



The Social & Economic Implications of Artificial Intelligence Technologies in the Near-Term
July 7th, 2016; New York, NY
<http://artificialintelligencenow.com>

WORKSHOP PRIMER: HEALTHCARE & AI

Contents

[Brief description](#)

[From expertise to data: AI histories and rapid advancements](#)

[Five challenges for healthcare & AI](#)

[1. How will AI impact production of new health research?](#)

[2. How will AI impact diagnostics and healthcare delivery?](#)

[3. How will AI impact patient relationships to healthcare providers?](#)

[4. How will AI interact with the economics of healthcare?](#)

[5. How will AI impact professional ethics for health providers?](#)

[Questions to consider](#)

Brief description

Artificial intelligence (AI) technologies represent a new frontier for healthcare.¹ AI in the health domain sits at the intersection of multiple fields (including computer science, the biological sciences, medicine and health policy) and implicates a diverse range of stakeholders (such as patients, clinicians, researchers, insurance providers, pharmaceutical manufacturers, biotech companies, and governments). Applications of machine learning, robotics, and related fields are generating the potential for significant medical insights and increasing the efficacy and efficiency of healthcare.

While there is much positive potential, the future impacts of AI in healthcare are subject to the same longstanding social and economic issues that have shaped both clinical care

¹ As Russell and Norvig point out, the history of artificial intelligence has not produced a clear definition of AI but rather can be seen as variously emphasizing four possible goals: “systems that think like humans, systems that act like humans, systems that think rationally, systems that act rationally.” In the context of this primer on healthcare, we are relying on the emphasis proposed by Russell and Norvig, that of intelligence as rational action, and that “an intelligent agent takes the best possible action in a situation.” Stuart J. Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, Englewood Cliffs, NJ: Prentice Hall, 1995: 27.

and medical research.² These issues include increasing costs, differential health outcomes, and serious obstacles to patient access, among others. AI systems present opportunities to address many of these issues, but also have the potential to exacerbate old problems and to create new ones.

As AI systems change the norms and expectations of what constitutes care, the outcomes may be beneficial or may result in new harms. AI systems translate social and medical problems into technical solutions. This translation is far from neutral or without unintended consequences. What is lost, altered, or interposed in the process of translation? What feedback loops might be created? Who gets to decide which forms of care are most valuable or necessary?

In the US in particular, the introduction of AI systems is also subject to the substantial political and economic dynamics of health insurance, medical institutions, and for-profit healthcare companies like pharmaceutical or medical device manufacturers. How do these dynamics influence the design and regulation of new AI systems? How should professional ethics take these shifts into account? How can the complex range of incentives embedded into AI systems be made clear to patients and practitioners?

From expertise to data: AI histories & rapid advancements

The application of AI and “computer-aided diagnostics” more broadly has been an area of investigation in medicine since the 1960s. A motivating goal of these early applications was to use automated expert systems to assist physicians in making clinical decisions. These “expert systems” encoded decision trees and other forms of explicit knowledge, usually translated from interviews with domain experts.³ The aim of integrating codified expertise into diagnostic processes was to standardize clinical decisions and minimize forms of clinical bias, including anchoring bias (relying on the first information received), availability bias (relating current information to a recent or memorable diagnosis), and premature closure (failing to consider alternatives).⁴

Fifty years hence, the need to mitigate clinical bias and ensure more accurate diagnoses remains a central goal of integrating AI into medicine and healthcare. In comparison to previous expert systems which relied on explicit clinical knowledge, the latest waves of AI

² For example, significant scholarship has critically examined how dimensions of race are implicated in unequal health care, health outcomes, and medical research. See Dayna Bowen Matthew, *Just Medicine: A Cure for Racial Inequality in American Health Care*, New York: NYU Press, 2015; Brian Smedley, Adrienne Stith, and Alan Nelson, eds., *Unequal Treatment: Confronting Racial and Ethnic Disparities in Health Care*, Committee on Understanding and Eliminating Racial and Ethnic Disparities in Health Care, Board on Health Sciences Policy, Institute of Medicine, Washington, D.C.: National Academies Press, 2002; Alondra Nelson, *The Social Life of DNA: Race, Reparations, and Reconciliation After the Genome*, New York: Beacon Press, 2016.

³ For a critical perspective on the design and use of expert systems in the medical field, see Diana Forsythe, *Studying those who study us: An anthropologist in the world of artificial intelligence*, Stanford: Stanford University Press, 2001.

⁴ For a brief discussion of several kinds of biases, see Steven Dilsizian and Eliot Siegel, “Artificial Intelligence in Medicine and Cardiac Imaging: Harnessing Big Data and Advanced Computing to Provide Personalized Medical Diagnosis and Treatment,” *Current Cardiology Reports* 16, no. 441 (2014): 1-8, <https://www.ncbi.nlm.nih.gov/pubmed/24338557>.

technologies take advantage of the affordances of big data sets, using complex statistical modeling and machine learning to generate insight from existing data. This is in part because the potential data sources have substantially increased with electronic health records (EHRs) and clinical and insurance databases, as well as patient-produced data from consumer devices and apps.⁵ This wealth of data has positioned AI technologies as providing the potential to improve healthcare not only in the context of diagnostics but also to improve the production, organization, and communication of medical information.

Providers of medical care do not always have sufficient time to sift through, analyze, and apply this wealth of data, and may have little time to spend with patients to explain their decisions. AI is gaining traction in this new data-rich and often time-poor context of healthcare, particularly as a way to take advantage of the opportunities that accompany more fine-grained health data.⁶

AI systems rely on processing these large datasets to provide novel insights, and offer the opportunity to improve the accuracy of medical decision-making and reduce incorrect diagnoses. However, the primacy of data implicit in these paradigms may lead to unintended consequences. For instance, health providers may come under increasing pressure to prioritize statistical data over other forms of analysis and care. Or clinicians and care-workers may be expected to rely on AI-driven diagnoses and accept the underlying, often opaque models without being able to challenge the analysis or outcome.

In addition, as AI systems continue to be incorporated into healthcare, new technological norms are established, such as comprehensive tracking and surveillance of bodies and health data. The implications of such granular and ongoing data-tracking raise important questions regarding core health values such as confidentiality, autonomy, and informed consent. It also has serious implications for how we understand and value health labor, such as the difference between those who perform research, diagnostic work, and care work.

⁵ Electronic Health Records (EHRs) are quickly becoming a part of dynamic digital systems that hold multiple forms of patient information, from providers' notes and clinical assessments to lab test results and medication prescriptions. EHRs can also include non-patient specific information such as clinical guidelines, medical research literature, and scientific news. For an example of prediction models based on such data see N. Razavian, S. Blecker, A.M. Schmidt, A. Smith-McLallen, S. Nigam, D. Sontag, "Population-Level Prediction of Type 2 Diabetes using Claims Data and Analysis of Risk Factors" *Big Data: Data and Healthcare Special Issue* (January 2016), <http://online.liebertpub.com/doi/pdf/10.1089/big.2015.0020>.

⁶ See for instance Deep Mind's collaboration with the Royal Free Hospital London on Streams, an app to assist in the detection of acute kidney injury. Deep Mind Health, "Supporting direct patient care with clinical apps," DeepMind, 2016, <https://deepmind.com/health/clinical-apps>.

Five challenges for healthcare & AI

1. How will AI impact production of new health research?

The use of AI in medical research is in its infancy, but the insights drawn from AI and big data are already being heralded as a potential revolution in the field of medical research. One of the most high-profile successes has been the development of a new drug aimed at pancreatic cancer, BPM 31510. The drug, which has been described as the first drug to be developed through AI, was produced by BERG utilizing AI to process and analyze vast amounts of data, including biological data as well as clinical and patient data.⁷ Based on these analyses, BERG was able to generate the first complete model of how pancreatic cancer functions, and in turn focused on developing a drug that could prohibit the cancerous cells from metabolizing energy and growing.⁸ Phase I trials, which began in 2013, have shown positive results for the drug's safety and efficacy.⁹

The vast amounts of data that can be analyzed by AI systems are transforming not only how and how much biological or clinical data can be processed, but also how medical research itself can be conducted.¹⁰ Projects such as the Allen Institute's "Semantic Scholar" allow researchers to search and locate relevant research through AI-driven insights.¹¹ One of the potentially transformative aspects of IBM Watson's integration into clinical decision-making is the capacity for Watson to access, process, and analyze thousands of new and existing studies that clinicians are unlikely to have read, rendering clinical decision-making more informed by the latest research.¹² However, the use of proprietary systems like Watson may have implications for the openness of medical research, and the ways in which new research can be evaluated and peer-reviewed.

The integration of AI into medical research holds out exciting prospects for developing new treatments and offers the possibility of tailoring medicine to specific individuals. Still, this research is subject to the existing limitations and biases that may affect research outcomes, including incomplete datasets that exclude certain minority populations and financial incentives that favor the development of certain drugs over others. Medical research data often presents itself as objective and universal, while in reality the findings may be partial, temporary, or specific to only some communities.¹³ As AI systems learn

⁷ Liz Harley, "Taking on cancer with AI: Niven Narain, BERG," *Front Line Genomics*, April 22, 2016, <http://www.frontlinegenomics.com/interview/3983/3983/>.

⁸ Laura Lorenzetti, "This Company Wants to Cure Pancreatic Cancer Using AI," *Fortune*, April 22, 2016, <http://fortune.com/2016/04/22/berg-pancreatic-cancer-artificial-intelligence/>.

⁹ Ibid.

¹⁰ See for instance, Scott Spangler et al., "Automated Hypothesis Generation Based on Mining Scientific Literature," 20th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, August 24th - 27th, 2014, <http://dl.acm.org/citation.cfm?id=2623667>.

¹¹ "Semantic Scholar," Allen Institute for Artificial Intelligence, <http://allenai.org/semantic-scholar/>.

¹² Carl Zimmer, "Enlisting a Computer to Battle Cancers, One by One," *New York Times*, March 27, 2014, http://www.nytimes.com/2014/03/27/science/enlisting-a-computer-to-battle-cancers-one-by-one.html?_r=0

¹³ For example, research recently published in the journal *PNAS* has raised attention to issues with inferential statistical models and software bugs that likely invalidate many thousands of research papers within a particular field of fMRI-based

and contribute to medical research, awareness of these normative and structural limitations are needed to appropriately interpret the “new knowledge” that is generated.

2. How will AI impact diagnostics and healthcare delivery?

Automated and artificial intelligence systems present an opportunity to mitigate or even prevent misdiagnoses. This would be a substantial contribution to the field of healthcare; the harm of incorrect diagnoses are well-known and significant, including patient morbidity and mortality, increased length of hospital stay, unnecessary testing, and increased health care and malpractice costs.¹⁴ For instance, AI-based approaches have demonstrated great potential to reduce the harms and costs associated with hospital readmissions for heart failure.¹⁵

AI systems will also have limitations. When AI systems are presented as expert authorities understood to be free from error or bias, they may be subject to less scrutiny.¹⁶ But opportunities to introduce new errors via AI are myriad, including data entry and other points of translation between physical and digital records. This is especially relevant for the transition to electronic health records (EHRs), upon which many AI health systems are based. While these systems may introduce or extend existing incorrect models, the use of big data sets also has the potential to mitigate the biases that may be inherent in the existing constitution of randomized control trials (RCTs) or other public health databases.¹⁷

Despite a desire to make healthcare accessible and affordable to all, substantial evidence shows that access to healthcare and health outcomes are unequally distributed, with poor, non-white, and female populations often systematically disadvantaged.¹⁸ The introduction of AI will not automatically correct these systemic inequalities, and has the potential to amplify them.

The history of researching and treating heart disease is a case in point. Over the last decade, medical institutions such as the American Heart Association have acknowledged

research. Neuroskeptic blog, “False-Positive fMRI Hits The Mainstream,” *Discover Magazine*, July 7, 2016, <http://blogs.discovermagazine.com/neuroskeptic/2016/07/07/false-positive-fmri-mainstream/>.

¹⁴ Bradford Winters et. al., “Diagnostic errors in the intensive care unit: a systematic review of autopsy studies,” *BMJ Quality & Safety* vol 21 iss 11 (2012) <http://qualitysafety.bmj.com/content/21/11/894.abstract?sid=647c745e-4b1f-45e4-8883-a700adfc0005>.

¹⁵ Mohsen Bayati, Mark Braverman, Michael Gillam, Karen M. Mack, George Ruiz, Mark S. Smith and Eric Horvitz, Data-Driven “Decisions for Reducing Readmissions for Heart Failure: General Methodology and Case Study,” *PLoS One* vol 9 no. 10 (2014), <http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0109264>.

¹⁶ Nicholas Carr, “Automation Makes Us Dumb,” *Wall Street Journal*, Nov 21, 2014 <http://www.wsj.com/articles/automation-makes-us-dumb-1416589342>.

¹⁷ For an overview of the potential limitations of data generated from RCTs see Peter M Rothwell, “Factors That Can Affect the External Validity of Randomised Controlled Trials,” *PLoS Clin Trials* vol 1 iss 1 (May 2006), <http://journals.plos.org/plosclinicaltrials/article?id=10.1371/journal.pctr.0010009>.

¹⁸ Kaiser Family Foundation, “Health Coverage by Race and Ethnicity: The Potential Impact of the Affordable Care Act,” Report, Washington, D.C., 2013, <http://kaiserfamilyfoundation.files.wordpress.com/2013/04/8423.pdf>; John Ayanian, “The Costs of Racial Disparities in Health Care,” *Harvard Business Review*, October 1, 2015, <https://hbr.org/2015/10/the-costs-of-racial-disparities-in-health-care>.

that "most heart disease research is done in men, so how we categorize it is based on men. We need more science in women."¹⁹ Numerous studies have provided evidence that standards around heart disease, from education to diagnosis to medication, have been based on male bodies to the detriment of the diagnosis and treatment of heart disease in women.²⁰ If "normal" or "standard" are equated with only a subset of people, then those who fall outside this subset may not benefit from treatments based on those standards. Again, AI systems have the potential to enable appropriately specialized care.²¹ Still, as machine learning gains a more central role in establishing and identifying what illness is, additional attention must be paid to the faulty assumptions that may underlie constructions of "normal" or "average" health, constructions that are no doubt reflected in the data used to train AI systems.

Many AI systems currently being developed are intended to operate in the context of international health. Devices like Peek (Portable Eye Examination Kit), currently being tested in countries such as Kenya, create opportunities to access comprehensive eye exams without a doctor being physically present, expanding access to healthcare.²² The potential benefits of these systems are immense, but this does not mean that their risks can be overlooked. For instance, in the case of Peek or other remote and AI-supported diagnostics, how can the integration of these apps into existing healthcare contexts reinforce rather than inhibit local capacity-building in medicine?

3. How will AI impact the relationship between patients and healthcare providers?

The existing and potential applications of AI technologies have profound implications for the constitution of care-work and what it means to care for sick or vulnerable bodies. AI systems may be placed as mediators of care work, or may be tasked with supplanting caregivers entirely. In this way, AI systems may change the relationships and norms between patients and doctors or other caregivers.

Common examples used to illustrate the potentials of AI technologies to replace or supplement human caregivers range from robotic surgery to virtual avatars to companion robots.²³ Such examples have catalyzed important debates about the social implications of delegating care and companionship to non-human agents.²⁴ What kinds of equivalences are being made when machines replace, not just augment, human

¹⁹ EurekaAlert, "American Heart Association makes first ever statement on female heart attacks," American Association for the Advancement of Science, January 25, 2016, http://www.eurekaalert.org/pub_releases/2016-01/m-aha012516.php.

²⁰ Vidhi Doshi, "Why Doctors Still Misunderstand Heart Disease in Women," *The Atlantic Monthly*, <http://www.theatlantic.com/health/archive/2015/10/heart-disease-women/412495/>; Richard J McMurray et al. "Gender disparities in clinical decision making." *JAMA* 266 no 4 (1991): 559-562.

²¹ For example see Steven Dilsizian and Eliot Siegel, "Artificial Intelligence in Medicine and Cardiac Imaging: Harnessing Big Data and Advanced Computing to Provide Personalized Medical Diagnosis and Treatment," *Current Cardiology Reports* 16, no. 441 (2014): 1-8, <https://www.ncbi.nlm.nih.gov/pubmed/24338557>.

²² Peek, "Peek Vision: What It does," <http://www.peekvision.org/what-it-does>.

²³ Laurel Riek, "Robotics Technology in Mental Healthcare," in D. Luxton (Ed.), *Artificial Intelligence in Behavioral Health and Mental Health Care*. New York: Elsevier, 2015: 185-203.

²⁴ Nick Bilton, "Disruptions: Helper Robots Are Steered, Tentatively, to Care for the Aging," *New York Times* May 19, 2013. <http://bits.blogs.nytimes.com/2013/05/19/disruptions-helper-robots-are-steered-tentatively-to-elder-care/>

professionals? When we deem a machine capable of “care,” what capacities are we assuming, and what definition of “care” are we using? Are these equivalences in the best interests of patients?

An example that highlights the potential risks and benefits of employing AI-based technologies to interface directly with patients is a system like AICure. AICure uses facial recognition, motion detection, and AI analytics to monitor patients for drug adherence. The app monitors participants in clinical trials and those in clinical care to ensure that they are taking their medication at the right times and in the appropriate amounts. Patients using AICure administer their pill or medication in front of the camera on a their smartphone.²⁵ The app records this information, collecting and disseminating the data to doctors, caregivers, and/or researchers. An interface available to those at the receiving end enables close patient tracking, and the ability to contact a patient if adherence is not being followed.

This technology has the potential to improve health outcomes by encouraging better medication adherence, to improve the accuracy of clinical trials, and to reduce pharmaceutical companies’ revenue losses due to non-adherence. The attractiveness of AICure to large medical institutions and research hospitals is easy to see. Standardized data collection, and “proof” of adherence can be gathered much more concretely and cheaply, since there is no need for a human to witness and verify adherence. The time and skill it takes to develop trust between a patient and doctor or to monitor patients within the context of a human relationship is no longer required.

Still, the potential short-term and long-term downsides to this technology must be considered. What kinds of patients might not adapt well to app and smartphone-based systems? How might such an app exacerbate clinical stereotyping and the labeling of certain individuals and populations as “treatment resistant” or “non-compliant?” In turn, how might such technologies and the surveillance and tracking they enable facilitate unfair discrimination against particular populations? What would be the impact if patients receive adherence scores, similar to FICO scores, that would presume to reflect their worthiness to participate in research studies or receive medical care?²⁶ An app like AICure may inadvertently obscure the range of factors that contribute to a lack of adherence. In presenting a technological solution to one aspect of the treatment adherence problem, other social and cultural aspects, including environmental, economic, or psychological factors, may be overlooked.

AI systems are also being integrated into the modes of communication and courses of action available to patients. The field of health communication is shifting from a public health frame to a personalized frame, with information tailored, through the analysis of

²⁵ AiCure, “AiCure: Advanced Medication Adherence Solutions,” <https://www.aicure.com/>.

²⁶ Tara Parker-Pope, “Keeping score on how you take your medicine,” *New York Times* June 20, 2011, <http://well.blogs.nytimes.com/2011/06/20/keeping-score-on-how-you-take-your-medicine/>.

big data, to individuals using social media or devoted devices, including health coach avatars. For instance, ChronologyMD, a pilot application, allows patients with Crohn's disease to track and record their symptoms daily and also provides relevant reminders for sleep, medication, appointments and other forms of self-care.²⁷ Patients have reported better management of their disease resulting in improved quality of life, and providers were able to administer more comprehensive care because patients were more prepared with detailed information during appointments.

This application, and others like it, are marketed as tools of self-empowerment, and for some, they are.²⁸ However, additional implications need to be taken into account. Will the development of such applications effectively shift the responsibility for care and monitoring from healthcare professionals to patients themselves? What kinds of patients are favored in this new dynamic, and how do patients not well-equipped to manage and maintain their own data fall through the cracks? Moreover, how do such apps disrupt the locus of clinical expertise? What new roles and responsibilities do the designers and developers of such apps take on, and how do the ethical responsibilities at the heart of the medical profession get integrated into these differing design and engineering contexts?

4. How will AI interact with the economics of healthcare?

The economics of healthcare are already subject to increasing pressures from government actors, insurance companies, health institutions, pharmaceutical companies, medical device manufacturers, employers, and many others. With ongoing efforts to reduce costs and at the same time promote growth, many of these stakeholders have turned to large-scale data collection and now to AI to help sustain their economic models of research and care.²⁹ Many such efforts are focused on laudable goals, such as enabling more accurate clinical decisions or faster analysis of clinical data. AI systems also provide the possibility to improve preventative care and decrease costs. Still, the resources to develop and maintain these information technology systems within hospitals remain unstable and unequally distributed.³⁰ Moreover, because economic expectations for AI will depend heavily on making an increasing amount of patient data available for mining and learning across various devices, platforms, and networks, the implications for “surveillance capitalism” to transform fundamental structures within the health economy

²⁷ Linda Neuhauser, Gary L. Kreps, Kathleen Morrison, Marcos Athanasoulis, Nikolai Kirienko, Deryk Van Brunt, “Using Design Science and Artificial Intelligence to Improve Health Communication: ChronologyMD Case Example,” *Patient Education and Counseling* 92 (2013): 211-217, <https://www.ncbi.nlm.nih.gov/pubmed/23726219>.

²⁸ For scholarly work on the range of self-quantification and health tracking systems, see Dawn Nafus (ed.), *Quantified: Biosensing Technologies in Everyday Life*, Cambridge, MA: MIT Press, 2016.

²⁹ Claudia Grossmann et al., *Clinical Data as the Basic Staple of Health Learning: Creating and Protecting a Public Good*, Workshop Summary, Institute of Medicine of the National Academies, Washington DC: National Academies Press, 2010, <http://www.ncbi.nlm.nih.gov/books/NBK54296/>.

³⁰ For instance, see American Hospital Association, *Adopting Technological Innovation in Hospitals: Who Pays and Who Benefits*, Report. American Hospital Association, Washington, DC., 2006, <http://www.aha.org/content/2006/pdf/061031-adoptinghit.pdf>.

are profound.³¹ The economic emphasis on surveillance and consumption of sensitive health data will only increase with the recent moves to promote evidence-based medicine and the Affordable Care Act's (ACA) shift from a fee-for-service to a pay-for-performance model.³²

Insurers may also feel increased pressure to justify cross-subsidization models in the context of AI systems. Despite prohibitions in the Genetic Information Nondiscrimination Act of 2008, there is already growing interest in using genetic risk information for insurance stratification.³³ In fact, differential pricing has become one of the standard practices for data analytics vendors, introducing new avenues to perpetuate inequality.³⁴

Such efforts raise significant ethical questions. For example, how should AI systems that utilize this data take up questions of ethics, privacy, and potential discrimination in health care outcomes? What happens when one's insurer not only seeks to review your EHRs but also to query your FitBit and health agent AI about your status before setting the price of your premium? How will these dynamics influence the design and regulation of new AI systems?

5. How will AI impact professional ethics for health providers?

The use of AI in health contexts can also present challenges to core values upheld within the ethical codes of medical professionals, such as confidentiality, dignity, continuity of care, avoiding conflicts of interest, and informed consent.³⁵ Patient privacy and confidentiality of care has emerged as a primary concern in the adoption of new medical technologies.³⁶ As noted above, the shift to ubiquitous tracking and self-surveillance through "smart" devices and apps may often push the limits of existing privacy protections, such as the Health Insurance Portability and Accountability Act (HIPAA).³⁷ For example, computer scientist Latanya Sweeney has demonstrated that the vast majority of Americans can be identified using only three pieces of perceived "anonymized data:" ZIP code, birthdate, and sex.³⁸ Risks that patients will be reidentified from granular-level

³¹ Shoshana Zuboff, 'Big other: surveillance capitalism and the prospects of an information civilization', *Journal of Information Technology*, Vol. 30, Issue 1 (March 2015): 75-89, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2594754.

³² David Blumenthal, Melinda Abrams, and Rachel Nuzum, "The Affordable Care Act at 5 Years." *New England Journal of Medicine* Vol. 372, Issue 25, (2015), 2453, <http://www.nejm.org/doi/full/10.1056/NEJMhpr1503614#t=article>.

³³ Yann Joly et al., "Life Insurance: Genomic Stratification and Risk Classification." *European Journal of Human Genetics* 22 No. 5 (May 2014): 575-79, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3992580/>.

³⁴ Jason Furman and Tim Simcoe, "The Economics of Big Data and Differential Pricing," The White House Blog, February 6, 2015, <https://www.whitehouse.gov/blog/2015/02/06/economics-big-data-and-differential-pricing>.

³⁵ American Medical Association, Code of Medical Ethics, "Opinion 10.01 - Fundamental Elements of the Patient-Physician Relationship", <http://www.ama-assn.org/ama/pub/physician-resources/medical-ethics/code-medical-ethics/opinion1001.page>.

³⁶ Helen Nissenbaum and Heather Patterson, "Biosensing in Context: Health Privacy in a Connected World," in Dawn Nafus (ed.), *Quantified: Biosensing Technologies in Everyday Life*, Cambridge, MA: MIT Press, 2016.

³⁷ Paul Ohm, "Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization," *UCLA Law Review* 57 (2010): 1701-1778, <http://ssrn.com/abstract=1450006>.

³⁸ Latanya Sweeney, "Simple Demographics Often Identify People Uniquely," Carnegie Mellon University, Data Privacy Working Paper 3, Pittsburgh: PA, 2000. See also, Bradley Malin, Latanya Sweeney and Elaine Newton, "Trail

data or have their identities, illnesses, or other health information predicted through proxy data will only increase as AI systems become integrated into health and consumer products.³⁹

Moreover, conflict of interest issues may arise in new and unanticipated ways if health AI providers, developers, and vendors fail to disclose all the relevant risks to patients posed by systems in which they have a financial stake. Even continuity of care might become an issue if the AI systems behind decision-support tools are kept proprietary to AI companies, and health providers are unable to share or transfer patient information or predictive analytics in the same way that Electronic Health Records (EHRs) are guaranteed to be portable with the patient.

Particularly thorny ethical issues may arise as AI emerges in “smart” medical devices. Currently, there are efforts to incorporate AI systems across multiple devices, from pacemakers to blood pressure and glucose monitors to activity trackers to multimodal bed sensors for the elderly to “smart” insulin pumps.⁴⁰ These AI applications promise many benefits including better remote patient monitoring.

However, the software that powers these devices is often proprietary, rather than open source (open to external scrutiny and auditing).⁴¹ A recently granted exemption to the Digital Millennium Copyright Act (DMCA) provides the opportunity to examine code for external medical devices; internal medical devices, however, are equally if not more important to examine.⁴² Experts have warned of grave security issues in the deployment of networked technologies across Internet of Things (IoT) devices, many focusing on medical devices as specifically problematic. A team at the University of South Alabama demonstrated these issues in 2015, successfully hacking the iStan brand networked pacemaker and “killing” a mannequin that had the pacemaker installed.⁴³

Re-identification: Learning Who You are From Where You Have Been,” LIDAP-WP12. Carnegie Mellon University, Laboratory for International Data Privacy, Pittsburgh, PA, March 2003, <http://dataprivacylab.org/dataprivacy/projects/trails/index3.html>. Latanya Sweeney, “Matching Known Patients to Health Records in Washington State Data,” Harvard University. Data Privacy Lab, White Paper 1089-1, Cambridge, MA, June 2013, <http://dataprivacylab.org/projects/wa/index.html>.

³⁹ Nicolas Terry, “Big Data Proxies and Health Privacy Exceptionalism,” *Health Matrix* 24 (2014): 65-108, <http://ssrn.com/abstract=2320088>. For how search data can predict health issues, see: Ryen W. White and Eric Horvitz, “Predicting escalations of medical queries based on web page structure and content,” in *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, ACM, 2010, 769-770. http://research.microsoft.com/en-us/um/people/horvitz/SIGIR_2010_page_structure.pdf

⁴⁰ G. P Moustiris., S. C. Hiridis, K. M. Deliparaschos, and K. M. Konstantinidis, “Evolution of Autonomous and Semi-Autonomous Robotic Surgical Systems: A Review of the Literature.” *The International Journal of Medical Robotics and Computer Assisted Surgery* 7, no. 4 (2011): 375–92, <http://onlinelibrary.wiley.com/doi/10.1002/rcs.408/abstract>.

⁴¹ For a discussion of the implications of such proprietary systems, see Frank Pasquale, *The Black Box Society: The secret algorithms that control money and information*, Cambridge: Harvard University Press, 2015.

⁴² Karen Sandler, Lysandra Ohrstrom, Laura Moy and Robert McVay, “Killed by Code: Software Transparency in Implantable Medical Devices,” Software Freedom Law Center, New York, 2010, <http://www.softwarefreedom.org/resources/2010/transparent-medical-devices.html>.

⁴³ William Bradley Glisson, Todd Andel, Todd McDonald, Mike Jacobs, Matt Campbell, and Johnny Mayr, “Compromising a Medical Mannequin,” *arXiv preprint, arXiv:1509.00065*, 2015, <https://arxiv.org/pdf/1509.00065v1.pdf>

Compounding the issue, regulatory agencies such as the FDA struggle to maintain sufficient expertise and resources to audit such devices for biases or inaccuracies.⁴⁴ With the addition of AI systems, such challenges will multiply in both complexity and cost. In particular, because AI tends to generate many of its decisions dynamically and in emergent patterns, even access to source code may not be sufficient to assess its impact on any particular medical device. Thus, the right of patients to be informed about the benefits and risks of these devices becomes a significant challenge in a world of AI-enabled medical technologies. Determining the dimensions of transparency and accountability required for effective regulation will be a crucial area of debate. How could patients and practitioners audit or contest AI-based clinical decisions? What responsibility do those who provide AI health products and services have to be accountable for health outcomes?

Questions to consider

- What are the goals of AI-based systems in healthcare? Given the range of customers for such systems (healthcare providers, governments, and insurers), what different outcomes are they optimizing for, and are these outcomes in the best interest of patients?
- What does the relative paucity of data for certain demographics (women, people of color) mean for the accuracy of AI diagnostic and decision-support technologies trained on available data? How will this impact existing health disparities?
- Data sharing and standardization is key for the expansion of AI in health. In June, 2016, the US Secretary for Health and Human Services announced a federal initiative aimed at facilitating health data sharing among industry leaders. What challenges are raised by these new sharing agreements? Are current data privacy and ethics frameworks built to handle sharing for purposes of AI?
- Who gets to decide which forms of care are prioritized in a world of AI diagnostics? What rights might doctors, patients, and other caregivers have to contest and alter decisions made by AI systems? How might such decisions be audited and understood by patients and doctors?
- How can patient confidentiality, autonomy, and informed consent be maintained in the face of the increasingly granular tracking required for precision medicine and its attendant AI systems?
- How does AI impact the relationship between patients, doctors, and other caregivers? What does it mean for a machine to “care”, and what impact will mechanized care work have on patients? What is the role of affective labor, empathy, and human connection in a world in which “care” is increasingly automated?

⁴⁴ Gregory Curfman and Rita Redberg, “Medical Devices — Balancing Regulation and Innovation,” *New England Journal of Medicine* 365, no. 11 (September 15, 2011): 975–77, <http://www.nejm.org/doi/full/10.1056/NEJMp1109094#t=article>.