to such technologies and the pertinent data (Carrier, 2011; Lemley and Shafir, 2011).

AI – and the data that fuels it – can be protected by various intellectual property rights (IPRs), typically involving a combination of long contracts, copyright, trade secrets/the law of confidence, and/or – in Europe – database rights, as well as may also comprise competition law and personal data integrity rights (Burbidge, 2017; Minssen and Schovsbo, 2018, p. 123). The result is that data are frequently the subject of litigation (Burbidge, 2017; Minssen and Schovsbo, 2018, p. 123). Thus, it has been suggested that more regulations for data-generating internet giants are necessary as well as that the new data economy requires a better approach to competition and antitrust rules (Burbidge, 2017; Minssen and Schovsbo, 2018, p. 123).

The combination of big data and IPRs creates challenges that need to be addressed such as access to data and ownership rights (Minssen and Pierce, 2018, p. 311). In particular, in cases of data mining and data analytics, various forms of IPRs might protect the references to or copying of databases and information (Minssen and Pierce, 2018, p. 311). However, users will need to rely on an exemption to IPR infringement where data is *not* licensed or owned (Minssen and Pierce, 2018, p. 311). This circumstance has led to vigorous disputes between stakeholders, especially data scientists and data “owners” (Minssen and Pierce, 2018, p. 311). Moreover, in the context of big data applications, there is a lot of misunderstanding about the nature, the availability, and legal effects of overlapping rights and remedies.

For example, copyrights might protect the software that helps to collect and process big data sets. However, due to the somewhat unstructured nature of the non-relational databases – a typical characteristic of big data sets and the material they contain – the traditional role and purpose of copyrights and the EU’s sui generis right in databases have been called into question (Gervais, 2019).

With regard to patents, recent case law in Europe (e.g., the German Federal Supreme court case on receptor tyrosine kinase and the UK *Illumina* case) and the U.S. (e.g., the landmark cases *Mayo*, *Myriad*, and *Alice*) might have an impact on precision medicine with its aim to better tailoring treatment to the need of patients in three areas, namely (1) biomarkers and nature-based products, (2) diagnostics, as well as (3) algorithms, big data, and AI (Aboy et al., 2019). In the U.S., recent patent law decisions made it harder – but not impossible – to obtain patent protection for precision medicine inventions, whereas in Europe, a less stringent standard of patent eligibility is applied such as for nature-based biomarkers (Aboy et al., 2019).

Drug companies will most likely use AI systems to expand their traditional drug patent portfolio (Gervais, 2019; Maloney, 2019). However, AI systems could also be used by competitors or patent examiners to predict incremental innovation or to reveal that a patent was ineligible for patent protection due to, for example, the lack of novelty or inventive step (Gervais, 2019; Maloney, 2019). Furthermore, trade secret law, in combination with technological protection measures and contracts, can protect complex algorithms, as well as data sets and sets of insights and correlations generated by AI systems (Gervais, 2019).

Some rights, such as copyrights and trade secrets, are becoming more and more crucial for the commercial protection of big data (Minssen and Pierce, 2018, p. 323). Other rights, such as patents, may not always be applicable, or they may be tactically used in novel ways (Minssen

and Pierce, 2018, p. 323). While more flexible data exclusivity regimes could perhaps address some of the issues posed by traditional IP protections for chemical and pharmaceutical products, it is clear that these developments raise considerable doctrinal and normative challenges to the IPR system and the incentives it creates in a variety of areas (Gervais, 2019; Minssen and Pierce, 2018). Moreover, the full effect and purpose of some IPRs (e.g., as data aggregators) are unclear in the context of big data innovation and need additional study (Minssen and Pierce, 2018, p. 323; Burk, 2015). The un/availability of such rights could not only lead to underinvestment in some areas due to a lack of incentives but also block effects for anticommons scenarios and open innovation in other areas (Minssen and Pierce, 2018, p. 323). Furthermore, the interaction between IPRs and data transparency initiatives and their possible impact on public-private partnerships or open innovation scenarios should be clarified (Minssen and Pierce, 2018, p. 323). For different technological applications, differentiated approaches and IPR user modalities will need to be taken into account and discussed (Minssen and Pierce, 2018, p. 323).

It becomes apparent that more data sharing is necessary in order to achieve the successful deployment of AIs in health care on a large scale. Stakeholders such as companies, agencies, and health care providers need to increasingly consider with whom they are going to collaborate and what data sets under what conditions they are going to share. Some stakeholders are reluctant and refuse to share their data due to, for example, a lack of trust, previous spending on data quality or the protection of commercial and sensitive personal data (Minssen and Schovsbo, 2018, p. 123). To resolve these tensions, legal frameworks would be desirable that promote and incentivize data sharing through, for example, data sharing intermediaries (Richter and Slowinski, 2019) and public-private partnerships, while ensuring adequate protection of data privacy. In cases where stakeholders such as companies act unfairly and collude to entirely

control a market where competition and access are essential for health care, the hope is that more refined competition and antitrust law tools can intervene. To serve this role, competition and antitrust law will need to become more future-oriented to better understand and predict the dynamics and developments of big data and AI in the health care sector. The value of data differs and often depends on multiple factors, including its usage and uniqueness (Kamenir, 2018, p. 2). For instance, diverse data that provides a multitude of signals appears to be more useful and thus valuable since ML is a dynamic experimentation process (Kamenir, 2018, p. 2). It could also be the case that particular combinations where patient data or other medical data is a crucial asset may result in market power if the data is unique and *not* replicable (Kamenir, 2018, p. 2).

5 Conclusion

In this chapter we have given an overview of *what AI is*, and have discussed the *trends and strategies* in the U.S. and Europe, thereby focusing on the ethical and legal debate of AI in health care and research. We have seen that the U.S. has taken a more free-market approach than Europe and that several AI products such as IDx-DR – the first FDA authorized autonomous AI diagnostic system – have already entered the U.S. market. According to one forecast, AI has the potential to contribute up to 13,33 trillion EUR to the worldwide economy in 2030, and the regions that are estimated to gain the most from AI are likely to be China and North America, followed by Southern Europe (European Commission, 2018d, pp. 2, 3). In contrast, Europe emerges as a global player in AI ethics. In particular, the European Commission's High-Level Expert Group on AI published *Ethics Guidelines for Trustworthy AI* in April 2019.

We have also discussed four primary *ethical challenges* that need to be addressed to realize the full potential of AI in health care: (1) informed consent to use, (2) safety and transparency, (3) algorithmic fairness and biases, as well as (4) data privacy. This has been

followed by an analysis of five *legal challenges* in the U.S. and Europe, namely (1) safety and effectiveness, (2) liability, (3) data protection and privacy, (4) cybersecurity, and (5) intellectual property law. In particular, it is crucial that all stakeholders, including AI makers, patients, health care professionals, and regulatory authorities, work together on tackling the identified challenges to ensure that AI will be successfully implemented in a way that is ethically and legally. We need to create a system that is built on *public trust* to achieve a desirable societal goal that AI benefits everyone.

Informed consent, high levels of data protection and privacy, cyber resilience and cybersecurity, algorithmic fairness, an adequate level of transparency and regulatory oversight, high standards of safety and effectiveness, an optimal liability regime for AIs are all key factors that need to be taken into account and addressed to successfully create an AI-driven health care system based on the motto "*Health AIs for All of Us*". In this regard, we not only need to rethink current regulatory frameworks and update them to the new technological developments. But it is also important to have public and political discussions centered on the ethics of AI-driven health care such as its implications on the human workforce and the society as a whole. AI has tremendous potential for improving our health care system, but we can only unlock its potential by already starting now to address the ethical and legal challenges facing us.

Acknowledgements

This article is forthcoming in *Artificial Intelligence in Healthcare*, 1st edition, Adam Bohr, Kaveh Memarzadeh (eds.), Chapter 12: Ethical and Legal Challenges of Artificial Intelligence-Driven Health Care, ISBN: 9780128184387, Copyright Elsevier (2020). This draft is current as of April 23, 2020, for the definitive version check out *Artificial Intelligence in Healthcare* when published. This research for this contribution was supported by a Novo Nordisk Foundation-grant for a scientifically independent Collaborative Research Programme in Biomedical Innovation Law (grant agreement number NNF17SA0027784).

References

- Aboy, M., Liddell, K., Crespo, C., Cohen, G., Liddicoat, J., Gerke, S., & Minssen, S. (2019). How does emerging patent case law in the US and Europe affect precision medicine? *Nature Biotechnology*, 37, 1118–1125. doi: 10.1038/s41587-019-0265-1.
- Accenture. Artificial Intelligence (AI): Healthcare’s New Nervous System. (2017). https://www.accenture.com/_acnmedia/pdf-49/accenture-health-artificial-intelligence.pdf Accessed 8 August 2019.
- Ada. Your personal health guide. (2020). <https://ada.com> Accessed 16 March 2020.
- AI HLEG. Ethics Guidelines For Trustworthy AI. (2019a). <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai> Accessed 8 August 2019.
- AI HLEG. A Definition of AI. Main Capabilities and Disciplines. (2019b). <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai> Accessed 16 March 2020.

AI HLEG. Policy and Investment Recommendations for Trustworthy AI. (2019c).

<https://ec.europa.eu/digital-single-market/en/news/policy-and-investment-recommendations-trustworthy-artificial-intelligence> Accessed 8 August 2019.

Arney, D., Senges, M., Gerke, S., Canca, C., Haaber Ihle, L., Kaiser, N., Kakarmath, S., Kupke, A., Gajeele, A., Lynch, S., & Melendez, L. A User-Focused Transdisciplinary Research Agenda for AI-Enabled Health Tech Governance. (2019).

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3385398 Accessed 8 August 2019.

Arterys. (2019). <https://www.arterys.com> Accessed 8 August 2019.

Babic, B., Gerke, S., Evgeniou, T., & Cohen, I. G. (2019). Algorithms on Regulatory Lockdown in Medicine. *Science*, 366, 1202-1204. doi: 10.1126/science.aay9547.

Bezair, J., Felton, K. W., Greve, P., & Allen, D. Hospital negligent credentialing liability. (2017). <https://www.willistowerswatson.com/en-US/Insights/2017/08/hospital-negligent-credentialing-liability> Accessed 8 August 2019.

Brown, J. (2018). IBM Watson Reportedly Recommended Cancer Treatments That Were ‘Unsafe and Incorrect’. *Gizmodo*. <https://gizmodo.com/ibm-watson-reportedly-recommended-cancer-treatments-tha-1827868882> Accessed 8 August 2019.

Burbidge, R. Medical Data in a twist – Technomed v Bluecrest. (2017).

<http://ipkitten.blogspot.com/2017/09/medical-data-in-twist-technomed-v.html> Accessed 8 August 2019.

Burk, D. L. (2015). Patents as Data Aggregators in Personalized Medicine. *Boston University Journal of Science and Technology Law*, 21.

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2597525## Accessed 8 August 2019.

- Burt, A. Is there a 'right to explanation' for machine learning in the GDPR? (2017).
<https://iapp.org/news/a/is-there-a-right-to-explanation-for-machine-learning-in-the-gdpr>
Accessed 8 August 2019.
- Carrier, M. A. (2009). *Innovation for the 21st Century: Harnessing the Power of Intellectual Property and Antitrust Law*. New York: Oxford Univ. Press.
- Cohen, I. G., Amarasingham, R., Shah, A., Xie, B., & Lo, B. (2014). The Legal and Ethical Concerns That Arise From Using Complex Predictive Analytics In Health Care. *Health Affairs*, 7, 1139–1147. doi: 10.1377/hlthaff.2014.0048.
- Cohen, I. G. (2018a). Petrie-Flom Center launches Project on Precision Medicine, Artificial Intelligence, and the Law (PMAIL). *Harvard Law Today*.
<https://today.law.harvard.edu/petrie-flom-center-launches-project-precision-medicine-artificial-intelligence-law-pmail> Accessed 8 August 2019.
- Cohen, I. G. (2018b). Is There a Duty to Share Healthcare Data? In I. G. Cohen, H. F. Lynch, E. Vayena, & U. Gassner (Eds.), *Big Data, Health Law, and Bioethics* (pp. 209–222). Cambridge: Cambridge University Press.
- Cohen, I. G., & Mello, Michelle M. (2018). HIPAA and Protecting Health Information in the 21st Century. *JAMA*, 320, 231–232. doi:10.1001/jama.2018.5630.
- Cohen, I. G., & Pearlman, A. (2018). Smart pills can transmit data to your doctors, but what about privacy? *New Scientist*. <https://www.newscientist.com/article/2180158-smart-pills-can-transmit-data-to-your-doctors-but-what-about-privacy> Accessed 8 August 2019.
- Congress. (2018). Summary: H.R.3359 – 115th Congress (2017-2018).
<https://www.congress.gov/bill/115th-congress/house->

- bill/3359?q=%7B%22search%22%3A%5B%22cybersecurity+and+infrastructure%22%5D%7D&r=1 Accessed 8 August 2019.
- Copeland, R. Google's 'Project Nightingale' Gathers Personal Health Data on Millions of Americans. (2019). <https://www.wsj.com/articles/google-s-secret-project-nightingale-gathers-personal-health-data-on-millions-of-americans-11573496790> Accessed 17 March 2020.
- Corti. A Co-Pilot for Medical Interviews. (2019). <https://corti.ai> Accessed 16 March 2020.
- Cossins, D. (2018). Discriminating algorithms: 5 times AI showed prejudice. *New Scientist*. <https://www.newscientist.com/article/2166207-discriminating-algorithms-5-times-ai-showed-prejudice> Accessed 8 August 2019.
- De Fauw, J., Ledsam, J. R., Romera-Paredes, B., Nikolov S., Tomasev N., Blackwell, S., Askham, H., Glorot, X., O'Donoghue, B., Visentin, D., Van Den Driessche, G., Lakshminarayanan, B., Meyer, C., Mackinder, F., Bouton, S., Ayoub, K., Chopra, R., King, D., Karthikesalingam, A., Hughes, C. O., Raine, R., Hughes, J., Sim D. A., Egan, C., Tufail, A., Montgomery, H., Hassabis, D., Rees, G., Back, T., Khaw, P. T., Suleyman, M., Cornebise, J., Keane, P. A., & Ronneberger, O. (2018). Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature Medicine*, 24, 1342-1350. doi: 10.1038/s41591-018-0107-6.
- Dinerstein v. Google. (2019). No. 1:19-cv-04311.
- Dutton, T. An Overview of National AI Strategies. (2018). <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd> Accessed 8 August 2019.
- European Commission. Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the

- Committee of the Regions. Artificial Intelligence for Europe. COM(2018) 237 final. (2018a). <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52018DC0237&from=EN> Accessed 8 August 2019.
- European Commission. Member States and Commission to work together to boost artificial intelligence “made in Europe”. (2018b). http://europa.eu/rapid/press-release_IP-18-6689_en.htm Accessed 8 August 2019.
- European Commission. Commission Staff Working Document. Liability for emerging digital technologies. *Accompanying the document* Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions. Artificial Intelligence for Europe. SWD(2018) 137 final. (2018c). <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52018SC0137&from=en> Accessed 8 August 2019.
- European Commission. USA-China-EU plans for AI: where do we stand? (2018d). https://ec.europa.eu/growth/tools-databases/dem/monitor/sites/default/files/DTM_AI%20USA-China-EU%20plans%20for%20AI%20v5.pdf Accessed 8 August 2019.
- European Commission. White Paper On Artificial Intelligence – A European approach to excellence and trust. (2020a). https://ec.europa.eu/info/sites/info/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf Accessed 16 March 2020.
- European Commission. Communication from The Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions. A European strategy for data. (2020b).

https://ec.europa.eu/info/sites/info/files/communication-european-strategy-data-19feb2020_en.pdf Accessed 16 March 2020.

European Commission. Report from the Commission to the European Parliament, the Council and the European Economic and Social Committee. Report on the safety and liability implications of Artificial Intelligence, the Internet of Things and robotics. (2020c).

https://ec.europa.eu/info/sites/info/files/report-safety-liability-artificial-intelligence-feb2020_en_1.pdf Accessed 16 March 2020.

Executive Office of the President. Memorandum for the Heads of Executive Departments and Agencies. M-18-22. (2018). <https://www.whitehouse.gov/wp-content/uploads/2018/07/M-18-22.pdf> Accessed 8 August 2019.

European Parliament. Parliament decides to postpone new requirements for medical devices. (2020) <https://www.europarl.europa.eu/news/en/press-room/20200415IPR77113/parliament-decides-to-postpone-new-requirements-for-medical-devices> Accessed 23 April 2020.

FDA. Summary K163253. (2017a).

https://www.accessdata.fda.gov/cdrh_docs/pdf16/K163253.pdf Accessed 30 April 2019.

FDA. Software as a Medical Device (SAMD): Clinical Evaluation. (2017b).

<https://www.fda.gov/media/100714/download> Accessed 8 August 2019.

FDA. Digital Health Innovation Action Plan. (2017c).

<https://www.fda.gov/downloads/MedicalDevices/DigitalHealth/UCM568735.pdf>
Accessed 8 August 2019.

FDA. FDA permits marketing of artificial intelligence-based device to detect certain diabetes-related eye problems. (2018a).

<https://www.fda.gov/newsevents/newsroom/pressannouncements/ucm604357.htm>

Accessed 8 August 2019.

FDA. DeNovo Summary DEN180001. (2018b).

https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN180001.pdf Accessed 8 August 2019.

FDA. DeNovo Summary DEN180005. (2018c).

https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN180005.pdf Accessed 8 August 2019.

FDA. FDA permits marketing of artificial intelligence algorithm for aiding providers in detecting wrist fractures. (2018d).

<https://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm608833.htm>
Accessed 8 August 2019.

FDA. Is The Product A Medical Device? (2018e).

<https://www.fda.gov/medicaldevices/deviceregulationandguidance/overview/classifyyourdevice/ucm051512.htm> Accessed 8 August 2019.

FDA. Changes to Existing Medical Software Policies Resulting from Section 3060 of the 21st Century Cures Act. (2019a). <https://www.fda.gov/media/109622/download> Accessed 17 March 2019.

FDA. Clinical Decision Support Software. Draft Guidance for Industry and Food and Drug Administration Staff. (2019b). <https://www.fda.gov/media/109618/download> Accessed 17 March 2020.

FDA. Digital Health Software Precertification (Pre-Cert) Program. (2019c).

<https://www.fda.gov/medicaldevices/digitalhealth/digitalhealthprecertprogram/default.htm> Accessed 8 August 2019.

FDA. Developing a Software Precertification Program: A Working Model (v.1.0). (2019d).

<https://www.fda.gov/media/119722/download> Accessed 18 March 2020.

FDA. Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML) - Based Software as a Medical Device (SaMD). (2019e).

<https://www.fda.gov/media/122535/download> Accessed 8 August 2019.

Fefegha, A. Racial Bias and Gender Bias Examples in AI systems. (2018).

<https://medium.com/thoughts-and-reflections/racial-bias-and-gender-bias-examples-in-ai-systems-7211e4c166a1> Accessed 8 August 2019.

Figure Eight. What is Training Data? (2020). <https://www.figure-eight.com/resources/what-is-training-data> Accessed 17 March 2020.

Finlayson, S. G., Bowers, J. D., Ito, J., Zittrain, J. L., Beam, A. L., & Kohane, I. S. (2019).

Adversarial attacks on medical machine learning. *Science*, 363, 1287–1289. doi: 10.1126/science.aaw4399.

FLI Team, AI Policy – United States. (2019). <https://futureoflife.org/ai-policy-united-states>

Accessed 8 August 2019.

FLI Team. State of California Endorses Asilomar AI Principles. (2018).

<https://futureoflife.org/2018/08/31/state-of-california-endorses-asilomar-ai-principles> Accessed 8 August 2019.

- Froomkin, A. M., Kerr, I., & Pineau, J. (2019). When AIs Outperform Doctors: Confronting the Challenges of a Tort-Induced Over-Reliance on Machine Learning. *Arizona Law Review*, *61*, 34–99.
- G20. G20 Ministerial Statement on Trade and Digital Economy. (2019).
[https://www.mofa.go.jp/files/000486596.pdf#targetText=a\)%20AI%20actors%20should%20respect,and%20internationally%20recognized%20labor%20rights](https://www.mofa.go.jp/files/000486596.pdf#targetText=a)%20AI%20actors%20should%20respect,and%20internationally%20recognized%20labor%20rights). Accessed 17 March 2020.
- Gerke, S. (2018). The EU's GDPR in the Health Care Context. *Bill of Health*.
<http://blog.petrieflom.law.harvard.edu/2018/05/30/the-eus-gdpr-in-the-health-care-context> Accessed 8 August 2019.
- Gerke, S., Kramer, D. B., & Cohen, I. G. (2019). Ethical and Legal Challenges of Artificial Intelligence in Cardiology. *AIMed Magazine*, *2*, 12–17.
- Gerke, S., Minssen, T., Yu, H., & Cohen, I. G. (2019). Ethical and Legal Issues of Ingestible Electronic Sensors. *Nature Electronics*, *2*, 329-334. <https://doi.org/10.1038/s41928-019-0290-6>.
- Gerke, S. Babic, B., Evgeniou, T., and Cohen, I. G. (2020a). The need for a system view to regulate artificial intelligence/machine learning-based software as medical device. *npj Digital Medicine*, *3*, Article number: 53. <https://doi.org/10.1038/s41746-020-0262-2>.
- Gerke, S., Yeung, S., & Cohen, I. G. (2020b). Ethical and Legal Aspects of Ambient Intelligence in Hospitals. *JAMA*, *323*, 601-602. doi:10.1001/jama.2019.21699.
- German Federal Government. Strategie Künstliche Intelligenz der Bundesregierung (2018).
<https://www.bmwi.de/Redaktion/DE/Publikationen/Technologie/strategie-kuenstliche->

- intelligenz-der-bundesregierung.pdf?__blob=publicationFile&v=8 Accessed 8 August 2019.
- Gervais, D. Big Data and Intellectual Property Law. (2019). http://www.ceipi.edu/en/news/piece-of-news/?tx_ttnews%5Btt_news%5D=10942&cHash=2f3a1ea32cb69dd6b7ed38659ef6e440 Accessed 8 August 2019.
- Goodman, B., & Flaxman, S. (2017). European Union Regulations on Algorithmic Decision Making and a “Right to Explanation”. *AI Magazine*, 38, 50–57. doi: 10.1609/aimag.v38i3.2741.
- Graham, C. NHS cyber attack: Everything you need to know about ‘biggest ransomware’ offensive in history. (2017). <https://www.telegraph.co.uk/news/2017/05/13/nhs-cyber-attack-everything-need-know-biggest-ransomware-offensive> Accessed 8 August 2019.
- House of Lords. AI in the UK: ready, willing and able? (2018). <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf> Accessed 8 August 2019.
- IBM. IBM Letter of Support for the 21st Century Cures Act. (2016). <https://www.ibm.com/blogs/policy/ibm-letter-support-21st-century-cures-act> Accessed 8 August 2019.
- IBM. IBM Watson for Oncology (2020). <https://www.ibm.com> Accessed 17 March 2020.
- ICO. RFA0627721 – provision of patient data to DeepMind. (2017a). <https://ico.org.uk/media/2014353/undertaking-cover-letter-revised-04072017-to-first-person.pdf> Accessed 8 August 2019.

- ICO. Royal Free – Google DeepMind trial failed to comply with data protection law. (2017b). <https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2017/07/royal-free-google-deepmind-trial-failed-to-comply-with-data-protection-law> Accessed 8 August 2019.
- IDx Technologies Inc. (2018). <https://www.eyediagnosis.net/idx-dr> Accessed 8 August 2019.
- IMDRF. Software as a Medical Device (SaMD): Key Definitions. (2013). <http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-131209-samd-key-definitions-140901.pdf> Accessed 16 March 2019.
- Kamenir, E. Themes and Takeaways from the FTC Hearings on the Intersection of Big Data, Privacy, and Competition. (2018). <https://www.competitionpolicyinternational.com/wp-content/uploads/2018/11/North-America-Column-November-2018-7-Full.pdf> Accessed 8 August 2019.
- Kaminski, M. A. (2018). The Right to Explanation, Explained. *U of Colorado Law Legal Studies Research Paper No. 18-24*, 1–25.
- Katz, M. L., & Shapiro, C. (1992). Product Introduction with Network Externalities. *The Journal of Industrial Economics*, 40, 55–83.
- Klugman, C. M., Dunn, L.B., Schwartz, J., & Cohen, I.G. (2018). The Ethics of Smart Pills and Self-Acting Devices: Autonomy, Truth-Telling, and Trust at the Dawn of Digital Medicine. *AJOB*, 18, 38–47. <https://doi.org/10.1080/15265161.2018.1498933>.
- Knight, W. (2019). Trump has a plan to keep America first in artificial intelligence. *MIT Technology Review*. <https://www.technologyreview.com/s/612926/trump-will-sign-an-executive-order-to-put-america-first-in-artificial-intelligence> Accessed 8 August 2019.

- Lemley, M. A., Shafir, Z. (2011). Who Chooses Open-Source Software? *The University of Chicago Law Review*, 78, 139–164.
- London, A. J. (2019). Artificial Intelligence and Black-Box Medical Decisions: Accuracy versus Explainability. *The Hastings Center Report*, 49, 15–21. doi: 10.1002/hast.973.
- Lusa. Barreiro Hospital disputes judicially fine of 400 thousand euros Data Commission. (2018). https://translate.google.com/translate?depth=1&hl=en&ie=UTF8&prev=_t&rurl=translate.google.com&sl=auto&sp=nmt4&tl=en&u=https://www.publico.pt/2018/10/22/sociedade/noticia/hospital-barreiro-contesta-judicialmente-coima-400-mil-euros-comissao-dados-1848479&xid=17259,15700022,15700124,15700149,15700186,15700190,15700201 Accessed 8 August 2019.
- Maack, M. M. Europe launches a heart arrack-detecting AI for emergency calls. (2018). <https://thenextweb.com/artificial-intelligence/2018/04/25/europe-launches-heart-attack-detecting-ai-emergency-calls> Accessed 8 August 2019.
- Maloney, D. AI Patent Trolls Now on the Job for Drug Companies. (2019). <https://hackaday.com/2019/01/30/ai-patent-trolls-now-on-the-job-for-drug-companies> Accessed 8 August 2019.
- Marr, B. (2017). First FDA Approval For Clinical Cloud-Based Deep Learning In Healthcare. *Forbes*. <https://www.forbes.com/sites/bernardmarr/2017/01/20/first-fda-approval-for-clinical-cloud-based-deep-learning-in-healthcare/#6af107d161c8> Accessed 8 August 2019.
- MDCG. (2019). Guidance on Qualification and Classification of Software in Regulation (EU) 2017/745 – MDR and Regulation (EU) 2017/746 – IVDR. <https://ec.europa.eu/docsroom/documents/37581> Accessed 18 March 2020.

- Mehta, N., & Devarakonda, M. V. (2018). Machine learning, natural language programming, and electronic health records: The next step in the artificial intelligence journey? *Journal of Allergy and Clinical Immunology*, *141*, 2019–2021.e1.
<https://doi.org/10.1016/j.jaci.2018.02.025>.
- Minssen, T., & Pierce, J. (2018). Big Data and Intellectual Property Rights in the Health and Life Sciences. In I. G. Cohen, H. F. Lynch, E. Vayena, & U. Gasser (Eds.), *Big Data, Health Law, and Bioethics* (pp. 311–323). Cambridge: Cambridge University Press.
- Minssen, T., & Schovsbo, J. (2018). Big Data in the Health and Life Sciences: What Are the Challenges for European Competition Law and Where Can They Be Found? In X. Seuba, C. Geiger, & J. Pénin (Eds.), *Intellectual Property and Digital Trade in the Age of Artificial Intelligence and Big Data* (pp. 121–130). CEIPI-ICTSD.
- Minssen, T., Gerke, S., Aboy, M., Price, N., and Cohen, I. G. (2020). Regulatory responses to medical machine learning. *Journal of Law and the Biosciences*, 1–18.
[doi:10.1093/jlb/ljaa002](https://doi.org/10.1093/jlb/ljaa002).
- Moorfields Eye Hospital. Next phase of Moorfields work with Google Health. (2019).
<https://www.moorfields.nhs.uk/news/next-phase-moorfields-work-google-health>
Accessed 16 March 2020.
- NTF. Liability for Artificial Intelligence and Other Emerging Digital Technologies. (2019).
<https://ec.europa.eu/transparency/regexpert/index.cfm?do=groupDetail.groupMeetingDoc&docid=36608> Accessed 18 March 2020.
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. *Science*, *366*, 447–453. doi:
[10.1126/science.aax2342](https://doi.org/10.1126/science.aax2342).

OECD. OECD Principles on AI. (2019). <https://www.oecd.org/going-digital/ai/principles>

Accessed 16 March 2019.

Pinsent Masons. New ‘digital’ pills pose data protection and cybersecurity challenges for drugs manufacturers and health bodies, says expert. (2017). <https://www.out-law.com/en/articles/2017/november/new-digital-pills-pose-data-protection-and-cybersecurity-challenges-for-drugs-manufacturers-and-health-bodies-says-expert>

Accessed 8 August 2019.

Price II, W. N. (2017). Artificial Intelligence in Health Care: Applications and Legal Implications. *The SciTech Lawyer*, 14, 10–13.

Price II, W. N. (forthcoming 2019). Medical AI and Contextual Bias, *Harvard Journal of Law & Technology*, available at SSRN:

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3347890 Accessed 8 August 2019.

Price II, W. N., & Cohen, I. G. (2019). Privacy in the Age of Medical Big Data. *Nature Medicine*, 25, 37–43. doi: 10.1038/s41591-018-0272-7.

Price II, W. N., Gerke, S., & Cohen, I. G. Potential Liability for Physicians Using Artificial Intelligence. *JAMA*, 322, 1765-1766. doi: 10.1001/jama.2019.15064.

Price II, W. N., Kaminski, M. G., Minssen, T., & Spector-Bagdady, K. (2019). Shadow Health Records Meet New Data Privacy Laws. How Will Research Respond to a Changing Regulatory Space? *Science*, 363, 448–450. doi: 10.1126/science.aav5133.

Richter, H., & Slowinski, P. R. (2019). The Data Sharing Economy: On the Emergence of New Intermediaries. *International Review of Intellectual Property and Competition Law*, 50, 4–29. <https://doi.org/10.1007/s40319-018-00777-7>.

- Roberts, J. L., Cohen, I. G., Deubert, C. R., & Lynch, H. F. (2017). Evaluating NFL Player Health and Performance: Legal and Ethical Issues. *University of Pennsylvania Law Review*, *165*, 227–314.
- Ross, C., & Swetlitz, I. (2017). IBM to Congress: Watson will transform health care, so keep your hands off our supercomputer. *STAT*. <https://www.statnews.com/2017/10/04/ibm-watson-regulation-fda-congress> Accessed 8 August 2019.
- Ross, C., & Swetlitz, I. (2018). IBM’s Watson supercomputer recommended ‘unsafe and incorrect’ cancer treatments, internal documents show. *STAT*. <https://www.statnews.com/2018/07/25/ibm-watson-recommended-unsafe-incorrect-treatments> Accessed 8 August 2019.
- Sharkey, N. (2018). The impact of gender and race bias in AI. *Humanitarian Law & Policy*. <https://blogs.icrc.org/law-and-policy/2018/08/28/impact-gender-race-bias-ai> Accessed 8 August 2019.
- Short, E. (2018). It turns out Amazon’s AI hiring tool discriminated against women. *Silicon Republic*. <https://www.siliconrepublic.com/careers/amazon-ai-hiring-tool-women-discrimination> Accessed 8 August 2019.
- Smart, W. (2018). Lessons learned review of the WannaCry Ransomware Cyber Attack. <https://www.england.nhs.uk/wp-content/uploads/2018/02/lessons-learned-review-wannacry-ransomware-cyber-attack-cio-review.pdf> Accessed 8 August 2019.
- The Economist. The world’s most valuable resource is no longer oil, but data. (2017). <https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data> Accessed 8 August 2019.

The Medical Futurist. FDA Approvals for Smart Algorithms in Medicine in One Giant Infographic. (2019). <https://medicalfuturist.com/fda-approvals-for-algorithms-in-medicine> Accessed 16 March 2019.

Topol, E. J. (2019). High-Performance Medicine: The Convergence of Human and Artificial Intelligence. *Nature Medicine*, 25, 44–56.

UK Department of Health and Social Care. New code of conduct for artificial intelligence (AI) systems used by the NHS. (2019). <https://www.gov.uk/government/news/new-code-of-conduct-for-artificial-intelligence-ai-systems-used-by-the-nhs> Accessed 8 August 2019.

UK Government. Industrial Strategy. Artificial Intelligence Sector Deal. (2018). https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/702810/180425_BEIS_AI_Sector_Deal__4_.pdf Accessed 8 August 2019.

UK Nuffield Council on Bioethics. Artificial intelligence (AI) in healthcare and research. (2018). <http://nuffieldbioethics.org/wp-content/uploads/Artificial-Intelligence-AI-in-healthcare-and-research.pdf> Accessed 8 August 2019.

Ultromics. (2019). <http://www.ultromics.com> Accessed 8 August 2019.

U.S. Department of Homeland Security. (2019). Cybersecurity. <https://www.dhs.gov/topic/cybersecurity> Accessed 8 August 2019.

U.S. Government. Preparing for the Future of Artificial Intelligence. (2016a). https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf Accessed 8 August 2019.

U.S. Government. The National Artificial Intelligence Research and Development Strategic Plan. (2016b).

- https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/national_ai_rd_strategic_plan.pdf Accessed 8 August 2019.
- U.S. Government. Artificial Intelligence, Automation, and the Economy. (2016c).
<https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF> Accessed 8 August 2019.
- U.S. Government. Accelerating America's Leadership in Artificial Intelligence. (2019).
<https://www.whitehouse.gov/articles/accelerating-americas-leadership-in-artificial-intelligence> Accessed 30 April 2019.
- Vincent, J. (2018). AI that detects cardiac arrests during emergency calls will be tested across Europe this summer. *The Verge*. <https://www.theverge.com/2018/4/25/17278994/ai-cardiac-arrest-corti-emergency-call-response> Accessed 8 August 2019.
- 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. *International Data Privacy Law*, 7, 76–99. doi:10.1093/idpl/ix005.
- Wachter, S., Mittelstadt, B., & Russell, C. (2018). Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. *Harvard Journal of Law & Technology*, 31, 842–887.
- Wahl, B., Cossy-Gantner, A., Germann, S., & Schwalbe, N. R. (2018). Artificial intelligence (AI) and global health: how can AI contribute to health in resource-poor settings? *BMJ Glob Health*, 3, e000798. doi: 10.1136/bmjgh-2018-000798.
- Wexler, R. (2018). Life, Liberty, and Trade Secrets: Intellectual Property in the Criminal Justice System. *Stanford Law Review*, 70, 1343–1429.

- White House. Summary of the 2018 White House Summit on Artificial Intelligence for American Industry. (2018). <https://www.whitehouse.gov/wp-content/uploads/2018/05/Summary-Report-of-White-House-AI-Summit.pdf> Accessed 30 April 2019.
- White House. Executive Order on Maintaining American Leadership in Artificial Intelligence. (2019). <https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence> Accessed 8 August 2019.
- White House. Draft Memorandum for the Heads of Executive Departments and Agencies. Guidance for Regulation of Artificial Intelligence Applications. (2020a). <https://www.whitehouse.gov/wp-content/uploads/2020/01/Draft-OMB-Memo-on-Regulation-of-AI-1-7-19.pdf> Accessed 16 March 2020.
- White House. American Artificial Intelligence Initiative: Year One Annual Report. (2020b). <https://www.whitehouse.gov/wp-content/uploads/2020/02/American-AI-Initiative-One-Year-Annual-Report.pdf> Accessed 16 March 2020.
- Working Party on the Protection of Individuals with Regard to the Processing of Personal Data. (2018). Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679. WP251rev.01.
- Yu, K. H., Beam, A. L., & Kohane, I. S. (2018). Artificial intelligence in healthcare. *Nature Biomedical Engineering*, 2, 719–731. <https://doi.org/10.1038/s41551-018-0305-z>.