



REVIEW

Ethics of AI in Pathology

Current Paradigms and Emerging Issues

Chhavi Chauhan* and Rama R. Gullapalli^{†‡}

From American Society of Investigative Pathology, Rockville,* Maryland; and the Departments of Pathology[†] and Chemical and Biological Engineering,[‡] University of New Mexico, Albuquerque, New Mexico

Accepted for publication
June 24, 2021.

Address correspondence to
Rama R. Gullapalli, M.D.,
Ph.D., University of New
Mexico, Room 333A, MSC08-
4640, Albuquerque, NM
87131. E-mail: rgullapalli@salud.unm.edu.

Deep learning has rapidly advanced artificial intelligence (AI) and algorithmic decision-making (ADM) paradigms, affecting many traditional fields of medicine, including pathology, which is a heavily data-centric specialty of medicine. The structured nature of pathology data repositories makes it highly attractive to AI researchers to train deep learning models to improve health care delivery. Additionally, there are enormous financial incentives driving adoption of AI and ADM due to promise of increased efficiency of the health care delivery process. AI, if used unethically, may exacerbate existing inequities of health care, especially if not implemented correctly. There is an urgent need to harness the vast power of AI in an ethically and morally justifiable manner. This review explores the key issues involving AI ethics in pathology. Issues related to ethical design of pathology AI studies and the potential risks associated with implementation of AI and ADM within the pathology workflow are discussed. Three key foundational principles of ethical AI: transparency, accountability, and governance, are described in the context of pathology. The future practice of pathology must be guided by these principles. Pathologists should be aware of the potential of AI to deliver superlative health care and the ethical pitfalls associated with it. Finally, pathologists must have a seat at the table to drive future implementation of ethical AI in the practice of pathology. (*Am J Pathol* 2021, ■: 1–11; <https://doi.org/10.1016/j.ajpath.2021.06.011>)

Health care is inherently data-centric, encompassing various data-generating subdomains such as insurance, pharmacy, administration, health care institutions, and different specialties of clinical practice.¹ Vast amounts of information are generated at each level of health care with the potential to provide unique insights into how medicine is practiced at scale.² Artificial intelligence (AI)-enabled clinical workflows have tremendously improved our ability to collect health care data.³ However, large-scale data analytics across these health care subdomains are lagging.^{4,5} Computational algorithms based on principles of machine learning and natural language processing are expected to automate big data analytics, identify patterns to improve our understanding of health care processes, and improve efficiencies of health care delivery.^{1,2,6}

Although data-generating sources within health care are vast, AI researchers tend to focus on the data generated in

the context of routine clinical work. Clinicians generate vast amounts of unstructured data (eg, clinical notes during patient encounters).¹ However, especially in the developed nations, clinicians rely heavily on radiology and pathology to guide the diagnosis, prognosis, therapeutics, and management of patients.^{7–10} Radiologists are adept in the use of technology (eg, computed tomography and magnetic resonance imaging), using it as a key driver of the practice of

Supported by the Centers of Biomedical Research Excellence (CoBRE) grant 1P20GM130422 from the National Institute of General Medical Sciences (NIGMS) (R.R.G.).

Both authors contributed equally to this work.

Disclosures: None declared.

This article is part of a mini-review series on the applications of artificial intelligence and deep learning in advancing research and diagnosis in pathology.

radiology.^{8–11} Technological innovation is playing an increasingly dominant role in the practice of pathology as well. Both radiology and pathology are image-intensive specialties that make extensive use of image data for patient care via specialist-generated interpretations. Although radiology is further along the path of digitization and management of medical images, pathology is increasingly moving along the same path.^{8,10–13} The heavy reliance on images and digitization makes these two health care specialties most attractive to AI researchers for testing emerging ideas of imaging AI research. Imaging AI algorithms have seen the greatest amount of research and advances over the past decade. The confluence of abundant imaging data, ever-increasing cheap and powerful computational capacity, and advancing algorithmic AI research make radiology and pathology prime targets for disruptive innovation of health care AI applications over the next decade.^{7,9,14,15}

The second arm of the practice of pathology (in addition to image-intensive anatomic pathology) is the area of clinical laboratory medicine. Automation in clinical laboratory medicine has been well underway for many decades, resulting in vastly improved efficiency in delivering patient test results. In emerging fields such as precision medicine, there is great interest in the use of genomics and other forms of *-omics* data for both diagnosis and prognosis, with the use of information at a molecular level.¹⁶ The new frontier of omics technologies is a true big data specialty with vast amounts of omics patient data generated in each encounter.^{17,18} The field of bioinformatics focuses on algorithmic computational methods to manage and interpret such omics data in various clinical settings. AI researchers are highly interested in using AI-based methods to understand omics data in the context of patient health care.^{16,19} Perhaps the ultimate challenge in the use of AI-enabled health care is to synthesize both imaging and genomics data from a patient to provide novel insights into clinical outcomes and management.^{17,18,20} Many such efforts are currently underway.

Although the potential of AI-based algorithms to effectively manage and interpret big data in health care is considerable, there are significant downsides to using such a powerful technology without the necessary ethical and moral safeguards.^{5,21–24} There is increasing unease with the unrestricted use of AI in health care, especially regarding ethical issues such as patient privacy, exacerbation of race and gender inequities, and patient safety outcomes. The broader field of AI ethics is focused on the use of AI technologies to ensure development in an ethically and morally appropriate manner to benefit society at large. Issues surrounding AI technology misuse have both common themes across specialties and also reflect more specific specialty-centric concerns. Thus, the development of AI ethical guidelines requires the participation of domain experts (eg, practicing pathologists) to develop specialty-centric guidelines for the ethical use of AI technologies.

Key participants in enabling ethical AI technologies include AI researchers, pathologists, clinicians, institutional administration, professional societies, and, finally, the patients themselves. The AI ethics paradigm and the participants involved in this interactive process of development are illustrated in Figure 1. Some of the key definitions associated with topics of AI and ethics discussed in this paper are listed in Table 1 to better inform the reader.

While writing this paper, an article was identified with a similar thematic focus examining the role of AI ethics in pathology and laboratory medicine.²⁵ The article by Jackson et al²⁵ discusses AI ethics from a traditional bioethicist's perspective, relating to the core principles of bioethics as laid down in the Belmont report (<https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report/index.html>, last accessed June 24, 2021). In contrast, the current review focuses primarily on AI ethics from the perspective of ongoing developments in the area of AI research and algorithmic decision-making (ADM), and discussing its potential impact on the future practice of pathology. Thus, these articles are complementary. Hopefully, this review will encourage the interested pathologists to become involved in this emerging area to help guide the future development of AI algorithms in the practice of pathology.

This review first discusses the issues of AI ethics from the perspective of an AI researcher interested in AI-enabled pathology. It then addresses issues of AI ethics and research pertinent to a practicing pathologist, including risks of AI to the practice of pathology. Finally, it discusses AI ethics issues relevant at a professional society or institutional level to guide the safe development and deployment of AI in pathology.

Ethical AI Study Designs in Pathology

This section reviews issues in AI pathology research based on the principles of ethical AI design. There are a multitude of imaging-based pathology AI studies underway, with many more to come, and it is worthwhile to remember that the key ethical issues involved must be reviewed before initiating such studies.

Inclusive AI Design and Bias

Pathology is highly data-centric, making use of both clinical and phenotypic (histomorphologic) data elements to enable the traditional practice of pathology. However, there is an increasing appreciation for the need to place a classical pathology expert diagnosis in a broader context by integrating additional patient data elements as part of routine diagnostic workflow (eg, molecular biomarker information at an individual and population level).^{26–28} There is also a need to include additional data elements such as lifestyle and socioeconomic data to improve research categories such

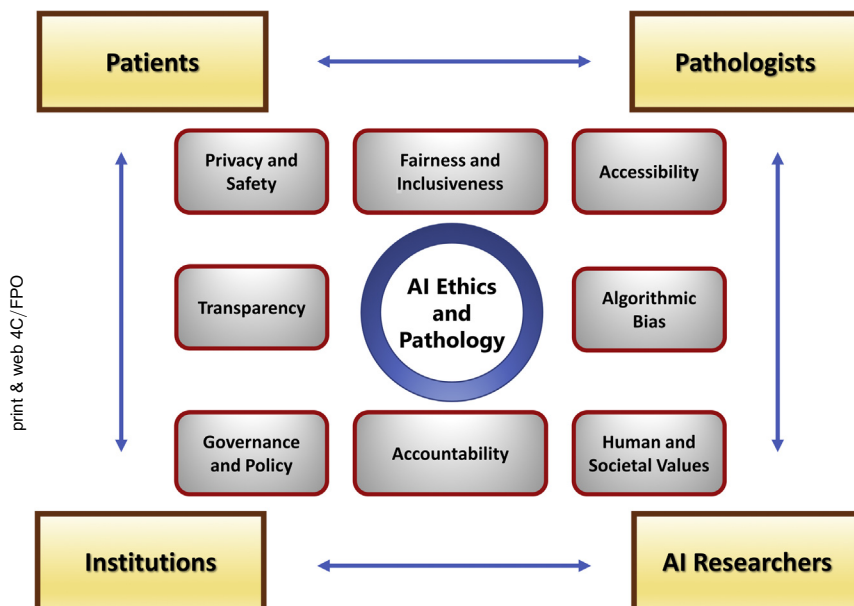


Figure 1 Key issues and participants involved in the study of artificial intelligence (AI) ethics in pathology. The various topics shown at the center of this diagram are the key emerging topics of significance to understand AI ethics in detail. These topics are common across various fields of study and also unique within the context of each specialty. Interactions between the various participants (patients, AI researchers, pathologists, and institutions) are central to developing a comprehensive framework of understanding AI ethics in pathology.

as cohort description, methods applied, and patient outcomes within the domains of pathology and precision medicine. Such complex, cross-domain data handling and integration are perfectly suited for the use of AI algorithms. A key strength of AI and ADM is their ability to enable integrative cross-domain analyses of diverse research data sets that are perhaps not amenable easily to a human mind. However, one must be careful while conducting such studies, as the complexity of these cross-domain data sets may lead to introduction of bias that can manifest in two major ways.^{29,30}

These biases include those inherent to the AI algorithm itself and biases arising within the data sets used for the purpose of training the AI algorithm. The former does not usually have ethical implications but rather is necessary to understand the inner-workings of an AI algorithm. For example, the popular *k*-means clustering algorithm works best with data form clusters that are roughly spherical and similar in size; however, sometimes they may not.³¹ This is not an ethical dilemma; it is simply a function of how the algorithm works. Similarly, deep learning (DL) networks typically have tunable variables (known as hyper-parameters) integral to the AI algorithm that must be assigned on an *ad hoc* basis by human beings as a part of the AI research study. These are conscious choices on the part of the AI researcher; however, in some cases, ingrained, and sometimes unconscious, biases may creep in with the potential to result in cascading downstream effects unforeseen by the AI researchers.

From an AI ethics point of view, the key critical issues relating to algorithmic bias occur mainly in the context of data sets used for a research study.^{29,30,32} Both sample choice and valuation play a role in this regard. For example, if a data set has category imbalance (ie, a study that is

composed mainly of adult white male subjects due to factors such as sample availability and socioeconomic factors of health care access), then the results of the AI algorithm trained on such data may not be accurate when implemented on the population as a whole. Results from such a homogeneous research study might inadvertently disadvantage a minority subpopulation. The second problem is related to underspecification (Table 1), which describes a phenomenon in which an AI-training data set is not provided with all of the necessary parameters. For example, if genetics of a population were a key factor in categorization of histology images, not including those details would lead to an incompletely trained AI model. Underspecification can thus lead to faulty correlations in predicting clinical outcomes. Another example: in trying to weigh the importance of various factors in determining the extent of disease, one might consider that a subpopulation that spends less on medical care might be a healthier one. However, it is just as likely, if not more so, that the subpopulation spending less on medical care is not healthier, but from a lower socioeconomic group, and cannot afford costly care. Issues such as these mandate a deep understanding on the part of the AI researcher seeking to train DL algorithms to improve the practice of pathology. It also behooves pathologists to be aware of such issues that may result in a skewed interpretation while using an AI algorithm in pathology practices.

In 2016, Arkansas approved the use of an algorithm-based program designed by InterRAI, a nonprofit coalition of health researchers from around the world, to determine the care hours needed by patients with limited mobility (The Verge, <https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy>, last accessed June 24, 2021). However, the AI algorithm was limited in its utility because it assigned variable scores for

Table 1 Brief Description of the Key Terms Used in This Review Explaining AI, Ethics, and Pathology

Terminology	Explanation
Artificial intelligence	Intelligence emulated or simulated by the use of technological means. Computational machinery is used to achieve intellectual autonomy and independence of thought similar to that seen in humans
Algorithm	A set of step-wise commands to accomplish a specific task/goal/objective. In AI, algorithms are the programming code that enable the functionality of an objective task and are key to emulating intelligence in an artificial manner
Bias	Discrimination in favor or against a set of outcomes in a particular setting. In AI ethics, this often refers to the ability of the AI algorithm to discriminate against individuals, groups, or populations based on the design of the original algorithm
Big data	Big data refers to data produced by an automated and repetitive technological process. Big data may be quantified in terms of the abundance of the data size generated. In pathology, some examples of big data are a digital pathology whole slide image repository, or databases containing complete blood counts across a population and time
Data privacy	The moral, legal, and ethical expectations to maintain confidentiality of data collected from either individuals or non-individual resources. In pathology, institutions responsible for the collection of patient laboratory data are tasked with the responsibility to ensure data privacy at individual and population levels
Data-shifts	A concept referring to the change in the data distribution between training and real-world data sets in AI algorithm development
Digital pathology	An emerging paradigm of pathology focused on digitization of traditional glass-based slides read by pathologists. Digitized slide data can be stored, viewed, and shared in real time, leading to enhanced efficiency of the sign-out process
Ethics	A branch of philosophy studying the concepts of right and wrong human behavior in a systematized manner
Machine learning	Computational algorithms that are capable of automated learning processes through iterative feedback of data without (or with minimal) human intervention
Underspecification	Failure to specify adequate details in the context of a training set of an AI algorithm

AI, artificial intelligence.

people with similar disabilities and made several erroneous decisions in calculating care hours needed, resulting in life-changing outcomes for hundreds of patients. There was no explanation offered to the patients, as the standards for use were neither clearly defined nor disclosed to all stakeholders. This precluded identification and rectification of these errors in a timely fashion. AI researchers must consider such high-level factors related to inclusivity and algorithmic bias while designing AI studies using pathology-based data.

Algorithmic bias, and statistical bias in general, is an ill-understood topic. It is perhaps unrealistic to expect a busy practicing pathologist to be well-versed in the various nuances of algorithmic and/or statistical bias. The solution to the issue of algorithmic bias may thus fall primarily on the shoulders of regulatory agencies [eg, the US Food and Drug Administration now identifies Software as a Medical Device (<https://www.fda.gov/medical-devices/digital-health-center-excellence/software-medical-device-samd>, last accessed June 24, 2021)] to approve sustainable AI-enabled workflows. However, pathologists must realize the potential of bias to manifest itself over long-term use of unregulated AI-enabled pathology workflows. AI researchers, as well as vendors, must partner with pathologists to obtain practical perspectives to identify and remedy potential sources of long-term bias in AI pathology algorithms. There are

multiple complementary efforts underway in professional pathology organizational committees [eg, Digital Pathology Association, the College of American Pathologists (CAP)] tasked with understanding the use of AI within the practice of pathology.

Race in Ethical AI Design

Another key variable to consider while designing AI workflows in pathology is the impact of race. AI model performance deterioration may occur due to issues related to data shifts, faulty correlations, and underspecification (Table 1) that limit the eventual utility of the AI algorithms in pathology classification studies. These limitations become even more striking in health care data incorporating complexities such as underlying physiological effects as well as genetic factors of disease predisposition. These complexities are further compounded when these data are analyzed without due consideration of race and ethnicity.

Although, historically, race has been considered to be a social construct without any biological basis, evidence suggests that it is tied to genetics.³³ Studies presented by AI researchers such as Joy Boulamwini and Timnit Gebru have made it evident that the AI systems and the defined parameters should be tested intersectionally with race, to determine their efficacy and broad utility beyond the use in

training data sets alone.³⁴ In light of the current global coronavirus disease 2019 (COVID-19) pandemic, the Centers for Disease Control and Prevention has placed emphasis on health equity considerations in racial and ethnic minorities (<https://www.cdc.gov/coronavirus/2019-ncov/community/health-equity/race-ethnicity.html>, last accessed June 24, 2021; and APM Research Lab, St. Paul, MN). The COVID Tracking Project confirms that COVID-19 has affected Black, Indigenous, Hispanic, and other minorities at higher rates (<https://covidtracking.com/race>, last accessed June 24, 2021). These recent findings make it imperative to include race and ethnicity data points in data sets used to train AI algorithms in the health care sector for broad applicability of any clinically facing AI model.

Optum, a subsidiary of insurance giant UnitedHealth Group, designed an application to identify high-risk patients with untreated chronic diseases. The automated categorization algorithm, however, was noted to discriminate against Black patients based on the cost of an individual patient's past treatments.³⁵ AI-based methods have been extensively deployed for cancer staging, especially in breast cancer. The US Food and Drug Administration has granted clearance to these applications without a need to publicly disclose how extensively their tools have been tested on people of color. Thus, AI-based applications have the potential to exacerbate disparities in clinical outcomes in breast cancer, a disease that is 46% more likely to be fatal for Black women.³⁶ This example provides an illustration of the potential of the cascading damaging effects a poorly researched and designed AI algorithm can have with real-world consequences on minority communities in particular.

Stakeholder Concerns with Regards to Consent and Awareness

In the absence of any defined guardrails and recommendations for the use of AI-based algorithms in clinical practice, issues related to patient consent and awareness have come to the forefront. Although, historically, there have been rigorous standards in place monitored by institutional review boards to protect patients with respect to data privacy, there seems to be a stunning lack of regulation and even basic guidelines on how to train an AI algorithm on available pathology data sets. For example, if a researcher trains an AI model on annotated data from a pathology slide archive and launches an independent AI company based on this model, do they own rights to all of the profits generated from it? How should the pathologists who worked on generating the annotation for each patient based on their expertise be compensated? This is an issue that has not yet been dealt with. Ethically, one would imagine that at least a part of the profit should be shared toward the continued maintenance of the pathology archive and/or support the department whose pathologists performed the initial annotation to generate the viable AI-based product.

Other pressing concerns also need a thorough discussion in the context of stakeholder concerns. For example, are patients made aware that an AI-based platform may have affected a pathologist's decision? Should the patients be informed about the use of AI-aided decision-making? Are the patients offered a choice to approve or reject the regimen based on their informed consent for the use of an AI-based model? Are the pathologists familiar with and do they completely understand the parameters and limitations of an AI-based algorithm? Are AI-based algorithms audited for bias, fairness, transparency, ethics, and risk mitigation, and if so, how often? Who are the governing bodies overseeing these audits?

In developing and underdeveloped countries, physicians often form informal consultation groups over various social media platforms and apps that often have loose privacy settings. These platforms and apps may allow data sharing to variable degrees within the platform/app as well as with third-party vendors. Such data-sharing practices put patient data safety and privacy at a high risk. There is often no regulation regarding how patient data are shared within an AI-enabled app that is often used without patient consent. How does one address concerns of de-identification of patient data and patient data privacy concerns in such cases? AI researchers must consider such questions from pathologist and patient perspectives while designing research studies before AI-enabled algorithms can become mainstream within the discipline of pathology.

Risks of AI in Pathology and to Pathologists — Real or Imagined?

Pathologists in current clinical practice are anxious to understand the scope and impact of AI algorithms. Questions such as “will AI replace pathologists?” and “how impactful is AI in enabling patient diagnoses?” are commonplace and reviewed often (with justified concern) by pathologists.^{37,38} This section examines some of these issues from a practicing pathologist's perspective with a focus on AI ethics and the important role pathologists have to play in the future of AI algorithm development.

An intensely studied aspect of AI ethics is risk assessment and evaluation of the dangers associated with implementation of AI. The popular press and social media are key drivers in fueling the imagination of the public about the future of AI, often in apocalyptic terms. Emerging AI-based technologies such as driverless cars, automated facial recognition, and AI-based deep fakes are indeed a cause for concern on the economic and moral outlook of our society as a whole.³⁹ However, in reality, AI is here to stay in one form or another with all of the attendant risks associated with it. A key requirement of DL, the latest iteration of AI, is the need for vast amounts of training data for eventual implementation. By training on large amounts of raw data, the algorithmic performance of DL workflows is much

better compared with that achieved by previous ADM approaches. In contrast to previous machine learning algorithms, DL algorithms are able to work on both structured (eg, laboratory data) and unstructured (eg, pathologist reports) data to create ADM models capable of attaining high levels of prediction accuracy.

A commonly held notion is that the laboratory-generated data influence 70% of all clinical decisions.⁴⁰ Although the accuracy of this claim is contested, it is undeniable that laboratory-generated data (both clinical and anatomic) comprise a significant portion of the quantitative data associated with a patient's electronic health record.⁴⁰ The ready availability of a significant amount of structured and unstructured clinical data in pathology archives and databases is highly attractive to AI researchers (and potential unscrupulous actors) to leverage laboratory-generated data for purposes of benefit (and harm), constituting a potential risk to patient safety. Pathologists, as custodians of laboratory data, will be at the center of heated debates on issues of patient data ownership for purposes of AI research in the future. The profession of pathology must be ready for this battle.

Assessing risk outcomes associated with AI research implementation is a key concern of AI ethicists and has attracted the attention of well-known luminaries such as Elon Musk. In fact, there are specific institutes dedicated to enabling such risk assessments and are populated by subject matter experts from different fields (eg, Machine Intelligence Research Institute, <https://intelligence.org>, last accessed June 24, 2021; Future of Humanity Institute, <https://www.fhi.ox.ac.uk>, last accessed June 24, 2021). As AI research into pathology accelerates, there is a significant opportunity for practicing pathologists to contribute toward such risk assessments to understand the broader implications of AI algorithms within pathology and patient management. Some of these risks (with a specific focus on pathology) are discussed here. We focus on two key issues related to AI risks in pathology: risk potential of AI to the pathologist workforce and risk impact of AI to the practice of pathology itself.

Underestimating the Risks of AI to Pathology

Research into medical applications of AI is proceeding rapidly. Radiology is perhaps at the forefront of many of these AI imaging initiatives. However, other case studies of AI health care applications such as AI-enabled natural language processing in clinical electronic health record text evaluation and AI-based analysis of whole slide images in digital pathology are also moving forward.^{1,41–43} Economic incentives of low-cost health care also drive commercial AI research efforts. Workforce automation is a key economic driver for many industry-driven AI initiatives to reduce the overall costs.

Traditionally, pathologists have relied heavily on experience and gestalt in rendering diagnoses. To the extent pathologists are unwilling to move on from simply rendering a

diagnosis based on visual impression, the risk of AI over the next decade to such practitioners is nonnegligible. AI techniques such as DL are beginning to outperform humans in certain image-based tasks, particularly those that involve a quantitation component. Pathologists must thus expand the scope of their practices and be more integrated within the overall clinical care of the patient. This includes adopting accessory techniques such as molecular, clinical, and epidemiologic data into providing comprehensive diagnoses for each patient. Future pathology practitioners must develop skills to synthesize information from multiple sources to provide integrated prognostic and even established high-level therapeutic National Comprehensive Cancer Network guidelines within their reports. Humans have the unique ability to synthesize information across different domains of knowledge with relative ease. In contrast, cross-domain integration is an ability that AI algorithms lack currently. It is this shortcoming of AI that human pathologists must capitalize on to improve the scope of pathology practice and stay uniquely relevant to the practice of medicine.

At the pathology practice level, it is also inevitable that the analytical instruments used by pathologists and laboratories will adopt increasing degrees of AI-enabled automation. Some examples of such AI-enabled automation would include providing first-pass oversight of the diagnostic algorithmic outputs, AI-based quality assurance of laboratory data, and automated assessments in high-volume clinical laboratory tests. In the majority of these instances, such developments would be driven primarily by the companies responsible for the instrument development. This may limit the impact an end-user pathologist or a laboratory has directly on the AI-enabled instrumentation development process. However, as an end-user customer, ideally, pathologists could significantly influence the adoption of such AI-enabled laboratory instrumentation into clinical practice by adopting due diligence for AI technologies and risk assessments based on principles of ethical AI. Thus, an awareness of issues surrounding algorithmic bias and ethical AI should be kept in the forefront while evaluating instruments and technologies in pathology practices. Active development of the guidelines of ethical AI and norms of assessment by professional organizations (eg, CAP) would raise awareness among pathologists that could help to mitigate the risks of adoption of AI-enabled technologies into clinical practice in the future. Pathologists and administrators must be aware of the subtle, unanticipated risks posed by AI algorithms in issues related to patient data privacy and the potential of AI-enabled technologies to (unconsciously) deviate from established pathology practice guidelines. The need for a wary and cautious eye on quality and process control by pathologists is unlikely to be automated away anytime soon.

Overestimating the Risks of AI to Pathology

AI has enabled some truly impressive advances in automation of narrowly defined tasks, particularly through the

use of DL-based approaches. However, it would be a stretch to say that AI in its current form is human-like in its capabilities. In fact, many experts are also of the opinion (perhaps pejoratively) that DL is no more than an extremely efficient statistical means to fit data. True human-like AI (also known as strong AI, AI-hard, or AI-complete), capable of emulating human-like awareness and decision-making capabilities, is unlikely to become a reality for decades to come.^{44,45} Within the context of pathology, an AI algorithm capable of replacing the skills of a highly trained human pathologist is also highly unlikely to come to pass anytime soon. Current DL techniques perform exceedingly well in addressing narrowly defined and well-formed questions in pathology with strict boundaries of performance.^{44,45} Pathologists must be wary of the hype and oversell commonly associated with AI research studies, while assessing the claims made by AI researchers. The hype associated with AI has been a well-known issue since the 1960s. AI research has passed through multiple boom and bust cycles when the ambitious goals of AI researchers failed to pan out.^{44,45} Although it feels as if the recent advances in computer hardware, networking, and data storage capacity may have allowed the field to turn a corner, DL techniques may still prove to be part of one such hype cycle playing itself out now. Although AI technologies are no doubt improving in each successive cycle of development, AI-driven technologies need to clear a very high bar before they are ready for widespread application in patient health care.

It is instructive to review the study by Frey and Osborne,⁴⁶ which assessed the impact of AI technology on workforce displacement in >700 professions in the United States. Physicians working in health care were the 15th least likely profession to be affected by AI automation, with an assigned low probability score of 0.0042. This finding is not unexpected, as the nature of work performed by physicians is complex, interactive, cross-domain, and multifaceted. Pathologists must thus actively seek skills and expertise to enable cross-domain relevancy in medicine. It is often discussed that, though AI may not replace pathologists, pathologists who know little about AI may be replaced in the future.

The laboratory workflow processes will likely see increased automation with incorporation of AI algorithms at various steps; however, the pace of change is expected to be at best incremental. Instrument vendors in health care are by nature cautious. Regulatory oversight by the US Food and Drug Administration will also ensure that evolution is gradual. Labor-intensive and uncompensated steps in laboratory workflow are the areas in which AI-enabled automation will likely be implemented first to improve overall process efficiencies. Equally, one may predict newer job opportunities within AI-enabled pathology to develop, deploy, and maintain these automated AI workflows in the future. However, the specifics of such development are not yet clear.

In summary, the profession of pathology and laboratory medicine must be well-informed of the potential risks

associated with AI. The only thing predictable about the AI-enabled future of pathology automation is that it is unpredictable. By adopting a proactive stance toward these technological developments within the AI space, pathologists can be at the forefront of mitigating the risks posed by AI while benefiting from its potential advantages. Awareness of issues of AI ethics will ensure that the balanced and informed viewpoint of pathologists is incorporated into the development of AI-enabled technologies in the field of pathology.

Institutional Frameworks to Enable Ethical AI in Pathology

Spending on health care worldwide was estimated at approximately \$9 trillion in 2014 and is projected to grow to approximately \$24 trillion by the year 2040.⁴⁷ In 2019, the United States spent approximately \$3.8 trillion, accounting for nearly 17.7% of its national gross domestic product. The United States also had the highest health care expenditures per person (approximately \$11,582) in the world, which is nearly double compared with that of the second most expensive country for health care (The American Medical Association, <https://www.ama-assn.org/about/research/trends-health-care-spending>, last accessed June 24, 2021). Naturally, there is an increasing push for automation to reduce health care costs. The economic incentives of lowering such costs are directly aligned with the potential of AI to help with this process.

Pathologists and laboratory medicine can expect to be squarely in the middle of the upcoming scramble to mine patient health care data at scale to enable AI workflows in health care. Although the ultimate value of such approaches remains to be determined, it will not impede the desire to acquire patient data for research and commercial purposes. The quantitative (clinical pathology) and (semi) structured (anatomic pathology) data formats are highly attractive to AI and machine learning researchers to assess the efficacy of AI algorithms in health care. An important emerging question pertains to the ownership of patient data. Ultimately, who owns patient data? Is it the patients themselves? Is it the institution where the data are held? Do pathologists who generate and curate the extensive data residing in the institutional databases have any intellectual property rights over the potential payoffs of collaborative AI algorithm development? Is it even ethical to consider patient data as something that is ownable with initiatives to mine it by using AI?

Conundrums such as these are likely to continue to confront pathology practice over the next decade. Professional organizations such as CAP and the American Society of Clinical Pathology have a key role to play in guiding the ethical development of AI in a manner that is appropriate to the practice of pathology. In the next subsections, a basic

framework of three key principles governing AI ethics within the specialty of pathology is described.

Transparency

A key element driving AI and ADM research ethics since the early 1960s is the idea of radical transparency.^{48,49} Algorithmic transparency may be defined as open availability of all information pertaining to the working of an AI algorithm. AI algorithmic transparency enables the interactive and relational assessment (moral and ethical) of the way an AI algorithm functions and, eventually, affects the end-user. Generally, transparency is defined as the robust and open availability of all information content related to the function and interactions within a system, both human and non-human.^{49,50} In health care—associated AI research, it is important to implement radical transparency due to the significant impact of ADM on a patient's health care.^{48,51} In pathology, an active area of research interest is the use of DL algorithms as agents to aid the anatomic pathologist's workflow. Questions such as "can DL AI algorithms identify a disease process as efficiently as a human pathologist?" underpin much of the AI research efforts into pathology. Ironically, however, much of the inner workings of DL algorithms remain unknown.³⁹ When training a DL algorithm on a set of images (anatomic pathology or otherwise) or any other data, the precise imaging data features used by a DL algorithm to identify and discriminate between different target categories is obscure. Thus, in essence, DL algorithms are currently a black box. In light of this, one must ponder when, if ever, we might be confident enough in trusting an ADM process with life-and-death decisions that are commonplace in the health care domain. Pathologists make such critical calls on a daily basis. Will we ever trust any such call made by an AI algorithm? This is an issue that remains to be answered. What is indisputable, however, is that AI and ADM algorithms exhibit an ever-increasing capability to identify and predict patterns to near human (or better) accuracies in a variety of fields.² It may well turn out that the performance characteristics of DL algorithms are so well established in the future that we may implicitly trust them without a detailed understanding of their inner workings. However, a transparent development process is the key to achieving general acceptance of the AI technologies in the pathology workflow.

Beyond the use of AI and ADM for mere diagnostic purposes, there are multiple pathology domains in which issues of AI transparency are likely to play an important role in the decade ahead. AI transparency will be critical while seeking to implement algorithmically enabled clinical workflow in a health institution. Although one may rely on the US Food and Drug Administration to provide overall regulatory supervision, it is institutions and pathologists who will be eventually responsible and liable for real-world outcomes. Thus, it is incumbent on the industry to provide transparency of the ADM process during the development

and implementation phases of an AI algorithm. Due to the potential for AI to constantly learn as a part of the workflow, developers and vendors must make two-way communication of the AI performance a routine part of their normal implementation protocols with end-users. As the AI and ADM protocols are upgraded iteratively, radical transparency must also be maintained to protect patient health.

Another scenario in which one can envision the need for transparency is in the data used to test and train as a part of the AI development process. AI trained on a local data set may not necessarily translate at a global implementation level. Transparency of such information will inform the pathologists of potential variability in the performance of the ADM and AI model in the local context and allow them to adjust accordingly. Equally, pathologists must be trained to recognize and deal with such issues. Over the next decade, as our understanding of the use and benefits of health care AI expands, our notions and expectations of AI and ADM transparency will also evolve in tandem.

Accountability

A natural corollary of transparency in AI is accountability. Accountability pertains to both human and non-human factors and their interactions alike. Additional details on this topic are provided in the article by Kroll.⁵² Current AI research initiatives in pathology focus on the eventual use of AI algorithms as an assistant in the normal pathology workflow. This may be attributed to the complexity of a normal pathology workflow process, and reflects the reluctance of AI researchers and industry to accept full accountability in the final decision-making process. Accountability is a shared transactional concept involving multiple entities such as the AI algorithm, the humans responsible for the health care decision process (eg, clinician and/or pathologists), and, finally, the institution implementing the AI-enabled workflow process. The eventual goal of a shared accountability process is to assign appropriate answerability as part of the normal health care delivery. Accountability may be either desired or undesired in the form of a reward for the beneficial outcome or blame for a non-beneficial outcome, respectively. However, the need for a formal accountability process in place when considering the implementation of AI-enabled workflows, including pathology is key. This includes formal documentation of accountability hierarchy at an institution and oversight detailing as to who is responsible for what and what outcomes are anticipated due to the implementation and use of algorithmic health care AI.^{49,51,53,54} In addition, periodic review and updating of the institutional AI accountability protocols are mandatory to reflect the current state of knowledge of the health care AI processes, which itself is constantly evolving. The black box nature of the neural networks was alluded to above (*Transparency*). In health care, in which patient safety is paramount, such a black box scenario of an AI algorithm mandates the need for

clearly defined hierarchies of accountability to ensure safe patient outcomes over the long term. In seeking to establish use of AI technologies within the field of pathology, one must establish the different tiers of accountability associated with the use of AI algorithms within an institution. This include the AI algorithm vendors, the human pathologists using AI technologies, and the institutions adopting AI themselves. Both physicians and institutions must be accountable for the use of AI in an ethical manner and share an equitable burden for the successful long-term use of AI/ADM in patient health care.^{49,51,53,54}

Governance

The third leg of the proposed framework for AI ethics in pathology relates to governance. The potential impact of AI and ADM tools creates enormous financial incentives for research and commercial interests in health care. AI has the potential to change the practice of health care delivery in the next decade, but it can also pose temptations for unethical entities to take advantage of shortcuts. Establishing and enforcing rules underpinning the governance of AI and policy will be critical to the moral and ethical implementation of AI and ADM in health care. Rules guiding the governance will thus need to be implemented at multiple scales: national, professional, and institutional. At a national level, the European Union has been proactive in instituting rules governing data ownership and privacy, in marked contrast to the practice in the United States, where regulation is highly lax. The rules implemented by the European Union in May 2018 are part of a framework known as the General Data Protection Regulation (<https://gdpr.eu>, last accessed June 24, 2021). This framework is directed toward countering the overwhelming power of the large tech giants that are at the leading edge of AI and ADM research. Similar policies are likely to be adopted globally over the next decade to ensure data privacy rights while harnessing the benefits of AI in a fair and egalitarian manner. More pertinent to the current review is the role of professional pathology societies in guiding the development and implementation of AI and ADM.

As domain experts, AI researchers, engineers, and scientists have been at the forefront of assessing the ethics aspects of AI and ADM and the wider impact of the technology. A majority of AI researchers and engineers are professionally affiliated with organizations such as the Institute for Electrical and Electronics Engineers, which has created a Global Initiative on Ethics of Autonomous and Intelligent Systems under the title Ethically Aligned Design (<https://ethicsinaction.ieee.org> and <https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e.pdf>, last accessed June 24, 2021). The science and ethics of AI are frequently represented at cutting-edge technical engineering conferences such as the Association for Advancement of AI and the ACM Conference on Fairness, Accountability and Transparency. However, there is increasing realization that

individuals (ie, engineers and AI researchers) who are involved in the development of the AI technology, also cannot simultaneously be the arbiters of the ethics of these technologies. Therefore, it will be necessary to incorporate the voices of the end-users of these technologies into the development process. For example, physicians may be well-positioned to understand the long-term impact of AI and ADM usage on patient outcomes within an AI-enabled clinical decision support system.^{29,42} Similarly, a radiologist may be better positioned to understand the weaknesses of an AI algorithm while incorporating clinical information into an AI-enhanced radiologic technique.¹⁴ Due to the interdisciplinary and transformative nature of AI and ADM technology, non-AI technical expert voices (physicians, lawyers, and lawmakers) are required to participate in building a framework of AI ethics to provide domain-specific expertise.^{5,55}

Our colleagues in radiology at the American College of Radiology have been at the forefront of launching AI-centric initiatives in the practice of radiology. ACR AI-LAB is an American College of Radiology initiative to educate radiologists in the use and implementation of AI tools.¹⁴ This course aims to empower radiologists with the basic knowledge to participate directly in the creation, validation, and use of health care AI. CAP also has similar early-phase initiatives within pathology AI. Through the creation of various AI and technology committees, the mandate of the CAP AI committee is to create a broad pathology-centric AI strategy. In addition, the committee aims to provide subject-matter expertise to CAP councils and enable the creation of AI laboratory standards in pathology. We propose that AI ethics must be a core component of the mandate of this committee moving forward. As more of AI and ADM technologies are incorporated into pathology workflows, pathologists must be at the table to guide the development of these technologies in a manner that aligns with the core principles of the profession of pathology. Finally, institutions employing pathologists (ie, universities and clinical practices) also must be actively engaged in helping to build a framework of AI ethics that is meaningful in a local context. This would involve initiatives such as the implementation of professional norms and of creating procedural guidelines for the adoption of ethical AI workflows within an institution. Effectiveness of ethical AI initiatives depends on the awareness of its importance, periodic review and oversight, and reinforcement of the saliency of this issue to employees (clinical and nonclinical health care workers) in general.

In summary, this review discusses three foundational frameworks to enable and guide ethical AI in specialty of pathology: transparency, accountability, and governance. These three core principles represent a starting point for adoption of AI- and ADM-based initiatives in pathology at an institutional level. As the science of AI evolves, pathologists must review and adopt additional measures to enable the ethical AI usage in pathology practice in a

manner that benefits patients. Also, as end-users of powerful data-centric technologies of AI and ADM and as custodians of structured patient data repositories within the field of medicine, pathologists are in a strong position to drive the adoption of ethical health care AI in the future.

Conclusions

AI research and the ethical implications of AI have become areas of great interest across various scientific fields, including health care. This review, highlights some of the ethical issues that a researcher needs to consider while conducting AI research in pathology. These include factors such as race, sex, and ethnicity, which play a key role in pathology AI research designs and outcomes. Multiple examples are described showing that neglecting these factors leads to a downstream exacerbation of existing inequities in health care delivery. This is particularly important in pathology, which serves as the big data repository of quantitative (and imaging) data measures of patient progress in the practice of medicine and has a central role in personalized patient care. Potential risk scenarios associated with the use of AI in the field of pathology are reviewed. Even though AI pathology is in its infancy, it behooves the profession to be aware of the potential risks posed by AI workflows to the practice of pathology. By improving awareness of such risks, pathologists can help guide the careful development of these technologies to benefit patients while minimizing potential downsides. Finally, three key foundational principles of AI ethics are discussed for the professional organizations to adapt to enable the development of AI within the field of pathology: transparency, accountability, and governance. This framework merely represents a starting point for the development of ethical AI in pathology at an institutional and organizational level. As the impact of AI-enabled workflows in pathology continues to increase over the next decade, more elements need to be added to the pathology AI ethics framework in accordance with the specific needs of the field of pathology. Pathologists have a critical role in enabling AI-based workflows in the laboratory and must have a seat at the table to guide the development and implementation of ethical AI and ADM within the practice of pathology. AI- and ADM-based workflows may create incredibly powerful new approaches for the practice of medicine. Pathologists must leverage this once-in-a-generation opportunity to be key drivers of this emerging paradigm shift within the practice of medicine.

Acknowledgment

We thank Dr. Stanley Cohen for critically reviewing select sections of this review.

References

1. Rajkomar A, Oren E, Chen K, Dai AM, Hajaj N, Hardt M, Liu PJ, Liu X, Marcus J, Sun M, Sundberg P, Yee H, Zhang K, Zhang Y, Flores G, Duggan GE, Irvine J, Le Q, Litsch K, Mossin A, Tansuwan J, Wang, Wexler J, Wilson J, Ludwig D, Volchenboum SL, Chou K, Pearson M, Madabushi S, Shah NH, Butte AJ, Howell MD, Cui C, Corrado GS, Dean J: Scalable and accurate deep learning with electronic health records. *NPJ Digit Med* 2018, 1:18
2. Rajkomar A, Dean J, Kohane I: Machine learning in medicine. *N Engl J Med* 2019, 380:1347–1358
3. Beam AL, Kohane IS: Big data and machine learning in health care. *JAMA* 2018, 319:1317–1318
4. Bukowski M, Farkas R, Beyan O, Moll L, Hahn H, Kiessling F, Schmitz-Rode T: Implementation of eHealth and AI integrated diagnostics with multidisciplinary digitized data: are we ready from an international perspective? *Eur Radiol* 2020, 30:5510–5524
5. Vayena E, Salathé M, Madoff LC, Brownstein JS: Ethical challenges of big data in public health. *PLoS Comput Biol* 2015, 11:e1003904
6. Cohen IG, Amarasingham R, Shah A, Xie B, Lo B: The legal and ethical concerns that arise from using complex predictive analytics in health care. *Health Aff (Millwood)* 2014, 33:1139–1147
7. Madabhushi A, Lee G: Image analysis and machine learning in digital pathology: challenges and opportunities. *Med Image Anal* 2016, 33:170–175
8. Niazi MKK, Parwani AV, Gurcan MN: Digital pathology and artificial intelligence. *Lancet Oncol* 2019, 20:e253–e261
9. Shen D, Wu G, Suk H-I: Deep learning in medical image analysis. *Annu Rev Biomed Eng* 2017, 19:221–248
10. Stenzinger A, Alber M, Allgäuer M, Jurmeister P, Bockmayr M, Budczies J, Lennerz J, Eschrich J, Kazdal D, Schirmacher P, Wagner AH, Tacke F, Capper D, Müller K-R, Klauschen F: Artificial intelligence and pathology: from principles to practice and future applications in histomorphology and molecular profiling. *Semin Cancer Biol* 2021, [Epub ahead of print] doi:10.1016/j.semcancer.2021.02.011
11. Miotto R, Wang F, Wang S, Jiang X, Dudley JT: Deep learning for healthcare: review, opportunities and challenges. *Brief Bioinform* 2018, 19:1236–1246
12. Abels E, Pantanowitz L, Aeffner F, Zarella MD, van der Laak J, Bui MM, Vemuri VN, Parwani AV, Gibbs J, Agosto-Arroyo E, Beck AH, Kozłowski C: Computational pathology definitions, best practices, and recommendations for regulatory guidance: a white paper from the Digital Pathology Association. *J Pathol* 2019, 249:286–294
13. Komura D, Ishikawa S: Machine learning methods for histopathological image analysis. *Comput Struct Biotechnol J* 2018, 16:34–42
14. Langlotz CP, Allen B, Erickson BJ, Kalpathy-Cramer J, Bigelow K, Cook TS, Flanders AE, Lungren MP, Mendelson DS, Rudie JD, Wang G, Kandarpa K: A roadmap for foundational research on artificial intelligence in medical imaging: from the 2018 NIH/R-SNA/ACR/The Academy Workshop. *Radiology* 2019, 291:781–791
15. Louis DN, Feldman M, Carter AB, Dighe AS, Pfeifer JD, Bry L, Almeida JS, Saltz J, Braun J, Tomaszewski JE, Gilbertson JR, Sinard JH, Gerber GK, Galli SJ, Golden JA, Becich MJ: Computational pathology: a path ahead. *Arch Pathol Lab Med* 2016, 140:41–50
16. Min S, Lee B, Yoon S: Deep learning in bioinformatics. *Brief Bioinform* 2017, 18:851–869
17. Diao JA, Wang JK, Chui WF, Mountain V, Gullapally SC, Srinivasan R, Mitchell RN, Glass B, Hoffman S, Rao SK, Maheshwari C, Lahiri A, Prakash A, McLoughlin R, Kerner JK, Resnick MB, Montalto MC, Khosla A, Wapinski IN, Beck AH, Elliott HL, Taylor-Weiner A: Human-interpretable image features derived from densely mapped cancer pathology slides predict diverse molecular phenotypes. *Nat Commun* 2021, 12:1613

18. Djuric U, Zadeh G, Aldape K, Diamandis P: Precision histology: how deep learning is poised to revitalize histomorphology for personalized cancer care. *NPJ Precis Oncol* 2017, 1:22
19. Castaneda C, Nalley K, Mannion C, Bhattacharyya P, Blake P, Pecora A, Goy A, Suh KS: Clinical decision support systems for improving diagnostic accuracy and achieving precision medicine. *J Clin Bioinforma* 2015, 5:4
20. Bhargava R, Madabhushi A: Emerging themes in image informatics and molecular analysis for digital pathology. *Annu Rev Biomed Eng* 2016, 18:387–412
21. Ananny M: Toward an ethics of algorithms. *Sci Technol Human Values* 2015, 41:93–117
22. Boddington P, O’Sullivan B, Wooldridge M: *Towards a Code of Ethics for Artificial Intelligence*. Cham, Switzerland, Springer, 2017
23. Brey PAE: Anticipatory ethics for emerging technologies. *Nano-Ethics* 2012, 6:1–13
24. Ienca M, Ferretti A, Hurst S, Puhani M, Lovis C, Vayena E: Considerations for ethics review of big data health research: a scoping review. *PLoS One* 2018, 13:e0204937
25. Jackson BR, Ye Y, Crawford JM, Becich MJ, Roy S, Botkin JR, de Baca ME, Pantanowitz L: The ethics of artificial intelligence in pathology and laboratory medicine: principles and practice. *Acad Pathol* 2021, 8:2374289521990784
26. Fu Y, Jung AW, Torne RV, Gonzalez S, Vöhringer H, Shmatko A, Yates LR, Jimenez-Linan M, Moore L, Gerstung M: Pan-cancer computational histopathology reveals mutations, tumor composition and prognosis. *Nature Cancer* 2020, 1:800–810
27. Kather JN, Heij LR, Grabsch HI, Loeffler C, Echle A, Muti HS, Krause J, Niehues JM, Sommer KAJ, Bankhead P, Kooreman LFS, Schulte JJ, Cipriani NA, Buelow RD, Boor P, Ortiz-Bruchle N-N, Hanby AM, Speirs V, Kochanny S, Patnaik A, Srisuwanakorn A, Brenner H, Hoffmeister M, van den Brandt PA, Jäger D, Trautwein C, Pearson AT, Luedde T: Pan-cancer image-based detection of clinically actionable genetic alterations. *Nat Cancer* 2020, 1:789–799
28. Moore DA, Young CA, Morris HT, Oien KA, Lee JL, Jones JL, Salto-Tellez M: Time for change: a new training programme for morpho-molecular pathologists? *J Clin Pathol* 2018, 71: 285–290
29. Gianfrancesco MA, Tamang S, Yazdany J, Schmajuk G: Potential biases in machine learning algorithms using electronic health record data. *JAMA Intern Med* 2018, 178: 1544–1547
30. Parikh RB, Teeple S, Navathe AS: Addressing bias in artificial intelligence in health care. *JAMA* 2019, 322:2377–2378
31. Cohen S: *The basics of machine learning: strategies and techniques*. Artificial Intelligence and Deep Learning in Pathology, Amsterdam, Elsevier, 2021. pp. 13–40
32. Rajkomar A, Hardt M, Howell MD, Corrado G, Chin MH: Ensuring fairness in machine learning to advance health equity. *Ann Intern Med* 2018, 169:866–872
33. Gannon M: Race is a social construct, scientists argue. *Sci Am* 2016, <https://www.scientificamerican.com/article/race-is-a-social-construct-scientists-argue/>; 2016. Accessed August 17, 2021
34. Buolamwini J, Gebru T: Gender shades: intersectional accuracy disparities in commercial gender classification. Proceedings of the 1st conference on fairness, accountability and transparency. Edited by Sorelle AF, Christo W. In Proceedings of Machine Learning Research: PMLR, 18; 2018. pp. 77–91
35. Obermeyer Z, Powers B, Vogeli C, Mullainathan S: Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 2019, 366:447–453
36. Pacilè S, Lopez J, Chone P, Bertinotti T, Grouin JM, Fillard P: Improving breast cancer detection accuracy of mammography with the concurrent use of an artificial intelligence tool. *Radiol Artif Intell* 2020, 2:e190208
37. Allen TC: Regulating artificial intelligence for a successful pathology future. *Arch Pathol Lab Med* 2019, 143:1175–1179
38. Jackups R Jr: Deep learning makes its way to the clinical laboratory. *Clin Chem* 2017, 63:1790–1791
39. O’Neil C: *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York, NY, Crown Publishing Group, 2016
40. Hallworth MJ: The ‘70% claim’: what is the evidence base? *Ann Clin Biochem* 2011, 48:487–488
41. Hoyt RE, Snider D, Thompson C, Mantravadi S: IBM Watson Analytics: automating visualization, descriptive, and predictive statistics. *JMIR Public Health Surveill* 2016, 2:e157
42. Miotto R, Li L, Kidd BA, Dudley JT: Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. *Sci Rep* 2016, 6:26094
43. Zarella MD, Bowman D, Aeffner F, Farahani N, Xthona A, Absar SF, Parwani A, Bui M, Hartman DJ: A practical guide to whole slide imaging: a white paper from the Digital Pathology Association. *Arch Pathol Lab Med* 2019, 143:222–234
44. Larson EJ: *The Myth of Artificial Intelligence: Why Computers Can’t Think the Way We Do*. Cambridge, MA, The Belknap Press of Harvard University Press, 2021
45. Mitchell M: *Artificial Intelligence: A Guide for Thinking Humans*. ed 1. New York, NY, Farrar, Straus and Giroux, 2019
46. Frey CB, Osborne MA: The future of employment: how susceptible are jobs to computerisation? *Technol Forecast Soc Change* 2017, 114: 254–280
47. Global Burden of Disease Health Financing Collaborator Network: Future and potential spending on health 2015–40: development assistance for health, and government, prepaid private, and out-of-pocket health spending in 184 countries. *Lancet* 2017, 389:2005–2030
48. Botkin JR: Transparency and choice in learning healthcare systems. *Learn Health Syst* 2018, 2:e10049
49. Diakopoulos N: Accountability in algorithmic decision making. *Commun ACM* 2016, 59:56–62
50. Diakopoulos N: *Transparency*. Edited by Oxford Handbook of Ethics and AI; 2020. pp. 197–213
51. Vayena E, Blasimme A, Cohen IG: Machine learning in medicine: addressing ethical challenges. *PLoS Med* 2018, 15:e1002689
52. Kroll JA: Accountability in computer systems In: *The Oxford Handbook of Ethics of AI*. Edited by Dubber MD, Pasquale F, Das S. In *The Oxford Handbook of Ethics of AI*. New York, NY: Oxford University Press, 2020. pp. 181–196
53. McCradden MD, Joshi S, Anderson JA, Mazwi M, Goldenberg A, Zlotnik Shaul R: Patient safety and quality improvement: ethical principles for a regulatory approach to bias in healthcare machine learning. *J Am Med Inform Assoc* 2020, 27:2024–2027
54. Wieringa M: What to account for when accounting for algorithms Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency; 2020. pp. 1–18
55. Vayena E, Tasioulas J: Adapting standards: ethical oversight of participant-led health research. *PLoS Med* 2013, 10:e1001402