Artificial Intelligence in Healthcare from a Policy Perspective

Monica Aggarwal, Christian Gingras, and Raisa Deber

Abstract The growth of Artificial Intelligence (AI) technologies in health care is driving a growing recognition among policymakers, businesses and researchers that there is a need for policies to address certain potential consequences of AI innovation. In this chapter, we provide insight on several policy implications and challenges relating to the impact of AI on accuracy, fairness and transparency, data privacy and consent, accountability, and workforce disruption. These issues include: monitoring of accuracy; minimizing bias and encouraging transparency, ensuring appropriate use, assessment of who is receiving the information and how it is being used, protecting privacy through data protection requirements, enactment of laws that defines accountabilities, establishment of policies for labour disruption; implementation of professional standards and codes of conduct; adapting educational training for clinicians; and determining what technologies will be insured and funded. Additional complexities arise when AI crosses geographic boundaries. The design, development and implementation of policy and regulation should be in conjunction with a diversity of stakeholders including product developers, researchers, patients, health care providers and policymakers.

Keywords Artificial intelligence · Policy · Regulation · Ethics · Algorithm bias · Privacy · Consent · Accountability · Human resources

1 Introduction

Artificial Intelligence (AI) is a branch of computer science concerned with the development of systems that can perform tasks that usually require human intelligence,

M. Aggarwal (✉)
Dalla Lana School of Public Health, University of Toronto, Toronto, Canada
e-mail: monica.aggarwal@utoronto.ca

C. Gingras
Innovative Health Care Management Solutions Inc., Toronto, Canada

R. Deber
Institute of Health Policy, Management and Evaluation, University of Toronto, Toronto, Canada

M. Househ et al. (eds.), *Multiple Perspectives on Artificial Intelligence in Healthcare*,
Lecture Notes in Bioengineering, https://doi.org/10.1007/978-3-030-67303-1_5

such as problem-solving, reasoning, and recognition (An Overview of Clinical Applications of Artificial Intelligence 2018). AI has significant prospect to fundamentally transform the delivery of health care. Despite the significant potential of AI, there are several policy challenges that need to be considered by policymakers as they embark on the AI journey.

Analyzing the policy implications is complex, because AI is not homogeneous (Scherer 2016), and the policy issues may vary accordingly. AI has been suggested for a wide variety of tasks, including but not restricted to assisting in health data management (including streamlining administrative processes to facilitate quality assurance); searching the medical literature in specialized domains; assisting in repetitive jobs (such as analyzing radiology images); smart algorithms to help interpret tests, improve diagnostics and generate targeted treatment pathway design; and patient empowerment (including allowing self-monitoring patient management) (Mesko 2017). The policy implications accordingly may vary depending on what the goals of the AI are, and who it is serving.

Policy can be defined as "a set of interrelated decisions taken by a political actor or group of actors concerning the selection of goals and the means of achieving them within a specified situation where these decisions should, in principle, be within the power of these actors to achieve" (Jenkins 1978). Policy makers can use a variety of policy instruments to accomplish this, which may include exhortation (providing information), expenditure (subsidizing activities), regulation, or public ownership (Doern and Phidd 1992). As these definitions recognize, there is likely to be significant variation in who would be responsible for these policy decisions, and the policy instruments they could use.

The growth of AI technologies in health care is driving the growing recognition among policymakers, businesses and researchers that there is a need for the establishment of policies to address the consequences of AI innovation. Several countries have released strategies to encourage the use and development of AI (Dutton 2018; OECD 2019). A number of approaches are being used to regulate AI, including: encouraging AI actors to develop self-regulatory mechanisms such as codes of conduct, accountability standards, ethical frameworks and best practices; and establishing public- and private-sector oversight mechanisms in the form of compliance reviews, audits, conformity assessments and certification schemes for AI applications (OECD 2019).

The purpose of this chapter is to provide insight to policymakers, researchers, businesses, clinicians, patients and caregivers on the policy implications and challenges relating to the impact of AI on such issues as: accuracy, fairness and transparency, data privacy and consent, accountability, and workforce disruption. Table 1 provides an overview of some challenges and opportunities.

Table 1 Challenges and opportunities

Challenges	Opportunities
Lack of universal definition of AI	Jurisdiction establish a consensus-based definition of AI amongst all stakeholders for the purpose of designing AI Policy and Regulation
Risk related to AI is unknown and algorithms are continually adapting and changing.	Basic "rules" anchored in Ethics are developed to allow for adaptability as AI risks and capabilities evolve
AI can discriminate due to algorithm bias or training data bias	Several approaches can be used to minimize the risk of discrimination. This includes: awareness building; funding development of representative datasets, organizational diversity policies and practices; recruitment of developers from diverse background; local and international standards (including post-market monitoring); technical solutions to detect and correct algorithmic bias; self-regulatory or regulatory approaches, and ethical governance and standards, and ethical auditing
Deep learning and Machine learning result in lack of transparency	Establish regulation and policies that articulates how transparency will be handled for consumers/patients.
Legal framework does not exist for who is accountable when harm is caused by autonomous AI applications	Laws must be developed in which there are multiple options for consideration: 1 – Establish AI as a "Person" under the law 2 – Introduce Enterprise Liability, assigning responsibility to all group involved in the creating and implementation of AI 3 – Modify duties of care of Health Professional to take into account AI and for them to exercise due care in its application
Privacy legislation is not well established around the globe. In the absence of laws and policies, significant investment may be invalidated once a framework is updated	Establish appropriate Policy and Regulation of AI to establish rules of engagement for the development of AI
AI challenges the traditional concept of consent	Establish guidelines for health care providers and private Companies on rules around the use of data and providing patients (or consumers) with information on the potential uses of their data
Fear of work displacement	Establish clear policies in the event that employment is displaced by AI function (i.e.: retraining programs, employment insurance, alternative taxation, etc…)
Adoption of AI in health care depend upon acceptance by health care professionals	Engage health care professional in discussions involving policy, product development and provide clinicians with education on the benefits and limitations of AI and how to use it.

2 Artificial Intelligence Policy

Ideally, AI policy would maximize AI innovation and benefits, and minimize its potential costs and risk. Achieving the appropriate balance is not obvious, and may depend on the priorities of different decision makers.

AI software is viewed by regulatory bodies such as Health Canada and the FDA as a medical device (Jaremko et al. 2019). Accordingly, an intended use statement must be submitted by the device manufacturer to receive approval (Jaremko et al. 2019). If approved, the regulatory body can place additional controls on the device to ensure safety. In this case, liability rests with the health care practitioner using that device (Jaremko et al. 2019). An important delineator in legal and regulatory risk assessments is whether AI acts independently (i.e., the software makes diagnostic or treatment decisions that are automatically implemented or that the human user is not able to evaluate) or whether it augments or supports clinical decision-making (i.e., the software makes recommendations but the final decisions are made by a clinician) (Sullivan and Schweikart 2019). However, current legal standards and doctrines regarding medical malpractice are not always clear on where responsibilities should lie when AI supports or autonomously delivers healthcare services (Sullivan and Schweikart 2019).

One question is who the intended user of the AI will be. Much of AI could be viewed as an extension of existing technology. If a physician orders diagnostic

testing (including imaging or laboratory tests), they would normally be returned with an interpretation of what these results mean. For such applications, similar regulatory controls would presumably exist, including ensuring that the test is being performed accurately, that the results are valid, and that the receiving provider understands the limitations of the test results and is responsible for communicating with the patients and ensuring that they understand the meaning of the results, and of the treatments that may be suggested. Such uses of AI do not represent significant new policy challenges.

To the extent that AI goes beyond such current testing, however, new issues may arise. One set of issues may result if the test results are provided to users other than clinicians. This may resemble such current examples as genetic tests provided to patients who order them on-line; there is a considerable literature about the potential risks to patients of receiving inaccurate information. Similarly, test results may be provided to employers (who may use them to discharge employees), insurers (who may use them to deny coverage or increase premiums), etc.

Another set of issues arises if the AI provider is not in the same jurisdiction as the recipient. While this can be advantageous (e.g., to patients in rural/remote areas without the infrastructure to provide such tests), it can also be problematic to the extent that it is unclear who will set and enforce the regulations to ensure that the tests are accurate, and that other ethical and regulatory issues are complied with.

Currently, there are two main approaches used for the regulation of AI that represent different balances between encouraging innovation, and avoiding risks. The European Union (EU) has adopted the "precautionary principle" (Thierer et al. 2017) approach which imposes limits or bans on certain applications due to their potential risks (Pesapane et al. 2018). The European regulatory regime is based on three directives on medical devices in which it requires manufacturers to ensure that the devices they produce are fit for their intended purpose and they comply with the requirements set out by the directives (Pesapane et al. 2018). This assessment can take place by the manufacturer or by a notified body, which is an independent accredited certification organization appointed by the EU Member States (Pesapane et al. 2018). On the other hand, the United States has adopted the "permissionless" innovation approach (Thierer et al. 2017; Pesapane et al. 2018) which permits experimentation with the expectation that issues will be addressed as they arise. The Food and Drug Administration (FDA) categorizes the medical devices into three classes, according to their uses and risks, in which the degree of regulation increases with more risk (Allen 2019). These approaches are hotly debated since the "precautionary principle" approach is seen to inhibit innovation and the "permissionless" approach is seen to increase risk of harm. The consensus appears to be that an ideal approach would be one that is a balance between these approaches.

Examples of policy issues in AI include: accuracy, fairness and transparency; data privacy and consent; accountability, and workforce disruption.

3 Accuracy, Fairness and Transparency

A substantial body of AI literature draws attention to the potential for bias by AI applications towards certain population sub-groups, which can result in discrimination, inequality and marginalization. In machine learning, algorithms rely on multiple data sets, or training data, that are used to make predictions about the 'correct' answer for the patient/client (An Overview of Clinical Applications of Artificial Intelligence 2018; Bathaee 2018). To the extent that this data is biased, incomplete or inaccurate, the AI can produce similarly biased results (An Overview of Clinical Applications of Artificial Intelligence 2018; Bathaee 2018). This can lead to decisions which can have a collective, disparate impact on certain groups of people even without the programmer's intention to discriminate (Lee et al. 2019).

One example is a recent study in a US hospital, that showed how the use of algorithms to identify primary care patients with the most complex needs (who would then be selected for the hospital's complex care program) discriminated against black patients (Obermeyer et al. 2019). The software attempted to predict patients' future health needs, but used their future health costs as a proxy for their health needs. Because Blacks generated lower cost due to structural inequalities in the health care system, they were less likely to be selected (Obermeyer et al. 2019). This example raises important policy questions about how we ensure data is representative so machine learning algorithms are generalizable, what mechanisms should be used to minimize discriminatory bias (e.g., antidiscrimination laws, consumer protection, industry standards), and what incentives should be in place to develop and adopt best practices? (Calo 2017).

The literature suggests several approaches to prevent algorithm discrimination. Industry standards can shape self-regulation, co-regulation and setting of regulatory requirements (OECD 2019; Lee et al. 2019). Ethical governance and standards can be used to clearly define the principles of 'fairness (OECD 2019). Building awareness of discriminatory practices (OECD 2019) and recruiting developers from diverse backgrounds permits representation of a range of populations (OECD 2019; Lee et al. 2019). Finally, simulation of predictions and using technical solutions to detect and correct algorithmic bias can be used before implementation (OECD 2019).

Many of these depend heavily upon the desire of the AI producers to ensure accuracy, rather than on the actions of regulators.

Another important policy issue arises from the lack of transparency with respect to the decisions made by deep learning technology. From a policy perspective, transparency focuses on how a decision is made, who participates in the process and the factors used to make the decision (OECD 2019). For example, some 'black box' machine learning models used in medical diagnosis are quite accurate at predicting the probability of a medical condition, but have been described as being too complex for humans to understand, which also means that errors are harder to detect (OECD 2019). There has been significant movement to make AI applications more explainable, but this can sacrifice accuracy if this requires reducing the variables to a set small enough for humans to understand (OECD 2019). In such cases, the potential

harms and benefits from these different types of models need to be weighed to see how we ensure that black-box algorithms are high quality and safe, and how much confidence we will place in treatment recommendations based on complex or 'black box' algorithms, particularly when new variables arise that may not be incorporated in that model.

In Europe, the General Data Protection Regulation (GDPR) provides individuals with the "right not to be subject to a decision based solely on automated means" (Nuffield Council on Bioethics 2018). The regulation also specifies that individuals should also be provided with meaningful information about how automated systems make their decisions (Nuffield Council on Bioethics 2018; Mowat Centre 2019). However, the scope and content of these restrictions—for example, whether and how AI can be intelligible—and how they will apply in the United Kingdom, remain uncertain and contested (Nuffield Council on Bioethics 2018). In Canada, the federal government has developed a set of guiding principles for the responsible use of AI and a Directive on Automated Decision-Making (Mowat Centre 2019).

4 Data Privacy and Consent

Because AI technologies involve the use of large datasets, there are also policy issues related to data privacy and consumer consent (Deane 2018). The expectations with respect to privacy varies around the world, particularly when these are anchored in cultural beliefs and moral judgments (Adler 1991). There are also differences in whether one is dealing with de-identified data that is used to construct the algorithms, or the personal data associated with an individual patient. A comparison of four commonly recognized healthcare privacy standards (Organisation for Economic Co-operation and Development Privacy Principles, Generally Accepted Privacy Principles, Personal Information Protection and Electronic Documents Act, Data Protection Act) indicates that all of these standards encompass principles that are premised on consent, collection, disclosure, access, security, quality, accountability, transparency, proportionality, notice and notification (Virtue and Rainey 2015).

A related set of policy issues relate to who is collecting (and using) the data. In some cases, regulations, policies and frameworks explicitly specify which entities are "covered" or "not covered" by these privacy rules. For example, under the US Health Insurance Portability and Accountability Act (HIPAA), physicians, health insurers, medical providers are "covered entities" while large companies such as Google, Apple are not. This means that a physician collecting a patient's data on heart rate will be subject to HIPAA but the same information collected by a private company such as Apple (e.g., via the Apple Watch), will not be (Price and Nicholson 2017). The EU is the only jurisdiction that has regulation via data protection legislation via the GDPR, which is applicable to all data regardless of who owns it (Forcier et al. 2019). The EU has also published new guidelines on developing ethical AI which include seven basic requirements; these include Privacy and Data Governance, which specifically guarantees privacy and data protection during the entire AI lifecycle (Commission

and Ethics Guidelines for Trustworthy AI 2019). In some instance, these regulations have been successful in addressing breaches in consent. For example, an AI program, Google DeepMind, was provided with patient records from Royal Free Hospital in the United Kingdom without patient consent. The information had sensitive information about HIV status, mental health history and abortion. The Royal Free argued during the trial that they had "implied consent" because the patients were aware that the app offered "direct care". The Information Commissioners Office (ICO) ruled that the deal was illegal but did not fine the hospital or Google (Duhigg 2012).

A related ethical issue with respect to privacy results if a predictive analytics model is used to create personal health information using information from individuals such as their location, purchase patterns, and/or internet access, without their consent or awareness (Deane 2018). In 2012, it was revealed that the Target stores in the US used big data and an AI algorithm to predict whether a customer was pregnant; the algorithm estimated due date based on the purchase habits associated with 25 products, and was used to send coupons for diapers and other pregnancy/parenting related coupons to these targeted consumers. When it was discovered that the enterprise was engaged in this activity, Target did not stop the practice but instead introduced additional random coupon offerings to the customer.(Reuters 2018) This was legal under HIPAA rules, because Target was not a "covered entity" as defined by the Act, but did present ethical issues related to consent, particularly if the consumers had not formally agreed to share their information with Target, and/or did not realize this information could be used to accurately predict a medical condition. This example also touches on personal data ownership and who owns it and how is it protected. For example, what would be th consequence if an employer discovered this information and discriminated against the individual by terminating their job, or if insurers changed coverage?

A related ethical issue that is relevant to the principles of consent, collection and disclosure and access is related to who is provided with the data? For example, there are examples of insurance companies that are moving towards interactive policy with "optional" fitness tracking in which refusing to participate in the voluntary program results in higher prices (framed in terms of not receiving discounts) (Caruana, et al. 2015). This example raises similar questions about what constitutes consumer consent, as well as what happens to the data. If AI data indicates that consumers are at high risk, their rates may rise, or they may become uninsurable. Can the data be deleted on the request of the consumer? Can the company use the information to predict clusters of high-risk consumers and adjust their rates? Should there be compensation to the consumers if their data is used by the insurer for economic gain? How do we prevent insurers from cherry-picking clients?

Many of these issues are not currently addressed in many privacy acts around the world. Given the cultural expectations with respect to privacy are locally driven, some policy analysts suggest that jurisdictions should develop their own local policy and regulatory framework, while others may propose more general frameworks. Issues that these frameworks would need to consider include which organizations would be included (e.g., health care providers? Insurers? Employers? Any organizations with health care data?), what mechanisms will be in place to ensure that product vendors

are creating AI applications that are aligned with privacy and consent rules and are complying with the policy and regulatory frameworks, and how consumers will be educated and informed by all data collecting organizations about how their data is being used.

5 Accountability

In most jurisdictions, there are regulatory structures in place to ensure that clinicians try to make accurate decisions. To the extent that clinicians receive the data from AI, they have some responsibility for evaluating its recommendations. However, the lack of transparency in 'black box' decision-making and its potential to cause medical errors may raise legal questions about what happens when a black-box AI system makes an erroneous diagnosis that results in harm to the patient? One study found that the use of machine learning to predict the risk of hospital attendants to develop pneumonia resulted in instructing physicians to send high-risk pneumonia patients home (Ardila et al. 2019). In this case, what happens if a patient dies because treatment was not provided? Who is legally responsible for this error? When should the responsibility be with the health care practitioner, health care organization, product vendor or the machine itself? Should this be a joint accountability? On the other hand, what are the implications for medical malpractice when a health care provider rejects diagnosis or recommendations from a machine?

The determination of liability regarding the use of the system and the user need further definition and clarification (Sullivan and Schweikart 2019; Reddy et al. 2019). Experts have offered possible solutions for current law or legal doctrines. One option for consideration is to implement AI personhood, which views the machine as an independent "person" under the law with duties who can then be sued directly for negligence claims (Sullivan and Schweikart 2019). In such instances, the AI system will be required to be insured and such claims will be paid out from the insurance. The second is to introduce common enterprise liability, which assigns responsibility to all groups involved in the use and implementation of the AI system (Sullivan and Schweikart 2019). The third solution is to modify the duties and standard of care of health care professionals using black-box AI that would require facilities and health care professionals to exercise due care in evaluating and implementing black-box algorithms (Sullivan and Schweikart 2019). Under this model, health care professionals are responsible for harm if they did not take adequate measures in properly evaluating the black-box AI technologies used in caring for the patient. Additional complications may arise if the AI is in a different jurisdiction, and hence not bound by the regulatory or legal requirements in place where the damage occurred.

6 Workforce Disruption

AI has the potential not only to be more accurate, but to work faster than humans. Several new studies have shown that computers can outperform physicians in cancer screenings and disease diagnoses (Rodriguez-Ruiz et al. 2019; Sharkey and Sharkey 2012). Others argue that AI can help streamline administrative processes, provide bots to help patients manage alone (e.g., reminding them to take their medicine), and better match patients with optimal treatment (Mesko 2017). There is a literature expressing concerns about whether AI will displace jobs for health care professionals by mastering tasks currently performed by people, and/or result in the employment of less skilled staff (SVayena et al. 2018). To the extent that AI is used to replace human contact, this may raise concerns (Secretary of State for Health and Social Care 2019). Others argue that this will free professionals from repetitive tasks and enable them to spend more of their time with patients (OECD 2019). Furthermore, AI is unlikely to have the capacity to understand emotions and show compassion, components that are foundational to the patient-health care professional relationship and heavily valued by patients and their families (Reddy et al. 2019). Given the potential impact to the workforce, it'll be important for governments to implement policies for managing this transition.

However, the fear of losing jobs can have implications for the adoption of AI by health care professionals. If there is a perception that health care professionals will be replaced, it is less likely that they will wish to adopt AI innovation. This raises ethical issues of whether medical establishments should be allowed to block AI technologies that are proven to be safer, better, or cheaper but may threaten jobs? Even if health care professionals adopt the 'black box' technology there is also the risk that reliance on a machine's decisions will reduce their skills or make them complacent, and might impact the patient-health care professional relationship if the clinician cannot explain the decision to the patient. This also raises concerns on the impact this will have on patient decision-making processes.

Another set of issues relate to who pays for these AI applications. To the extent that these applications are developed by for-profit industries seeking to maximize profits, there is a market for services provided directly to patients (and/or employers and insurers), many of which will not be covered by insurance. This category of applications is also less likely to undergo scrutiny by clinicians to assess their accuracy. At present, they may not be subject to regulatory processes. There is also the issue of who will pay for AI technology in health care organizations and physician offices, whether insurers would only pay for AI driven recommendations, and, if AI technology reduces the time spent by physicians to make treatment decisions, whether it should impact their compensation model.

As the industry develops AI applications, it will be important to maintain trust, which may require involving clinicians and patients in their design and development. Revision of professional standards and codes of conduct to accommodate changes from AI may also be required, as well as modification of education and training systems to skill and re-skill health care professionals to work in this new environment

(Dutton 2018). Policymakers will need to determine what AI technologies will be insured and funded. In addition, patient literacy with respect to the limitations of AI will also be important (Reddy et al. 2019).

7 Conclusions

The need for regulation of AI will continue to grow as more and more AI technologies are released in health care. Regulatory policy will need to balance the risk of stifling innovation by overregulation with the risk of harm caused by under-regulation. AI policy will need to focus on regulation that: monitors the accuracy of the recommendations proposed by the AI application, ensures that it is being used appropriately, minimizes bias and encourages transparency, assesses who is receiving the information and how it is being used, protects privacy through data protection requirements, enacts laws that clearly define accountabilities, establishes policies for labour disruption; implements professional standards and codes of conduct; adapts educational training to skill health care professionals; and determines what AI technologies will be insured and funded for clinicians. To the extent that these AI applications cross geographic boundaries, there are also questions about who will regulate them, and how. Development of regulation needs to be informed in conjunction with a diversity of stakeholders including product developers, researchers, patients, health care providers and policymakers.

As jurisdictions develop regulatory frameworks, it will be imperative that all stakeholders across sectors are engaged in the development and review of regulation and compliance requirements for new digital healthcare technologies.

References

Adler MJ (1991) Desires right & wrong: the ethics of enough. Macmillan Publishing Company

Allen B (2019) The Role of the FDA in ensuring the safety and efficacy of artificial intelligence software and devices. J Am Coll Radiol 16(2):208–210

An Overview of Clinical Applications of Artificial Intelligence (2018) CADTH, Ottawa (CADTH issues in emerging health technologies; issue 174). Retrieved from: https://www.cadth.ca/sites/default/files/pdf/eh0070_overview_clinical_applications_of_AI.pdf

Ardila D, Kiraly AP, Bharadwaj S, Choi B, Reicher JJ (2019) End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. Nat Med 25:954–961

Bathaee Y (2018) The artificial intelligence black box and the failure of intent and causation. Harvard J Law Technol 31(2). https://jolt.law.harvard.edu/assets/articlePDFs/v31/The-Artificial-Intelligence-Black-Box-and-the-Failure-of-Intent-and-Causation-Yavar-Bathaee.pdf. Accessed 13 Sep 2019

Calo R (2017) Artificial intelligence policy: a primer and roadmap. https://lawreview.law.ucdavis.edu/issues/51/2/Symposium/51-2_Calo.pdf. Accessed 01 Sept 2019

Caruana R et al (2015) Intelligible models for healthcare. In: Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining, pp 1721–30. http://people.dbmi.columbia.edu/noemie/papers/15kdd.pdf

Deane M (2018) AI and the future of privacy. https://towardsdatascience.com/ai-and-the-future-of-privacy-3d5f6552a7c4. Accessed 15 Oct 2019

Doern GB, Phidd RW (1992) Canadian public policy: Ideas, structure, process, 2nd edn. Nelson, Toronto, ON

Duhigg C (2012) How company learn your secrets. The New York Times Magazine. https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html. Accessed 12 Oct 2019

Dutton T (2018) An overview of national AI strategies. Retrieved from https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfdOECD (2019), Artificial Intelligence in Society, OECD Publishing, Paris. https://doi.org/10.1787/eedfee77-en

European Commission (2019) Ethics guidelines for trustworthy AI. (https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai). Accessed 12 Oct 2019

Forcier MB, Gallois H, Mullan S, Joly Y (2019) Integrating artificial intelligence into health care through data access: can the GDPR act as a beacon for policymakers? J Law Biosci, pp 317–335. https://doi.org/10.1093/jlb/lsz013

Jaremko JL, Azar M, Bromwich R, Lum A, Alicia Cheong LH, Gilbert M et al (2019) Canadian association of radiologists white paper on ethical and legal issues related to artificial intelligence in radiology. Can Assoc Radiol J 70(2):107–118. https://doi.org/10.1016/j.carj.2019.03.001. Epub 5 Apr 2019

Jenkins WI (1978) Policy analysis: a political and organizational perspective. Martin's Press, New York, St

Lee NT, Resnick P, Barton G (2019) Algorithmic bias detection mitigation: Best practices and policies to reduce consumer harm. https://www.brookings.edu/research/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/. Accessed 20 Sept 2019

Mesko B (2017) A guide to artificial intelligence in healthcare

Mowat Centre (2019) Governing the future: creating standards for artificial intelligence and algorithms. https://munkschool.utoronto.ca/mowatcentre/governing-the-future-creating-standards-for-artificial-intelligence-and-algorithms/. Accessed 12 Oct 2019

Nuffield Council on Bioethics (2018) Artificial intelligence in healthcare and research. https://nuffieldbioethics.org/wp-content/uploads/Artificial-Intelligence-AI-in-healthcare-and-research.pdf. Accessed 30 Oct 2019

Obermeyer Z, Powers B, Vogeli C, Mullainathan S (2019) Dissecting racial bias in an algorithm used to manage the health of populations. Science 366(6464):447–453. https://doi.org/10.1126/science.aax2342

OECD (2019) Artificial intelligence in society. OECD Publishing, Paris. https://doi.org/10.1787/eedfee77-en

Pesapane F, Volonté C, Codari, M, Sardanelli F(2018). Artificial intelligence as a medical device in radiology: ethical and regulatory issues in Europe and the United States. Insights Imag 9(5):745–753

Price W, Nicholson II (2017) Artificial intelligence in health care: applications and legal implications. SciTech Lawyer 14 (1). https://repository.law.umich.edu/cgi/viewcontent.cgi?article=2932&context=articles

Reddy S, Allan, S, Coghlan S, Cooper P (2019) A governance model for the application of AI in health Care. J Am Med Inf Assoc, pp 1–7. https://doi.org/10.1093/jamia/ocz192

Reuters T (2018) All John Hancock life insurance policies to include fitness incentives. CBC. https://www.vmmed.com/blog/all-john-hancock-life-insurance-policies-to-include-fitness-incentives/. Accessed 12 Oct 2019

Rodriguez-Ruiz A, Lång K, Gubern-Merida A, Broeders M, Gennaro G et al (2019) Stand-alone artificial intelligence for breast cancer detection in mammography: comparison with 101 radiologists. JNCI J Natl Cancer Inst 111(9):916–922

Scherer MU (2016). Regulating artificial intelligence systems: risks, challenges, competencies and strategies. Harvard J Law Technol 29(2). http://jolt.law.harvard.edu/articles/pdf/v29/29HarvJLTech353.pdf. Accessed 12 Nov 2019

Secretary of State for Health and Social Care (2019) Topol review: preparing the healthcare workforce to deliver the digital future. https://pharmafield.co.uk/healthcare/the-topol-review-preparing-the-healthcare-workforce-to-deliver-the-digital-future/. Accessed 15 Oct 2019

Sharkey A, Sharkey N (2012) Granny and the robots: ethical issues in robot care for the elderly. Ethics Inf Technol 14:27–40

Sullivan HR, Schweikart SJ (2019) Are current tort liability doctrines adequate for addressing injury caused by AI. AMA J Ethics 21(2): E160–166. https://doi.org/10.1001/amajethics.2019.160

SVayena E, Blasimme A, Cohen IG (2018) Machine learning in medicine: addressing ethical challenges. PLoS Med 15(11): e1002689. https://doi.org/10.1371/journal.pmed.1002689

Telegraph (2016, 4 May) Royal free breached UK data law in 1.6m patient deal with Googl's DeepMind. https://www.theguardian.com/technology/2016/may/04/google-deepmind-access-healthcare-data-patients

Thierer A, O'Sullivan A, Russel R (2017) Artificial intelligence and public policy. Mercatus Research Paper. Available via https://www.mercatus.org/system/files/thierer-artificial-intelligence-policy-mr-mercatus-v1.pdf

Virtue T, Rainey J (2015) HCISPP study guide. Syngress Publishing