


# The Ethics of Artificial Intelligence in Pathology and Laboratory Medicine: Principles and Practice

Brian R. Jackson, MD, MS<sup>1,2</sup> , Ye Ye, PhD<sup>3</sup>, James M. Crawford, MD, PhD<sup>4</sup>, Michael J. Becich, MD, PhD<sup>3</sup>, Somak Roy, MD<sup>5</sup>, Jeffrey R. Botkin, MD<sup>6</sup>, Monica E. de Baca, MD<sup>7</sup> and Liron Pantanowitz, MD, MHA<sup>8</sup>

## Abstract

Growing numbers of artificial intelligence applications are being developed and applied to pathology and laboratory medicine. These technologies introduce risks and benefits that must be assessed and managed through the lens of ethics. This article describes how long-standing principles of medical and scientific ethics can be applied to artificial intelligence using examples from pathology and laboratory medicine.

## Keywords

ethics, artificial intelligence, machine learning, algorithms, privacy, big data

Received October 1, 2020. Received revised November 24, 2020. Accepted for publication December 28, 2020.

## Introduction

Artificial intelligence (AI) is transforming society and health care. Expanded access to computing power and large digital data sets has created the ideal conditions for technologic and business innovation not seen since the industrial revolution. There is great enthusiasm for the potential of these AI tools to transform and improve health care.<sup>1</sup> This is reflected in efforts by pathology and laboratory medicine professionals both to enhance the practice of pathology and laboratory medicine and to advance medical knowledge based on the data that we generate.

Artificial intelligence encompasses a wide range of machine learning (ML), deep learning, and other analytic tools derived from statistics and computer science. In most cases, AI application developers use real-world data sets to “train” their applications to generate the desired output. Applications are ideally validated using separate real-world data sets to assess the accuracy and generalizability of the AI output. In pathology AI, for example, a training or validation data set might consist of digitized microscopic images together with the associated diagnoses as assessed by human expert pathologists.

There are growing concerns regarding unintended negative impacts of these technologies.<sup>2</sup> In some respects, AI applications have proliferated faster than social norms and regulations have been able to evolve in response to such innovation. Serious

<sup>1</sup> Department of Pathology, University of Utah School of Medicine, Salt Lake City, UT, USA

<sup>2</sup> ARUP Laboratories, Salt Lake City, UT, USA

<sup>3</sup> Department of Biomedical Informatics, University of Pittsburgh School of Medicine, Pittsburgh, PA, USA

<sup>4</sup> Department of Pathology and Laboratory Medicine, Donald and Barbara Zucker School of Medicine at Hofstra/Northwell, Hempstead, NY, USA

<sup>5</sup> Division of Pathology, Cincinnati Children’s Hospital Medical Center, Cincinnati, OH, USA

<sup>6</sup> Department of Pediatrics, University of Utah School of Medicine, Salt Lake City, UT, USA

<sup>7</sup> Pacific Pathology Partners, Seattle, WA, USA

<sup>8</sup> Department of Pathology, University of Michigan, Ann Arbor, MI, USA

## Corresponding Author:

Brian R. Jackson, Department of Pathology, University of Utah School of Medicine, 500 Chipeta Way, MC933, Salt Lake City, UT 84108, USA.

Email: brian.jackson@aruplab.com



**Table 1.** Ethical Principles Relevant to AI.

Core principle	Description
Patient autonomy (respect for persons)	Acknowledging patients have decision-making capacity Stewardship and use of personal health information Informed use of AI for clinical care
Beneficence and nonmaleficence	Realistic prospect of benefit for the patient Minimize potential patient harm
Justice	Equitable distribution of costs, risks, and benefits across populations
Scientific inquiry	Creating knowledge and sharing knowledge
Validation	Analytical validity, clinical utility, safety of intended use
Scientific integrity	Scientific design, adherence to research best practices, transparency, managing competing interests

Abbreviation: AI, artificial intelligence.

questions have been raised about privacy, safety, and fairness. In the case of health care, social and ethical expectations are particularly high, as reflected both in popular culture messaging and health care regulations. It is therefore important that we create systems, processes, and pipelines to ensure the ethical development and use of AI in health care.<sup>1,3</sup>

Pathology and laboratory medicine represent a large and important setting for employing health care AI. Clinical laboratories performing in vitro diagnostic tests (including histopathologic diagnosis) constitute one of the largest single sources of objective and structured patient-level data within the health care system.<sup>4-6</sup> Pathologists and other laboratory professionals have deep expertise in the science of both specimen analysis and the application of laboratory results to patient care. They are also accountable for the safe and reliable application of diagnostic testing. Lastly, pathologists and laboratory professionals are highly expert in managing systems, both strategically and operationally. They are experts in applying novel technologies for the delivery of safe, high-quality health care at a health system level. It is thus natural that pathologists and other laboratory professionals should play a leading role in pathology AI from both an innovation perspective and a stewardship/governance perspective. This principle applies not just to AI applied within a pathology setting, such as histopathological image analysis, but rather to all health care AI that relies on laboratory data. Examples include genetic and genomic analysis, clinical predictive analytics based on routine laboratory findings, and continuity-of-care applications. It is also important for laboratorians to understand that there are key differences between the medical and tech cultures. For example, technology companies tend to prioritize nimbleness as reflected in Facebook's former motto "Move fast and break things,"<sup>7</sup> whereas in health care, the premise is to be more cautious and proceed without doing harm. The tech industry prioritizes speed in bringing a minimum viable product to market with just enough features to satisfy early customers, whereas medicine and academia prefer products with validated claims based on published peer-reviewed data.

This article is intended to provide guidance to pathologists and other laboratory professionals for the ethical development, validation, and implementation of medical AI applications in

pathology and laboratory medicine. This includes, but is not limited to, the following:

- Stewardship of patient data
- Development of software applications
- Validation of applications for clinical use
- Scientific study and publication of AI applications
- Development of institutional policies and processes
- Management of external business relationships

The first portion of this document explains key ethical principles and illustrates how they apply to AI within pathology and laboratory medicine. The second portion of this document covers accountability related to personal, institutional, and commercial entities.

It is hoped that this document will inform development of organizational policies and procedures related to AI, including those related to engagement with external business partners. It is further hoped that the principles in this document may be of use to regulatory agencies as they consider new approaches to legal oversight of these technologies.

## Ethical Foundations

Ethics in medicine, scientific research, and computer science all have deep academic roots. This article does not attempt to provide a comprehensive history but rather a summary of these intellectual traditions. The foundational principles of medical ethics as articulated by Beauchamp and Childress are autonomy, beneficence, nonmaleficence, and justice<sup>8</sup> (Table 1). Autonomy means that each person has the right to make decisions related to their own body and their own personal information. Beneficence means that systems and actions need to have a reasonable expectation of benefit to the patient and of not harming the patient (nonmaleficence). Justice means that systems need to attend to the fair distribution of benefits and risks across all impacted populations.

Scientific ethics include the values of creating knowledge, sharing knowledge broadly and openly, and scientific integrity (rigor and honesty). The ethics of clinical research is a subset of medical ethics. The Belmont Report in 1979 articulated the basic principles of clinical research ethics in the United States,

those being respect for persons, beneficence, and justice.<sup>9</sup> Properly governed science contributes to respect for persons, beneficence, and justice, not just through minimizing risks of harm or injustice but also by expanding our knowledge regarding the impact of technologies on individuals and populations.

Computer science, with its subdomain of AI, is a much younger discipline than either medicine or the natural sciences, but nonetheless has a strong tradition of ethical thought.<sup>10</sup> An illustration of this can be found in recent statements on AI in the world economy from the Organisation for Economic Co-operation and Development and Group of Twenty. These statements address respect for human-centered values, explainability and transparency of AI application, robustness, and safety and security.<sup>11,12</sup> (Explainability refers to methods and techniques that allow humans to understand why an AI system generates its particular results.)

These ethical concepts are interdependent. The need for patient autonomy is greatest in precisely those aspects of an AI system that present greatest potential for harm. Similarly, justice is concerned with fair distribution of both benefits and harms. Thus, in order to effectively address autonomy and justice, we need to first have a deep understanding of the potential benefits and harms of a system. These in turn require scientific knowledge, explainability, and transparency. And finally, none of this works well without accountability. A number of ethical codes have been developed in medicine, science, and computer science, most of which emphasize professionalism as a mechanism to create ethical cultures that are enforced by individual practitioners and their professional societies.<sup>13-15</sup> In addition to holding individuals accountable for ethical actions, it is important that we also create robust accountability mechanisms for corporations and other organizations.

## Respect for Persons

Respect for persons, or autonomy, addresses an individual's ability to make decisions regarding what happens with both their physical body and their personal health information. Autonomy is sometimes, but not always, addressed through informed consent. With regard to AI, autonomy must be addressed at two levels: the personal health data used to develop, train, and validate AI systems, and the application of AI systems in patient care.

### *Use of Personally Identifiable Health Information for AI Development*

Most AI systems require large data sets for training and validation, and in medical settings, these typically come from aggregated personal health information. The initial collection and management of such information is integral to health care treatment, payment, and operations (TPO) and is outside the scope of this document. Repurposing personal health information for use in developing an AI system is a secondary use. Just because an individual has consented to data collection required for TPO

does not mean that they would agree to have their data used for other activities. For example, a 2017 survey of United Kingdom residents asked the question "How comfortable would you be with your personal medical information being used to improve healthcare?" These survey respondents were informed that it was impossible to guarantee 100% data security. Interestingly, 49% responded that they were either "not very comfortable" or "not comfortable at all."<sup>16</sup> Further, such preferences should not be assumed to be uniform across different populations. Another 2019 survey of patients in the US Veterans Health Administration (VHA) showed significant racial differences in desire to control access to their personal health information. More black patients than white patients wished to require consent for health record access by non-VHA providers.<sup>17</sup>

Most patients are largely or entirely unaware of common secondary uses of their health data and biological materials. In a focus group study by Botkin and colleagues, participants were provided a relatively detailed description of secondary uses including the potential benefits, risks, and safeguards. The study found that the large majority of participants were supportive of secondary uses if they were informed of the practices and had the opportunity to opt out.<sup>18</sup> It is possible, however, that patients may be less comfortable with purely commercial uses of their personal health information, such as for targeted marketing, in the absence of explicit patient benefit.<sup>19</sup> While in some respects these privacy considerations apply equally to all secondary uses of data, including AI development, quality improvement, and research activities, AI development can introduce additional risks related to the large number of patients in some data sets as well as the frequent involvement of outside technology companies.

In medical AI, health care data sets are commonly de-identified prior to use by developers. Indeed, under US Health Insurance Portability and Accountability Act (HIPAA) law, a data set stripped of certain identifiers (names, dates, etc) is no longer subject to HIPAA privacy restrictions.<sup>20</sup> This is not, however, always sufficient to satisfy the ethical requirement for patient autonomy. Some patients might expect transparency regarding the uses of their de-identified data, along with the ability to choose whether or not their data is included in those uses. Further, de-identification is not the same as permanent anonymization. Even when a de-identified record within a data set may not be identifiable by a human reader, it may be possible for software to reidentify the underlying individual by cross-referencing against other available data sets.<sup>21</sup>

A strict application of respect for persons might require informed consent on the part of patients before their personal information is used for training AI systems, whether de-identified or not. This could severely limit progress in medical AI due to the difficulty in obtaining large-volume, fully consented data and tracking what patients had consented to or not to regarding the use of their data. A second approach could be to provide an opt-out mechanism, such as the "right to be forgotten," within the European Union's General Data

Protection Regulation,<sup>22</sup> whereby an organization distributing or using an AI application could be required to remove an individual's personal information from a data set based upon a patient's request (regardless of whether that individual had previously provided consent for use of their data).<sup>22,23</sup> This approach is obviously dependent on public transparency regarding the algorithms and data sets used for their creation and a health organization's ability to remove nonconsented or de-consented patient data from their data sets.

Preserving strict patient autonomy in the case of individuals who do not wish to have their personal information used for algorithm development, then, raises significant technical and practical challenges. At a community level, medical record stewards (eg, health systems) as well as AI developers should have a high level of public transparency in order to inform the kinds of discussions necessary to balance autonomy with the desire for technologic progress.<sup>24</sup> Federated learning is a mechanism for training AI systems on patient data without actually sharing data itself and thus reduces the risk of a privacy violation.<sup>25</sup> Because of HIPAA concerns and intellectual property concerns, many academic centers partnering with commercial AI companies are favoring "bringing algorithms to the data, rather than data to the algorithm."<sup>26</sup> Hence, many researchers have begun developing and testing AI systems within their firewalls. In this way, it is easier for health care systems to control their own data. Going forward, it will be interesting to determine how institutions navigate ongoing sharing of their data to continually enhance adaptive algorithms.

Ultimately, individual patient control over whether their data are used for AI system development is something that must be taken seriously. It could be argued that patient control is particularly important in cases where there is high risk that patient information could be reidentified or otherwise lead to patient harm. The conditions for which the need for patient consent can be waived need clear articulation at the institutional level and within the health care industry.

### *Patient Autonomy Related to the Use of AI in Clinical Care*

A second domain for autonomy is making decisions about the use of AI systems in an individual's care. As AI becomes more widespread in health care, questions may arise about whether and when a patient might need to be informed about, and potentially consent to, the use of a particular AI application in their care.<sup>27,28</sup> For example, would a pathologist using AI to assist in the interpretation of a microscopic image be covered under the existing consent that allows the pathologist to do the interpretation? Or should the patient be informed about the extent to which software rather than the human pathologist is contributing to the diagnosis? Answering this question requires knowledge about the nature and extent of risks that are introduced by the software as opposed to by the human pathologist. The same issue applies when AI systems based on laboratory data are used by frontline clinicians. Does the practice of medicine require disclosure of when AI is included in the

medical decision-making process? Both from an ethical and medicolegal perspective, this question remains unanswered as there is essentially no case law on liability involving medical AI.<sup>29</sup> Although courts have traditionally deemed it impossible for machines to have legal liability, the practicing physician remains liable regardless of how medical decisions are made. Hence, it is the physician who must retain the ability to choose between a range of competing options, as might be presented by AI algorithms.<sup>30</sup>

## **Beneficence and Nonmaleficence**

In order for an AI application to satisfy this requirement, it must have a realistic prospect of benefit for the patients affected by the application while minimizing potential harm. "Affected" includes not only the patients whose medical care decisions are directly influenced by the AI system but also the patients whose data had been used to build or train the system.

Determining the benefits and harms of an AI system in a real-world setting is not a trivial task. Unlike mechanical medical devices with more easily testable relationships between inputs and outputs, AI systems tend to have highly complex and often nonlinear performance. Thorough analytical and clinical validation prior to implementation is thus critical, as is postimplementation monitoring. In the case of systems developed by an external vendor, the health system should have sufficient expertise in AI to critically review a vendor's validation and verification materials and in many cases should conduct independent verification.

## **Benefits of AI**

One general category of AI benefit is automation of certain types of human tasks. These tasks can then be performed at high levels of speed and reproducibility, thus augmenting and freeing human experts to focus on tasks that require human judgment. For example, a human pathologist might visually determine a microscopic region of interest on a slide (eg, focus of invasive carcinoma), and an AI application could compute both the area and the number of particular features within that region (eg, number of mitoses per mm<sup>2</sup>). AI systems can also identify potential pathologic patterns such as cancer within a microscopic image, which a pathologist can then review to confirm or refine the diagnosis.<sup>31</sup> In addition, AI can now successfully detect genotypic patterns (eg, RNA expression or microsatellite instability) from digitized slides of routinely stained tumors (ie, computational staining) without performing actual ancillary testing (eg, immunostaining or molecular assay) that might greatly benefit cancer diagnostics, prognostics, and prediction of therapeutic response.<sup>32,33</sup>

Artificial intelligence can also be used to discover patterns within data that may not be apparent to a human. One example is predicting survival or other outcomes based on AI image analysis.<sup>34</sup> In the laboratory medicine domain, AI applied to telemonitoring can help improve blood glucose levels in patients with diabetes.<sup>35</sup> When applied to whole populations,

ML can be used to understand interactions between health and social conditions and thus lead to more effective and efficient public health programs.<sup>36</sup> Artificial intelligence and big data (including social media data and online query data) have enabled better prediction of epidemics<sup>37</sup> and detection of food poisoning cases.<sup>38</sup> Artificial intelligence-enabled mobile health is providing yet another new way of collecting data and delivering tailored health interventions.<sup>39</sup>

The field of radiology has long pursued development of evidence to determine whether the performance of a radiologist is equal to, better, or possibly even worse without-or-with computer-aided detection of imaging findings and/or computer-aided diagnostic interpretation, so as to optimize the use cases for AI.<sup>40</sup> Such evidence is beginning to emerge in pathology as well.<sup>41</sup> Once an AI technology has been validated and demonstrated to improve outcomes (and assuming the existence of high-quality evidence to support this), it raises the question of whether it would be unethical not to apply it in clinical practice.

## Potential Harms

### *Harms Related to Privacy and Security*

Privacy and security, while not identical, are closely related. Breaches of privacy generally arise from unauthorized access to personally identifiable information (PII), while security refers to controls that prevent access by unauthorized individuals. The harms from failure in in these two domains largely overlap. The simple fact of PII, including protected health information (PHI), being accessed contrary to a person's wishes can be considered a harm by itself. Examples of secondary harms can include social stigma, negative career impacts due to disclosure of embarrassing information, loss of health insurance, and identity theft leading to financial harm. Privacy risks are not specific to AI development—they apply whenever health information is aggregated and used for other purposes as well—but are addressed here because AI development is expanding the need for large corpora of health information.

With regard to AI development and use, the first and most obvious way in which privacy can be compromised is a breach of a database containing personally identifiable health information. Unfortunately, these are more frequent and severe than one might expect.<sup>42</sup>

Privacy can also be compromised as a result of inadequate de-identification of medical records. Medical records including structured data, textual data, image data, and genomic data are a key resource for AI development. In the case of pathology AI, this includes sets of electronic pathology reports and digitized microscopic images, including whole slide images. In many cases, these data sets are stripped of identifiers prior to use by developers, particularly when the developer is a separate organization or company. A number of software tools have been developed for de-identification.<sup>43-45</sup> The radiology informatics community has likewise studied ways to de-identify medical images.<sup>46-48</sup> However, it is not always clear what

constitutes “adequate” de-identification. Part of the challenge is that personally identifying details can be found in many different places within highly heterogeneous medical records. In 2019, the University of Chicago Medical Center was sued (NB: the suit was later dismissed) for sharing hundreds of thousands of patients' records with Google, allegedly containing sensitive information such as dates of service scattered throughout records.<sup>49,50</sup> The value of de-identification is also muddied by the growing technologic capability to cross-reference across multiple large data sets in order to re-identify previously de-identified data.<sup>21,51</sup> The possibility that specific image content in pathology records is itself PII and may thereby itself represent a reidentifiable data bears scrutiny.<sup>52</sup> This issue also is under discussion in the field of medical imaging.<sup>53</sup>

### *Harms Related to the Application of AI Systems in Clinical Care*

The potential of AI to reduce medical errors and improve healthcare quality is dependent on the performance of individual applications and their effective deployment in clinical practice settings. Use of AI that produces erroneous or misleading results can be worse than not using that AI at all, particularly if trust in the technology leads users to bypass traditional human-based systems for diagnosis and treatment. It is important that the role of AI in producing or influencing diagnostic information be documented within medical records in order to allow investigation of potential problems.

The question is not just whether a system produces inaccurate results but also whether biases are present that could distort results in particular cases.<sup>54</sup> A classic example of hidden bias was observed in the development of ML models in the late 1990s to risk-stratify patients with pneumonia. The models were trained using real-world data from the University of Pittsburgh Medical Center.<sup>55</sup> Although all the models generally performed well, they misclassified a subset of patients (those with comorbid asthma) as being lower than average risk, when in fact they were at higher than average risk. On further investigation, it was conjectured that a likely cause may have been a heightened level of care being provided to pneumonia patients known to have asthma.<sup>56</sup> Thus, there appears to have been an unrecognized bias in the clinical data used to train the models.

Additional ethical risks can arise from implementation issues.<sup>57</sup> For example, AI systems may increase interpersonal distance between patient and physicians, thus interfering with effective communication.<sup>36</sup> It is also unclear how AI will impact trainees in academic medical centers. For example, if residents are permitted to rely on AI assistance, will they be as well trained as those trainees who worked without AI?

## Justice

Justice requires that AI systems support and promote equitable distribution of costs, risks, and benefits across different populations. Fairness applies to each of the phases of AI system

development, from assembly of data sets to training AI systems, to clinical use, to sales and distribution through business channels. These risks and benefits apply to data (eg, related to privacy), AI system performance, and downstream economics. Many of the issues that were raised in the previous section on beneficence and maleficence can also be seen as justice issues when viewed through the lens of populations rather than individuals.

When collecting data sets for training and validation purposes, both the risks and benefits associated with health data collection need to be shared equitably across different populations, giving attention to not disadvantaging marginalized groups. For example, Google was criticized in 2019 when some of their contractors were accused of photographing homeless people in exchange for \$5 gift cards in order to expand Google's facial image database of dark-skinned individuals. The targeting of homeless people in this case raised concerns about exploitation.<sup>58</sup> In many cases of health care AI development, previously existing data sets are used, for which patients were not asked for their permission, let alone offered compensation for use of their data. In a 2019 survey, US residents were asked about donating their genetic information for hypothetical research uses. Only 12% of respondents indicated willingness to provide it for purely altruistic reasons; most expressed willingness if they were financially compensated. Respondents also expressed a higher willingness to donate their data if governance policies provided them with more control over downstream uses, such as restricting data sales to third parties, the right to delete data, and being asked for permission for each specific future use.<sup>59</sup>

A second set of justice considerations comes into play at the point at which an AI system is applied in clinical operations. At this point, the concern is whether that system produces equitable clinical impact across populations. A system could very easily have high global accuracy (eg, as measured by area under a receiver operating characteristic curve) yet have poor or even negative accuracy for subgroups that represent a minority of the population.

Although some have hoped that AI could be more "objective" than human clinicians and thus reduce such bias, the dependency on real-world data to train AI systems makes this problematic. One reason is that historical medical data may be less representative for minorities or socially marginalized groups.<sup>60-63</sup> Using these models for clinical decision support may further increase disparity. For example, an algorithm-based screening tool that was widely used to alert primary care doctors to high-risk patients for resource allocation was found to systematically discriminate against Black patients.<sup>64</sup> It is thus important to conduct subgroup analysis to evaluate the performance across race, gender, and socioeconomic groups before implementing AI in a clinical setting.<sup>65</sup> This is particularly important in the case of AI techniques such as deep learning which result in poorly explainable "black-box" algorithms in which bias is not transparently identifiable.<sup>66</sup> A potential technical mechanism to address these issues is to build "fairness" measures into the training

function.<sup>67</sup> For example, in the resource allocation algorithm mentioned above, a fairness measure might be the difference between the proportions of white versus racial minority patients classified as high risk.

Artificial intelligence has particular potential to benefit low- and middle-income countries, where a combination of limited higher education resources and international brain drain can lead to critical shortages in human medical expertise. Ideally, the knowledge of frontline health workers could be augmented with AI to better serve these populations. This will not happen, however, unless it is a priority for developers of medical AI systems. It requires, for example, that system design and training data sets be appropriate for the environment of intended use, so that system performance can be valid for that setting.<sup>68</sup>

## Advancing Science

Scientific advancement can be seen as an ethical good. Artificial intelligence in pathology holds great promise in advancing our fundamental knowledge about disease processes so critical to cancer progression,<sup>32</sup> chronic disease evolution, and infectious disease management. Advancing science through deeper interrogation of both anatomic and clinical pathology data holds great promise in significantly improving our understanding of human disease so critical to the advancement of our discipline.<sup>69</sup> As pathology "big data" grows, harnessing the morphologic data from anatomic pathology, large volume of data from clinical pathology, genomic data from our molecular pathology laboratories, and microbiomic data from infectious diseases will create a significant opportunity to advance science and improve health. Moreover, the applicability of AI to pathology diagnostics may provide valuable insights into how the human pathologist performs diagnostic tasks and thereby offer opportunities for tutoring and possible performance improvement. Collectively, it is critical that these data be shared with the research community in a FAIR (findable, accessible, interoperable, and reusable) manner.<sup>70</sup> The role of employing AI in assisting pathology informatics with extracting knowledge from this big data is becoming a key role for modern pathology departments.

Scientific inquiry should furthermore be recognized as inextricably intertwined with the ethical issues discussed in this article. A corollary is that academic institutions and their faculty and staff have a central role to play in facilitating and promoting ethical AI. Improving accuracy, reducing bias, expanding utility, ensuring safety, increasing transparency, and encouraging accountability are all ethical responsibilities and are complementary to the pursuits of autonomy, beneficence, and justice. Scientific norms help to advance these interests. Perhaps the most fundamental scientific norm is research integrity. This includes use of truthful and verifiable methods in all stages of research and reporting as well as adherence to rules, regulations, guidelines, and commonly accepted professional codes and norms.<sup>71</sup>

Another norm is that research findings should be disseminated broadly in order to allow other scientists to critique and

**Table 2.** Mechanisms to Assure Adherence to Ethical Principles.

Organizational level	Examples
Individual accountability	Professional codes of ethics and codes of conduct
Organizational accountability	Institutional policies and procedures for conflict of interest, management of external business relationships, transparency to stakeholders
Regulatory accountability	Government regulation, regulation by professional entities

replicate the results. Pathology AI applications, regardless of whether they originate in an academic or commercial setting, should be made available for study by independent researchers who can share their findings in the form of articles and conference presentations.<sup>72</sup> This ethic of open publication can be in conflict with the intellectual property management practices common in the commercial software and biotechnology industries or in technology transfer companies established by academic institutions. However, this must not be allowed to create barriers to adequately studying pathology AI algorithms.

## Accountability Mechanisms

Achieving ethical processes and outcomes across the health care system requires well-developed systems for accountability at all applicable organizational levels (Table 2). In other words, there must be mechanisms to hold actors (both individuals and organizations) responsible for how they acquire data as well as develop, validate, implement, and use AI systems. This also requires that AI systems be developed in highly transparent and explainable ways, so that they can be tested and audited.

### Individual Accountability

Professional codes of ethics have historically been the primary mechanism to codify and enforce social norms within medicine and other professions. Most people are aware of the Hippocratic Oath (including its modern descendants such as the Declaration of Geneva<sup>13</sup> and the American Medical Association Code of Ethics<sup>14</sup>) that articulates the ethical obligations of physicians. The actual implementation of these codes is quite distributed, not only by medical boards but also by educational and training institutions as they indoctrinate students and trainees while occasionally expelling violators. A limitation of general medical professional codes is that while they address physicians' responsibilities to protect patients' privacy, none of them address AI-specific issues such as the use of de-identified data for research purposes as well as the exploration and potential commercialization of AI for biomedicine.

A number of health care professional societies have made steps toward filling this gap. In 2018, the American Medical Association (AMA) issued a policy report on clinical AI. It included the principles that clinical AI should reflect best clinical practices, be transparent and reproducible, avoid bias and

disparities, and protect patients' privacy.<sup>73</sup> The United Kingdom National Health Service (NHS) has likewise issued a Code of Conduct for data-driven health and care technologies, of which the principles include understanding users' needs, clearly defining the expected outcomes and benefits, lawful data processing, transparency, and evidence of safety and effectiveness.<sup>74</sup>

The American Medical Informatics Association has a code of ethics that describes the application of medical ethical principles to practitioners and researchers working in health care information technology settings.<sup>75</sup> In the specialty of radiology, several professional societies collaborated to produce a statement on ethics of AI in radiology.<sup>76</sup> It stated that AI usage should promote well-being, minimize harm, and ensure that the benefits and harms are distributed among stakeholders in a just manner. It also stated that developers of radiology AI systems are responsible to make them transparent and dependable and to minimize bias and that radiologists should still remain ultimately responsible for patient care in AI-enabled settings.

All individuals involved in the development of medical AI systems, regardless of whether they are clinicians, software engineers, or entrepreneurs, should have personal familiarity with medical ethics and its special application to AI. Their projects and activities should be subject to applicable institutional ethics accountability mechanisms. Ideally this will include an institutional review board (IRB) or equivalent oversight body. Such review boards can assess social or scientific value, scientific validity, fair selection of participants, acceptable risk–benefit ratio, informed consent, and consideration for participants' welfare and rights.<sup>77</sup> Currently, however, IRB oversight is limited to settings where AI activities are structured as research intended for scientific publication, as opposed to purely clinical, quality improvement or business activities.

Individuals using AI systems in clinical practice should be sufficiently familiar with the performance of these systems in order to be able to render a professional judgment on the risks and benefits to individual patients. Clinicians must speak up in the event of concerns regarding an AI system. Clinicians should also consider when and how to inform patients when a diagnosis is either made or heavily influenced by an algorithm.<sup>76</sup> Transparency of the derivation of diagnostic information is important for maintenance of a data trail for eventual quality assurance. It remains to be worked out, however, just how that information is best communicated.

### Organizational Accountability

Although medical ethics has historically been conceptualized mainly as a set of obligations for individuals, our modern world is increasingly driven by the actions of organizations. In health care, these organizations are often extremely large and include hospitals, medical device companies, pharmaceutical companies, and insurance companies. Some of these are for-profit, some are nonprofit, and some have a hybrid model with both nonprofit and for-profit components. Regardless of profit

motive, the ethical obligations of all organizations in this space are ultimately the same as the obligations on individuals. It is therefore necessary to articulate ethical responsibilities that exist at an organizational level, and not just at the level of individual actors. This includes organizational policies, procedures, and formal accountability processes to enforce these values both internally and with external business partners.<sup>78</sup>

At a macro level, there are currently two dominant external accountability mechanisms for corporations within market economies. The first is accountability to investors and other shareholders for financial performance, mediated by financial accounting standards and reporting. The second is governmental, that is, regulatory. This dyad can lead business leaders to believe that their ethical responsibilities come down to maximizing shareholder financial value while minimizing regulatory penalties. In recent decades, a third perspective has been championed in some academic business circles, namely that of stakeholder capitalism.<sup>79,80</sup> This perspective holds that companies can enhance long-term sustainability, including financial performance, by addressing the interests of nonowner stakeholders such as workers, customers, and the communities within which they operate. Because all of health care is under the umbrella of medical ethics, the single most important stakeholder for any health care company is the patient.

At an organizational governance level, stakeholder capitalism can be implemented in several ways. Some organizations have chosen to include social benefit in their articles of incorporation, which serve to set expectations and clarify legal obligations vis-à-vis shareholders. US law recognizes “benefit corporations” as one such category; some other countries have similar structures.<sup>81</sup> Even corporations that are explicitly for-profit can build strategies that use ethical behavior to build trust with customers and communities and thereby enhance business success. In 2019, the Business Roundtable, a group made up of CEOs of some of the world’s largest for-profit corporations, issued a statement that the purpose of business is to benefit all major stakeholders including customers, employees, suppliers, and communities in addition to shareholders.<sup>82</sup> Just as publicly held organizations issue regular reports of their financial performance, organizations (both public and private) should issue regular reports of social performance.<sup>83</sup>

An important consequence of stakeholder capitalism is that academic and clinical organizations can and should seek opportunities for ethical accountability in their dealings with external AI developer and supplier companies.<sup>84</sup> The first step is transparency in corporation activities. Particularly when collaborations involve sharing of medical records, the interest of public accountability requires public access to terms of the agreements. This runs counter to the secrecy that is common within much of the software industry. For example, a number of data sharing agreements between Google and several major health systems only came to light following the actions of whistleblowers because the contracts had been set up as confidential.<sup>85,86</sup> Another useful case study involved DeepMind’s agreements with the UK NHS. At the outset, DeepMind

**Table 3.** Summary of Principles for the Ethical Development and Use of AI in Pathology and Laboratory Medicine.

- 
1. Developers of AI systems should proactively inform patients and the public of how their data are collected and used to develop and validate their systems.
  2. Clinical organizations and developers should provide for informed individuals to control whether and how their personal data are used in the development of pathology AI systems.
  3. Developers, validators, and implementers of pathology AI systems should ensure that their systems provide measurable benefit to patients and/or populations, while minimizing risks and harms.
  4. Developers, validators, and implementers of pathology AI systems should ensure that both the development processes and the developed systems themselves promote fair treatment across all populations. This includes fair distribution of benefits, risks, and harms across all populations whose data are used for development of these systems, as well as those impacted by the use of such systems.
  5. Developers and implementers of pathology AI systems should ensure that their systems are sufficiently transparent and auditable to ensure that the above principles are being followed. They should also establish formal auditing processes and provide for public transparency of any findings (eg, through scientific publication).
  6. Developers, validators, and implementers of pathology AI systems should follow scientific norms of broad knowledge sharing and research integrity.
  7. Developers, validators, and implementers of pathology AI systems should establish formal oversight mechanisms, akin to institutional review boards, to ensure accountability to these ethical principles.
  8. Organizations engaged in developing, validating, implementing, selling, or purchasing pathology AI system should hold each other accountable to this set of ethical principles through formal mechanisms such as contracts. This includes requirements of transparency, auditability and auditing, and prohibitions against reidentification and other misuse.
- 

published these agreements openly on their website. After Google acquired DeepMind in 2014, though, they began to transfer DeepMind’s contracts to Google Health, which does not make its data sharing agreements available to the public.<sup>87</sup>

In addition to public transparency of contracts, organizations can build ethical expectations into those same contracts. Examples include prohibitions on attempts at reidentifying PHI, requirements to make systems available to academic researchers for audit and publication in academic journals, and consideration of public interest with regard to pricing and intellectual property rights.<sup>86</sup>

## Role of Regulatory Agencies

Government regulation is one important mechanism for enforcing ethical behavior. In the United States, for example, the FDA regulates some AI applications (but not others) as medical devices. The Health and Human Services Office of Civil Rights enforces rules related to patient information privacy and security. Regulatory agencies such as these will need to develop



new processes to address the patient safety and ethical issues introduced and/or complicated by AI.<sup>88,89</sup>

Ethics, however, should not be conflated with regulatory compliance. The process of developing regulations is slow and deliberative, and there can be a considerable time lag between the emergence of new technologies and comprehensive laws to govern them. In addition, creative individuals and organizations can often find technically legal mechanisms (ie, loopholes) for ethically questionable actions. Conflating ethics and regulatory compliance might actually give moral cover to loophole-seeking behavior.<sup>90,91</sup> A better approach is to treat regulatory compliance as a subset of ethical behavior. Individuals and organizations should also seek opportunities to engage with regulatory agencies, professional societies, and trade groups during the deliberation phase of designing new rules.<sup>92,93</sup> It is important that individuals and organizations consciously and explicitly represent the interests of patients and populations in these settings, rather than focusing mainly on their own parochial interests.

## Conclusion

Artificial intelligence is an increasingly powerful set of technologies with potential to advance diagnostic pathology and laboratory medicine for the benefit of patients. However, AI brings a complex mix of benefits, risks, and costs. Maximizing the benefits while minimizing risks and costs requires managing the technology within an ethical framework. Pathologists and other laboratory professionals, along with their clinical and academic organizations as well as potential industry partners, have an obligation to promote ethical AI development, validation, and implementation both within their own organizations and with external partners (Table 3).


## Declaration of Conflicting Interests

The author(s) declared the following potential conflicts of interest with respect to the research, authorship, and/or publication of this article: Liron Pantanowitz declares that he is on the medical advisory board of Ibex, an AI start-up company. Michael J. Becich is a founder and has equity in SpIntellx, a spatial analytics pathology AI start-up company with an explainable AI solution. In addition, Dr Becich has an industry sponsored research agreement with Owkin (French, UK, and US AI company) unrelated to work referenced in this manuscript.

## Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

## ORCID iD

Brian R. Jackson  <https://orcid.org/0000-0001-7902-3990>

## References

1. Matheny M, Israni ST, Ahmed M, Whicher D. *AI in Health Care: The Hope, the Hype, the Promise, the Peril*. National Academy of Medicine. Published 2019. Accessed November 22, 2020. <https://nam.edu/artificial-intelligence-special-publication>.
2. Vellido A. Societal issues concerning the application of artificial intelligence in medicine. *Kidney Dis (Basel)*. 2019;5:11-17. doi:10.1159/000492428
3. Wiens J, Saria S, Sendak M, et al. Do no harm: a roadmap for responsible machine learning for health care. *Nat Med*. 2019;25:1337-1340.
4. Hallworth MJ. The ‘70% claim’: what is the evidence base? *Ann Clin Biochem*. 2011;48(pt 6):487-488.
5. Becich MJ. Information management: moving from test results to clinical information. *Clin Leadersh Manag Rev*. 2000;14:296-300.
6. Forsman RW. Why is the laboratory an afterthought for managed care organizations? *Clin Chem*. 1996;42:813-816.
7. Teneja H. The era of “Move fast and break things” is over [digital article]. *Harv Bus Rev*. Published 2019. Accessed November 22, 2020. <https://hbr.org/2019/01/the-era-of-move-fast-and-break-things-is-over>
8. Beauchamp TL, Childress JF. *Principles of Biomedical Ethics*. 5th ed. Oxford University Press; 2001.
9. *The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research*. The Commission; 1978.
10. *ACM Code of Ethics and Professional Conduct*. Association for Computing Machinery. Published 2018. Accessed November 22, 2020. <https://www.acm.org/code-of-ethics>
11. *Recommendation of the Council on Artificial Intelligence*. Organization for Economic Co-operation and Development. Published 2019. Accessed November 22, 2020. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
12. *G20 Ministerial Statement on Trade and Digital Economy*. The Group of Twenty. Published 2019. Accessed November 22, 2020. <https://www.mofa.go.jp/files/000486596.pdf>
13. *Declaration of Geneva*. World Medical Association. Published 2017. Accessed November 22, 2020. <https://www.wma.net/policies-post/wma-declaration-of-geneva/>
14. *Code of Medical Ethics*. American Medical Association. Published 2019. Accessed November 22, 2020. <https://www.ama-assn.org/delivering-care/ethics/code-medical-ethics-overview>
15. A code of ethics to get scientists talking. *Nature*. 2018; 555:5. doi:10.1038/d41586-018-02516-x
16. Fenech M, Strukelj N, Buston O. Ethical, social and political challenges of artificial intelligence in health. Wellcome Trust. Published 2018. Accessed November 22, 2020. [http://futureadvocacy.com/wp-content/uploads/2018/04/1804\\_26\\_FA\\_ETHICS\\_08-DIGITAL.pdf](http://futureadvocacy.com/wp-content/uploads/2018/04/1804_26_FA_ETHICS_08-DIGITAL.pdf).
17. Turvey CL, Klein DM, Nazi KM, et al. Racial differences in patient consent policy preferences for electronic health information exchange. *J Am Med Inform Assoc*. 2020;27:717-725. doi:10.1093/jamia/ocaa012
18. Botkin JR, Rothwell E, Anderson R, Stark LA, Mitchell J. Public attitudes regarding the use of electronic health information and residual clinical tissues for research. *J Community Genet*. 2013;5:205-213.
19. Tully MP, Hassan L, Oswald M, Ainsworth J. Commercial use of health data—a public “trial” by citizens’ jury. *Learn Health Sys*. 2019;3:e10200. doi:10.1002.lrh2.10200

20. *Guidance Regarding Methods for De-identification of Protected Health Information in Accordance with the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule*. U.S. Department of Health and Human Services. Published 2015. Accessed November 22, 2020. <https://www.hhs.gov/hipaa/for-professionals/privacy/special-topics/de-identification/index.html>
21. Rocher L, Hendrickx JM, de Montjoye Y. Estimating the success of re-identifications in incomplete datasets using generative models. *Nat Commun*. 2019;10:3069. doi:10.1038/s41467-019-10933-3
22. *Right to Erasure*. United Kingdom Information Commissioner's Office. Published 2018. Accessed November 22, 2020. <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/individual-rights/right-to-erasure/>
23. Botkin JR. Transparency and choice in learning healthcare systems. *Learn Health Sys* 2017;2:e10049. doi:10.1002/lrh2.10049
24. Demuro PR, Petersen C. Managing privacy and data sharing through the use of health care information fiduciaries. In: Nøhr C, Kuziemsky C, Marcilly R, eds. *Context Sensitive Health Informatics: Sustainability in Dynamic Ecosystems*. IOS Press; 2019: 157-162.
25. Sheller MJ, Edwards B, Reina GA, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Sci Rep*. 2020;10:12598.
26. Kohli MD, Summers RM, Geis JR. Medical image data and datasets in the era of machine learning: Whitepaper from the 2016 C-MIMI Meeting dataset session. *J Digit Imaging*. 2017;30:392-399.
27. Chang K-C, Hsieh P-H, Wu M-Y, et al. Usefulness of machine learning-based detection and classification of cardiac arrhythmias with 12-lead electrocardiograms. *Can J Cardiol*. 2020;37:94-104. doi:10.1016/j.cjca.2020.02.096
28. Schiff D, Borenstein J. How should clinicians communicate with patients about the roles of artificially intelligent team members? *AMA J Ethics*. 2019;21:E138-145.
29. Price WN, Gerke S, Cohen IG. Potential liability for physicians using artificial intelligence. *J Am Med Assoc*. 2019;322:1765-1766. doi:10.1001/jama.2019.15064
30. Chung J, Zink A. Hey Watson—can I sue you for malpractice? Examining the liability of artificial intelligence in medicine. *Asia Pac J Health Law Ethics*. 2018;11:51-80.
31. Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*. 2017;542:115-118.
32. Schmauch B, Romagnoni A, Pronier E, et al. A deep learning model to predict RNA-Seq expression of tumours from whole slide images. *Nat Commun*. 2020;11:3877. doi:10.1038/s41467-020-17678-4
33. Montalto MC, Edwards RJ. And they said it couldn't be done: predicting known driver mutations from H&E slides. *Pathol Inform*. 2019;10:17. doi:10.4103/jpi.jpi\_91\_18
34. Courtiol P, Maussion C, Moarii M, et al. Deep learning-based classification of mesothelioma improves prediction of patient outcome. *Nat Med*. 2019;25:1519-1525.
35. Andres E, Meyer L, Zulfiqar A-A, et al. Telemonitoring in diabetes: evolution of concepts and technologies, with a focus on results of the more recent studies. *J Med Life*. 2019;12:203-214.
36. Blasimme A, Vayena E. The Ethics of AI in biomedical research, patient care and public health. In: Dubber MD, Das S, Pasquale F, eds. *Oxford Handbook of Ethics of Artificial Intelligence*. Oxford University Press; 2020.
37. Al-garadi MA, Khan MS, Varathan KD, Mujtaba G, Al-Kabsi AM. Using online social networks to track a pandemic: a systematic review. *J Biomed Inform*. 2016;62:1-11.
38. Harris JK, Hawkins JB, Nguyen L, et al. Using twitter to identify and respond to food poisoning: The Food Safety STL Project. *J Public Health Manag Pract*. 2017;23:577-580.
39. Wahl B, Cossy-Gantner A, Germann S, Schwalbe NR. Artificial intelligence (AI) and global health: how can AI contribute to health in resource-poor settings? *BMJ Global Health*. 2018;3:e000798.
40. Jorritsma W, Cnossen F, van Ooijen PMA. Improving the radiologist-CAD interaction: designing for appropriate trust. *Clin Radiol*. 2015;70:115-122.
41. Pantanowitz L, Hartman D, Qi Y, et al. Accuracy and efficiency of an artificial intelligence tool when counting breast mitoses. *Diagn Pathol*. 2020;15:80. doi:10.1186/s13000-020-00995-z
42. Jiang JX, Bai G. Types of information compromised in breaches of protected health information. *Ann Int Med*. 2020;172:159-160.
43. Thomas SM, Mamlin B, Schadow G, McDonald C. A successful technique for removing names in pathology reports using an augmented search and replace method. *Proc AMIA Symp*. 2002; 777-781.
44. Gupta D, Saul M, Gilbertson J. Evaluation of a deidentification (De-Id) software engine to share pathology reports and clinical documents for research. *Am J Clin Pathol*. 2004;121:176-186.
45. Beckwith BA, Mahaadevan R, Balis UJ, Kuo F. Development and evaluation of an open source software tool for deidentification of pathology reports. *BMC Med Inform Decis Mak*. 2006;6:12.
46. Vcelak P, Kryl M, Kratochvil M, Kleckova J. Identification and classification of DICOM files with burned-in text content. *Int J Med Inform*. 2019;126:128-137.
47. Monteiro E, Costa C, Oliveira JL. A machine learning methodology for medical imaging anonymization. *Annu Int Conf IEEE Eng Med Biol Sci*. 2015;1381-1384.
48. Tafti AP, Baghaie A, Assefi M, Arabnia HR, Yu Z, Peissig P. OCR as a service: an experimental evaluation of Google Docs OCR, Tesseract, ABBYY FineReader, and Transym. In: Bebis G, Boyle R, Parvin B, et al, eds. *Advances in Visual Computing*. Lecture Notes in Computer Science. Springer; 2016:10071.
49. Wakabayashi D. Google and the University of Chicago are sued over data sharing. *The New York Times*. Published June 26, 2019. Accessed November 22, 2020. <https://www.nytimes.com/2019/06/26/technology/google-university-chicago-data-sharing-law-suit.html>
50. Cohen IG, Mello MM. Big Data, Big Tech, and protecting patient privacy. *JAMA*. 2019;322:1141-1142.
51. Sweeney L. Only you, your doctor, and many others may know. *Technol Sci*. Published 2015. Accessed November 22, 2020. <https://techscience.org/a/2015092903>

52. Bukowski M, Farkas R, Beyan O, et al. Implementation of eHealth and AI integrated diagnostics with multidisciplinary digitized data: are we ready from an international perspective? *Eur Radiol.* 2020; 30:5510-5524. doi:10.1007/s00330-020-06874-x
53. Lotan E, Tschider C, Sodickson DK, et al. Medical imaging and privacy in the era of artificial intelligence: myth, fallacy, and the future. *J Am Coll Radiol.* 2020;17:1159-1162. doi:10.1016/j.jacr.2020.04.007
54. Parikh RB, Teeple S, Navathe AS. Addressing bias in artificial intelligence in health care. *JAMA.* 2019;322:2377-2378.
55. Cooper GF, Aliferis CF, Ambrosino R, et al. An evaluation of machine learning methods for predicting pneumonia mortality. *Artif Intell Med.* 1997;9:107-138.
56. Caruana R, Lou Y, Gehrke J, Koch P, Sturm M, Elhadad N. Intelligible models for healthcare: predicting pneumonia risk and hospital 30-day readmission. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* 2015:1721-1730.
57. DeCamp M, Lindvall C. Latent bias and the implementation of artificial intelligence in medicine. *J Am Med Inform Assoc.* 2020; ocaa094.
58. Otis GA, Dillon N. Google using dubious tactics to target people with 'darker skin' in facial recognition project: sources. *New York Daily News.* Published October 2, 2019. Accessed November 22, 2020. <https://www.nydailynews.com/news/national/ny-google-darker-skin-tones-facial-recognition-pixel-20191002-5vxpqgwknffnvbmy5eg7epsf34-story.html>
59. Briscoe F, Ajunwa I, Gaddis A, McCormick J. Evolving public views on the value of one's DNA and expectations for genomic database governance: results from a national survey. *PLoS One.* 2020;15:e0229044.
60. Zou J, Schiebinger L. AI can be sexist and racist—it's time to make it fair. *Nature.* 2018;559:324-326.
61. Raji ID, Buolamwini J. Actionable auditing: investigating the impact of publicly naming biased performance results of commercial AI products. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, Honolulu, HI, January 27-28, 2019. Association for Computing Machinery; 2019:429-435.
62. Leavy S. Gender bias in artificial intelligence: the need for diversity and gender theory in machine learning. *Proceedings of the 1st International Workshop on Gender Equality in Software Engineering.* Gothenburg, Sweden: Association for Computing Machinery. 2019:14-16.
63. Venkatraman V. Bias in the machines. *New Sci.* 2020;247:30.
64. 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.
65. Chen IY, Szolovits P, Ghassemi M. Can AI help reduce disparities in general medical and mental health care? *AMA J Ethics.* 2019; 21:167-179.
66. Price WN. Medical malpractice and black-box medicine. In: Cohen G, Gasser U, Vayena E, Lynch HF, eds. *Big Data, Health Law, and Bioethics.* Cambridge University Press; 2018.
67. Kearns M, Roth A. *The Ethical Algorithm.* Oxford University Press; 2020:57-93.
68. Mehta MC, Katz IT, Jha AK. Transforming global health with AI. *N Engl J Med.* 2020;382:791-793.
69. Louis DN, Feldman M, Carter AB, et al. Computational pathology: a path ahead. *Arch Pathol Lab Med.* 2016;140:41-50. doi:10.5858/arpa.2015-0093-SA
70. Wilkinson MD, Dumontier M, Aalbersberg IJ, et al. The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data.* 2016;3:160018. doi:10.1038/sdata.2016.18
71. National Academies of Sciences, Engineering and Medicine. *Fostering Integrity in Research.* The National Academies Press; 2017.
72. Van Calster B, Wynants L, Timmerman D, et al. Predictive analytics in health care: how can we know it works? *J Am Med Inform Assoc.* 2019;26:1651-1654.
73. American Medical Association. Augmented intelligence in health care. Published 2018. Accessed November 22, 2020. <https://www.ama-assn.org/system/files/2019-01/augmented-intelligence-policy-report.pdf>
74. UK Department of Health & Social Care. Code of conduct for data-driven health and care technology. Published 2021. Accessed January 21, 2021. <https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology/initial-code-of-conduct-for-data-driven-health-and-care-technology>
75. Petersen C, Berner ES, Embi PJ, et al. AMIA's code of professional and ethical conduct 2018. *J Am Med Inform Assoc.* 2018; 25:1579-1582.
76. Geis JR, Brady AP, Wu CC, et al. Ethics of artificial intelligence in radiology: summary of the joint European and North American multisociety statement. *Radiology.* 2019;293: 436-440.
77. Emanuel EJ, Wendler D, Grady C. What makes clinical research ethical? *JAMA.* 2000;283:2701-2711.
78. Silva HP, Lehoux P, Miller FA, Denis J-L. Introducing responsible innovation in health: a policy-oriented framework. *Health Res Policy Syst.* 2018;16:90.
79. Freeman RE. *Strategic Management: A Stakeholder Approach.* Cambridge University Press; 2010.
80. Kramer MR, Pfitzer MW. The ecosystem of shared value. *Harv Bus Rev.* October 2016:80-89.
81. Hiller JS. The Benefit Corporation and corporate social responsibility. *J Bus Ethics.* 2013;118:287-301.
82. Harrison JS, Phillips RA, Freeman RE. On the 2019 business roundtable "Statement on the Purpose of a Corporation". *J Manage.* 2020;46:1223-1237.
83. Lydenberg S, Rogers J, Wood D. From transparency to performance: industry-based sustainability reporting on key issues. *The Initiative for Responsible Investment at Harvard University.* Published 2010. Accessed November 22, 2020. [https://iri.hks.harvard.edu/files/iri/files/iri\\_transparency-to-performance.pdf](https://iri.hks.harvard.edu/files/iri/files/iri_transparency-to-performance.pdf)
84. Wachter RM, Cassel CK. Sharing health care data with digital giants: overcoming obstacles and reaping benefits while protecting patients. *JAMA.* 2020;323:507-508.

85. Copeland R, Mattioli D, Evans M. Inside Google's quest for millions of medical records. *Wall Street Journal (East Ed)*. January 11, 2020.
86. Robbins R. Contract offers unprecedented look at Google deal to obtain patient data from the University of California. *Stat News*. February 26, 2020. Accessed November 22, 2020. <https://www.statnews.com/2020/02/26/patient-data-contract-google-university-of-california/>
87. Vaughan A. Google is taking over DeepMind's NHS contracts—should we be worried? *New Sci*. September 27, 2019.
88. Allen TC. Regulating artificial intelligence for a successful pathology future. *Arch Pathol Lab Med*. 2019;143:1175-1179.
89. McCradden M, Joshi S, Anderson JA, et al. Patient safety and quality improvement: ethical principles for a regulatory approach to bias in healthcare machine learning. *J Am Med Inform Assoc*. 2020;ocaa085.
90. Parker CE, Rosen RE, Nielsen VL. The two faces of lawyers: professional ethics and business compliance with regulation. *Georget J Leg Ethics*. 2009;22:201-249.
91. Weeks EA. The ethical health lawyer. *J Law Med Ethics*. 2007; 35:320-324.
92. Abels E, Pantanowitz L, Aeffner F, et al. Computation pathology definitions, best practices, and recommendations for regulatory guidance: a white paper from the Digital Pathology Association. *J Pathol*. 2019;249:286-294.
93. Leavitt M, McKeown R. *Finding Allies, Building Alliances: 8 Elements that Bring—and Keep—People Together*. Jossey-Bass; 2013.