

February 15, 2019

D. Bruce Hoffman
Director
Bureau of Competition
U.S. Federal Trade Commission
600 Pennsylvania Avenue, NW
Washington, DC 20580

Re: Federal Trade Commission: Competition and Consumer Protection Implications of Algorithms, Artificial Intelligence, and Predictive Analytics Hearing, Docket ID FTC-2018-0101

Dear Director Hoffman:

On behalf of the physician and medical student members of the American Medical Association (AMA), I appreciate the opportunity to submit comments in response to the Federal Trade Commission's (FTC) hearing on *Competition and Consumer Protection Implications of Algorithms, Artificial Intelligence, and Predictive Analytics*. The AMA applauds your leadership in pursuing a fact-finding agenda to assess the implications of advances in technology as applied to an array of sectors including health care. In addition to the comments below, the AMA has submitted comments to the FTC in response to the hearing concerning *Consumer Privacy and Data Security*. The AMA strongly supports the inclusion of practicing physicians and physician organizations in the policy discussions concerning these systems and the varied applications in health care. We welcome further engagement with the FTC to ensure consumer protection and competition are addressed consistently across industries and, where appropriate, health care specific policies are also adopted.

Standard Terminology

The AMA urges the FTC and other federal agencies to use the term "augmented intelligence" instead of "artificial intelligence" (without regard to whether the AI systems are intended to assist or to function autonomously). Artificial intelligence constitutes a host of computational methods that produce systems that perform tasks normally requiring human intelligence. These computational methods include, but are not limited to, machine image recognition, natural language processing, and machine learning. However, in health care the more appropriate term is "augmented intelligence" (referred to hereafter as AI), reflecting the enhanced capabilities of human clinical decision making when coupled with these computational methods and system.

The AMA strongly urges the FTC to partner with other federal agencies to collaborate early and often to establish to the greatest extent possible shared definitions concerning AI systems. Over a two-year period of extensive discussions with physicians, innovators, policymakers, regulators, and experts from various disciplines, we found that the lack of clearly defined terms at the outset of policy discussions has often led to confusion and a mismatch in expectations and a general lack of mutual understanding. Often, stakeholders are discussing varied AI systems with different risks, costs, intended uses, and necessary

conditions for successful deployment. The FTC hearing to advance the deployment of AI systems in health care—along with the President’s Executive Order, *Maintaining American Leadership in Artificial Intelligence*, and the work underway by the National Science Foundation, National Institutes of Health, Food and Drug Administration, and Centers for Medicare & Medicaid Services—provides forums to collaborate and represent opportunities for cross-agency synergies as applied to health care AI.

AMA Policy on Health Care Augmented Intelligence

The AMA’s House of Delegates adopted new policy on health care AI in June 2018. The policy reflects many of the themes and views of the physician panelist selected to present on the application of AI systems to clinical care during the FTC hearing. Specifically, the AMA welcomes the important transformations and improvements to health care AI systems promise to provide. The AMA agrees that it is essential that clinical applications of AI systems are independently clinically validated, deployment conditions clearly specified, and clinicians have the information needed to understand the clinical limitations and risks—even where a system is fully autonomous. And, the AMA strongly supports the position that where AI systems are fully autonomous in the delivery of clinical care, developers must take appropriate steps to address their assumption of liability such as obtaining medical liability insurance. The AMA does not support shifting risk and liability to patients and clinicians where they are not able to mitigate the risks (or even know what they are).

In brief, in accordance with our policy, the AMA is committed to promoting the integration of the views and expertise of practicing physicians into the development, design, validation, and implementation of health care AI. As noted during the hearing and as laid out in AMA policy, it is important to promote the development of thoughtfully designed, high-quality, clinically validated health care AI that:

- is designed and evaluated in keeping with best practices in user-centered design, particularly for physicians and other members of the health care team;
- is transparent;
- conforms to leading standards for reproducibility;
- identifies and takes steps to address bias and avoids introducing or exacerbating health care disparities including when testing or deploying new AI tools on vulnerable populations; and
- safeguards patients’ and other individuals’ privacy interests and preserves the security and integrity of personal information.

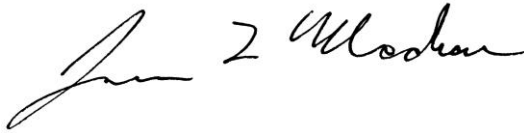
The AMA is also committed to encouraging education for patients, physicians, medical students, other health care professionals, and health administrators to promote greater understanding of the promise and limitations of health care AI. In addition, the AMA is focused on addressing other legal implications of health care AI, such as, issues of liability, intellectual property, and appropriate professional and governmental oversight for safe, effective, and equitable use of and access to health care AI.

Enclosed is a background brief along with AMA policy that provides a fuller discussion of some of these issues. The AMA appreciates the opportunity to comment and we welcome engaging further on this topic.

D. Bruce Hoffman
February 15, 2019
Page 3

If you have questions, please contact Shannon Curtis, Assistant Director of Federal Affairs at 202-789-8510 or Shannon.Curtis@ama-assn.org.

Sincerely,

A handwritten signature in black ink, appearing to read "Jim L Madara". The signature is written in a cursive style with a large initial "J" and a stylized "L".

James L. Madara, MD

Enclosure



Augmented intelligence in health care*

Interest in augmented intelligence (AI) and its potential to dramatically impact medicine is growing rapidly among Congress, federal agencies, and other health care stakeholders. As a leader in American medicine, our American Medical Association (AMA) is uniquely positioned to ensure that the evolution of AI in medicine benefits patients, physicians, and the health care community. This report contains baseline policy to guide AMA's engagement with a broad cross-section of stakeholders and policymakers to ensure that the perspective of physicians in various practice settings informs and influences the dialogue as this technology develops.

Ensuring the appropriate implementation of AI in health care will require that stakeholders forthrightly address challenges in the design, evaluation, implementation, and oversight of AI systems. Through its strategic partnerships and collaborations, the AMA has the capacity to help set priorities for health care AI; integrate the perspective of practicing physicians into the design, validation, and implementation of high-quality, clinically valuable health care AI; and promote greater understanding of the promise and limitations of AI across the health care community. A strong tradition of advocacy well positions our AMA to explore the legal implications of the emerging technologies of AI in health care and advocate effectively for appropriate professional and governmental oversight for safe, effective, equitable use of and access to health care AI.

AMA policy

As a leader in American medicine, our American Medical Association (AMA) has a unique opportunity to ensure that the evolution of augmented intelligence (AI) in medicine benefits patients, physicians, and the health care community. To that end our AMA will seek to:

- Leverage its ongoing engagement in digital health and other priority areas for improving patient outcomes and physicians' professional satisfaction to help set priorities for health care AI.
- Identify opportunities to integrate the perspective of practicing physicians into the development, design, validation, and implementation of health care AI.
- Promote development of thoughtfully designed, high-quality, clinically validated health care AI that:
 - is designed and evaluated in keeping with best practices in user-centered design, particularly for physicians and other members of the health care team;
 - is transparent;
 - conforms to leading standards for reproducibility;
 - identifies and takes steps to address bias and avoids introducing or exacerbating health care disparities including when testing or deploying new AI tools on vulnerable populations; and
 - safeguards patients' and other individuals' privacy interests and preserves the security and integrity of personal information.
- Encourage education for patients, physicians, medical students, other health care professionals, and health administrators to promote greater understanding of the promise and limitations of health care AI.
- Explore the legal implications of health care AI, such as issues of liability or intellectual property, and advocate for appropriate professional and governmental oversight for safe, effective, and equitable use of and access to health care AI.

* Content derived from Augmented Intelligence (AI) in Health Care (Annual Meeting 2018)

What is health care AI?

Computational methods and techniques for data analysis have been evolving for decades. [1,2] A number of these methods have come to be known collectively as “artificial intelligence.” Artificial intelligence constitutes a host of computational methods that produce systems that perform tasks normally requiring human intelligence. These computational methods include, but are not limited to, machine image recognition, natural language processing, and machine learning. However, in health care a more appropriate term is “augmented intelligence” (AI), reflecting the enhanced capabilities of human clinical decision making when coupled with these computational methods and systems.

In December 2017, Senators Maria Cantwell (D-WA), Todd C. Young (R-IN), and Edward Markey (D-MA) and U.S. Representatives John Delaney (D-MD) and Pete Olson (R-TX) introduced S. 2217/H.R. 4625, “Fundamentally Understanding the Usability and Realistic Evolution (FUTURE) of Artificial Intelligence Act of 2017.” The legislation defines “general AI” as computational methods that produce systems that exhibit intelligent behavior at least as advanced as a human across the range of cognitive, emotional, and social behaviors. In contrast, the bill defines the term “narrow AI” as computational methods that address specific application areas, such as playing strategic games, language translation, self-driving vehicles, and image recognition. Thus, these AI methods and tools for the foreseeable future are better characterized as narrow AI that augments human intelligence (augmented intelligence).

At a February 2018 U.S. House of Representatives Government Oversight Committee Subcommittee on Information Technology hearing, three national experts testified that general AI is decades away. Consistent with the foregoing, in response to a 2016 Request for Information on Artificial Intelligence issued by the White House Office of Science and Technology Policy, a technology company stated that it is “guided by the term ‘augmented intelligence’ rather than ‘artificial intelligence’” and noted further that “[i]t is the critical difference between systems that enhance and scale human expertise rather than those that attempt to replicate all of human intelligence.” [3]

Software algorithms developed using these evolving methods and techniques, coupled with proliferating sources of data (datasets) pertinent to health and medicine, offer the promise of new and more powerful



ways to augment human intelligence and expertise in health care.

The American College of Radiology (ACR), which has been at the leading edge of health care AI, addressed its promise in comments to the White House Office of Science and Technology Policy in 2016:

AI could offer various benefits to medical imaging in the future, including augmenting the capabilities of radiologists to enhance their efficiency and accuracy, as well as reducing costs by improving the appropriateness and cost-effectiveness of medical imaging utilization. The use of AI and machine learning in health care in general could be best applied to the areas of precision medicine, predictive analytics, and outcomes assessments. AI can streamline health care workflow and improve triage of patients (especially in acute care settings), reduce clinician fatigue, and increase the efficiency and efficacy of training. Moreover, shortages of medical experts to meet the needs of vulnerable and underserved populations in domestic and international settings could potentially be relieved, in part, by AI [4].

Prime AI applications include clinical decision support, patient monitoring and coaching, automated devices to assist in surgery or patient care, and management of health care systems [5]. AI in health care holds out the

prospect of improving physicians' ability to establish prognosis [6], as well as the accuracy and speed of diagnosis [6,7,8], enabling population-level insights to directly inform the care of individual patients [9], and predicting patient response to interventions [10]. The number of empirical studies of AI applications in medicine is growing rapidly [2].

What's next in health care AI?

Commercial entities are driving rapid evolution in AI across the board. In health care, the next three to five years will be marked by efforts to scale AI options involving patient-centered wearables that support clinical care, improved tools for diagnosis and physician training, and health system initiatives to improve patient care and clinical decision support [11]. The following are early examples of such efforts.

Wearable AI

Wearable monitoring devices that can transmit patient data are evolving rapidly. For example, one company has developed the Cardiogram application which is designed to work with the built-in infrared heart rate sensor of the Apple Watch to detect hypertension and sleep apnea. In a study carried out with the University of California–San Francisco that involved over 6,000 patients, the application and its machine learning system, DeepHeart, was able to detect hypertension and sleep apnea with 82 percent and 90 percent accuracy, respectively [12]. Rapid innovation is expected on this front propelled by coverage of payers, including Medicare, of remote patient monitoring and management.

New tools for diagnosis and physician training

The utilization of machine learning algorithms to enhance clinical decision making is increasing, but emerging systems take such support a step further. For example, the Human Diagnosis Project (Human Dx), organized as a tandem 501(c)(3) nonprofit and public benefit corporation, and created with and led by the medical community, allows attending physicians to ask for assistance on difficult medical cases from an online community of physicians all over the world. Responses from the medical community are combined with help from machine learning to create a synthesized collective assessment for each case. This collective insight is designed to augment clinical decision making with machine intelligence, providing useful information to physicians and patients who may not otherwise have access to specialist expertise. Human Dx also provides a platform for medical education through its Global Morning Report teaching cases. Today, residents from

over 40 percent of U.S. internal medicine residency programs have access to these cases. Human Dx vets the quality of responses by comparing how physicians solve reference training cases in order to calculate a quantitative measure of reasoning called Clinical Quotient, which is now being vetted in conjunction with the Johns Hopkins School of Medicine.

Health systems and data analytics

Applying AI to health system data to improve care is another area of rapid evolution. The University of Pittsburgh Medical Center (UPMC) has launched a system-wide effort to reduce hospital readmissions and enhance clinical decision making while a patient is receiving care. UPMC has applied machine learning to claims data to predict a patient's risk of readmission before the patient arrives. A second algorithm uses laboratory and clinical metrics extracted from clinical records to update the risk prediction every 15 minutes over the course of the patient's admission. Before discharge, if the risk prediction's two models are in conflict, UPMC uses unsupervised machine learning to come up with a set of rules that dictate which model takes precedence to inform clinician discharge decisions [13].

These three relatively nascent efforts are designed to scale, but will require significant additional research and real world testing. However, they illustrate the types of initiatives beyond condition-specific efforts to enhance clinical decision support that could produce significant improvements in health care. Notably, these efforts have active engagement and support of clinicians and seek to address medical challenges and problems identified by clinicians.

Federal engagement with AI

AI has surfaced as a public policy issue at the federal level in a relatively short period of time. In 2016, the White House Office of Science and Technology hosted several public meetings on a range of public policy issues addressing AI along with a public request for information regarding potential policy directions. In Congress, the U.S. Senate Commerce Committee held a hearing titled "The Dawn of Artificial Intelligence" at which the Department Chair for Genomic Medicine at MD Anderson Cancer Center highlighted the clinical applications of AI and discussed policy implications.

Shortly thereafter, the 21st Century Cures Act was passed by Congress and became law in December 2016. The Act included provisions modifying the U.S. Food and Drug Administration's (FDA) oversight of software as a

medical device, which has implications for a number of current AI computational methods. The FDA is now actively evaluating whether a new oversight framework is needed for software as a medical device, a precursor to future oversight models.

The bipartisan “FUTURE of Artificial Intelligence Act,” introduced in December 2017, provides for the establishment of a Federal Advisory Committee on the Development and Implementation of Artificial Intelligence. The legislation, if passed, would be the first effort at the federal level to provide a forum for consideration of AI public policy. In 2018, additional legislation has been introduced, and additional congressional hearings held on AI generally, with health care applications receiving particular attention.

Achieving the promise of AI in health care


Fulfilling the promise that “combining machine learning software with the best human clinician ‘hardware’ will permit delivery of care that outperforms what either can do alone” [14] will require that stakeholders forthrightly address challenges in the design, evaluation, implementation, and oversight of AI systems in health care. In the first instance, stakeholders across the board, not the least among them patients and physicians, must hold realistic expectations for the roles AI tools can and cannot play. Machine learning is only one of the AI computational methods and raises particularly thorny challenges. However, many of the public policy issues (including transparency and intellectual property) and clinical issues that will need to be addressed apply to other AI computational methods that are more common currently, such as natural language processing.

Designing and evaluating health care AI

There is a popular tendency to see AI as, at best, a form of neutral, “objective” decision making, a pristine mathematical process that takes only “the facts” into account, independent of human judgment [15,16,17]. The statistical process of AI specifically seeks to derive a rule or procedure from a body of data that explains that data or is able to predict future data [18]. An AI derived algorithm “is only as good as the data it works with” [19,20]. The data sets on which AI algorithms are trained are created by human agents and are imperfect.

The research, patient care, and insurance records available as training data sets for health care AI can be highly variable, reflecting the different purposes for and processes by which they were created [1,21]. Clinical trials systematically include or exclude participants with certain characteristics; patient charts and insurance records capture information only from those individuals who have access to the health care system and rarely contain information about exposure to environmental toxins. Different data sets focus on different kinds of information to the exclusion of other possible data points, and records capture and preserve information with varying degrees of accuracy.

One of the most significant implications for end users of AI systems is that these systems sets can, invisibly and unintentionally, “reproduce and normalize” the biases of their training data sets [16,17]. In health care, the result can be models that “reflect the conditions only of the fortunate” and yield “an aggregate understanding of health and illness that fundamentally excludes the marginalized” [21] in a way that risks exacerbating existing health disparities. Minority populations can be disadvantaged in the context of AI systems in a second way as well in that “by definition, there is proportionately less data available about minority



predictions,” while the accuracy of decision making, a proxy for fairness, will be higher for majority groups [17]. Addressing fairness is essential, even if doing so may be costly for developers when it requires them to seek more complex decision rules [17].

Design issues also encompass how a model is evaluated, as well as relationships between the dataset used to train an algorithm and the dataset used to evaluate the algorithm. In the first instance, evaluation criteria must be clinically relevant and evaluation should be representative of how the algorithm will be applied in practice [22]. For example, evaluating a model to predict risk of hospital-acquired infection over the entire course of a patient’s admission more accurately predicts how the model would be used and would perform in practice [22]. For predictive models, developers must evaluate “how far in advance the algorithm identifies positive cases.” [22] From a clinician’s perspective, the critical concern is “predicting events early enough for a relevant intervention to influence care decisions and outcomes.” [14] Ensuring that all examples in the training dataset are earlier in time than all examples in the evaluation set helps avoid misleading results by limiting the possibility that training data could otherwise reflect structural changes in hospital population, clinical protocols, electronic health record (EHR) systems, or other factors that occurred over time [22].

Developers also have a responsibility to ensure that their work is transparent and can be reproduced by others [23,24]. Proposed guidelines for essential components of publications reporting development of predictive machine-learning algorithms include not only rationale and objectives, but, importantly, the setting, prediction problem, relevant data, and a description of the building of the predictive model [23]. Authors should also provide information about the final model and its performance, and discuss the clinical implications of the work, its limitations, and unexpected results. Scholars have further recommended creating open repositories for long-term storage, archiving, and access to datasets and code to enable replication of published findings [24].

Furthermore, the AMA’s work in the area of EHRs reveals that to be useful and accepted in practice, AI systems need to be developed and evaluated in keeping with best practices in user-centered design [25]. The focus must be on users’ needs and usability should be tested by participants who are demographically representative of end users [26].

Health care AI and patient privacy


Commitment to protecting the confidentiality of patient information is central to medicine’s professional ethos. In this respect, AI poses a significant challenge where traditional strategies of notification and consent are no longer adequate [18]. Nor are anonymization, deletion of data, or distinguishing metadata sufficiently robust protections in the context of massive complex data sets [18,20] when machine-learning algorithms can identify a record “easily and robustly” from as few as three data points [20].

The ease of re-identification means that, in important respects, traditional expectations for health care privacy are simply no longer attainable. This significantly raises the bar on the task of ensuring the security and integrity of data. Among proposed technical solutions to the dilemma of privacy in large data sets are “blockchain-style” technology to secure data and track access or data auditing systems that allow secure verification of the contents of large data structures, such as those being explored by DeepMind Health in the UK [1]. Researchers at the University of Pennsylvania have explored the creation of publicly sharable simulated datasets that limit possible re-identification as another approach to protecting data privacy [27]. The recent revelation that the data mining firm Cambridge Analytica siphoned private data from 50 million Facebook users to target them for political campaigns raises confidentiality and privacy questions across the spectrum of digital platforms that collect and curate data. While this report establishes policy that underscores the necessity to safeguard individuals’ privacy interests and preserve the security and integrity of personal information, the Board recognizes the importance of this issue and will continue to assess our policy as our AMA engages in the public debate and discourse on protecting patient information.

Implementing health care AI

The AMA’s ongoing engagement with digital health offers insights for understanding, from physicians’ perspectives, what is at stake in integrating AI systems into the delivery of health care. The organization’s recent survey of 1,300 physicians about barriers to adoption of digital health technologies suggests that physicians are most receptive to digital health tools they believe can be integrated smoothly into their current practice, will improve care, and will enhance patient-physician relationships [28]. Coverage for liability, assurance that data privacy is protected, linkage to their EHR, and billing/reimbursement are key considerations.

Earlier AMA research into physician professional satisfaction found that frustrations with EHRs, especially



usability issues, were a major source of dissatisfaction in physicians' professional lives [29]. The findings led the AMA to identify priorities for ensuring usability in EHR systems, including, among other considerations, ensuring that EHRs are designed to meet the cognitive and workflow needs of physicians, that they support team-based care, promote coordination of care, focus on reducing cognitive workload instead of focusing simply on data collection, and incorporate end user feedback into designing and improving EHR systems [25].

AMA policies addressing the use of telemedicine similarly stress the importance of minimizing disruptive effects on patient-physician interactions, ensuring that technologies promote quality of care and safety, and, importantly, establishing mechanisms to monitor the impact of an innovation both to identify and address adverse consequences and to identify and encourage dissemination of outcomes [30,31].

To reap the benefits for patient care, physicians must have the skills to work comfortably with health care AI. Just as working effectively with EHRs is now part of training for medical students and residents [32], educating physicians to work effectively with AI systems, or more narrowly, the AI algorithms that can inform clinical care decisions, will be critical to the future of AI in health care.

Physicians need to understand AI methods and systems sufficiently to be able to trust an algorithm's predictions—or know how to assess the trustworthiness and value of an algorithm—as a foundation for clinical recommendations. The challenge may be more easily met with advances in “explainable AI,” that is, algorithms

that can “explain” to users why a particular prediction is made [33,34]. Technology to predict the risk of 30-day readmission for cardiac patients being tested by Boston-based Partners Connected Health provides clinicians with a readmission prediction score and identifies the top factors contributing to that score, providing information that is actionable for clinicians [35].

A leadership role for the AMA

A component of the AMA's strategic work in 2018 and beyond has been to provide the physician perspective across health care technology sectors by promoting improved usability of and productive access to data used in medical decision making as well as respect for the patient-physician relationship. As our AMA implements this component of its strategic plan, the Board of Trustees has observed a rapidly growing interest in augmented intelligence (AI) technology in health care. In 2018, the AMA Council on Long Range Planning and Development (CLRPD) provided the Board with a primer on the history, definitions and components, and the status of AI in health care that offered a high-level look at this rapidly evolving area and its potential to dramatically impact medicine. The AMA Council on Legislation (COL) and CLRPD have observed increased interest in AI by Congress, federal agencies, and other health care stakeholders. To form a clearer understanding of the expected impact of AI technologies for patients and physicians, as well as key stakeholders who are influencing legislation and regulation in this area, the COL has met with physician experts immersed in the development and clinical integration of various health care AI technologies.

The AMA has adopted a base-level of policy on health care AI to guide AMA's engagement with a broad cross-section of stakeholders and policymakers in order to ensure that the perspective of physicians in various practice settings informs and influences the dialogue as this technology develops.

To realize its potential to support improved patient care and health outcomes and enhance physician professional satisfaction, the health care AI enterprise should be informed and guided by the expertise, experience, and leadership of physicians and organized medicine in developing and implementing these tools. Physicians are well positioned to advocate for health care AI solutions that support healthier lifestyles and reduce disease burden, improve access to care, enhance diagnostic accuracy, inform individually tailored treatment plans, and improve patient self-management,

adherence, and health outcomes. Physicians are likewise well placed to apply their experience to drive improved design and implementation of health care AI that will strengthen clinicians' relationships with patients; enhance communication among the health care team and between team members, patients, and family members; simplify the coordination of care; minimize administrative burdens; and help the health care team to better deliver care to those patients and populations in greatest need.

In addition to the work of COL and CLRPD, at the 2017 Interim Meeting all seven AMA councils met jointly with AI experts to discuss issues in health care AI. Likewise, the AMA's ongoing engagement with key stakeholders from across the spectrum of clinical care, health care administration, implementation science, and AI product development enables the organization to play a distinctive role in contributing to the overarching vision for health care AI in the U.S.

Through its strategic partnerships and collaborations, the AMA has the capacity to offer the insight that is critical to the development of clinically sound AI systems that will enhance the quality of care and sustain the integrity of patient-physician relationships. The AMA's strong tradition of advocacy positions the organization to promote meaningful oversight of AI as it is integrated into clinical practice.

Conclusion

Patients, physicians, and the health care system in the U.S. face enormous challenges in the combined impact of a rapidly aging population, a relative decline in the working population that reduces revenue essential for safety net programs [36], and persistent high costs of care that will strain the nation's ability to support affordable, accessible, high quality care. With the engagement of physicians to identify needs and set priorities for design, development, and implementation, health care AI can offer a transformative set of tools to help patients, physicians, and the nation face these looming challenges. Given the number of stakeholders and policymakers involved in the evolution of AI in health care, it is important that our AMA not only adopt a base level of policy to guide our engagement, but equally continue to refine our policy as an organization to ensure that the perspective of physicians in various practice settings informs and influences the dialogue as this technology develops.

References

1. JASON. *Artificial Intelligence for Health and Health Care*. The MITRE Corporation, 2017. Available at https://www.healthit.gov/sites/default/files/jsr-17-task-002_aiforhealthandhealthcare12122017.pdf. Accessed February 19, 2018.
2. Chang A. *AIMed: Artificial Intelligence in Medicine: Analytics, Big Data, Cloud and Cognitive Computing, Databases, and Deep Learning in Healthcare and Medicine*, 2017.
3. IBM Research. *Preparing for the Future of Artificial Intelligence*, 2016. Available at <http://www.research.ibm.com/cognitive-computing/ostp/rfi-response.shtml>. Accessed January 30, 2018.
4. American College of Radiology. Request for information on artificial intelligence; comments from the American College of Radiology; July 22, 2016. Available at https://cqcengage.com/acradiology/file/ywodvc80AHO/acr_comments_ostp_rfiai_7222016.pdf. Accessed February 29, 2018.
5. Stanford University. *Artificial Intelligence and Life in 2030*. September 2016. Available at <https://ai100.stanford.edu/2016-report>. Accessed February 6, 2018.
6. Obermeyer Z, Emanuel EJ. Predicting the future—big data, machine learning, and clinical medicine. *New Engl. J. Med* 2016;375:1216–1219.
7. Darcy AM, Louie AK, Roberts LW. Machine learning and the profession of medicine. *JAMA*. 2016;315(6):551–552.
8. Golden JA. Deep learning algorithms for detection of lymph node metastases from breast cancer: helping artificial intelligence be seen. *JAMA*. 2017;318(22):2184–2186.
9. Mandl KD, Bourgeois FT. The evolution of patient diagnosis from art to digital data-driven science. *JAMA*. 2017;318(19):1859–1860.
10. Redlich R, Opel N, Grotegerd D, et al. Prediction of individual response to electroconvulsive therapy via machine learning on structural magnetic resonance imaging data. *JAMA Psychiatry*. 2016;73(6):557–564.
11. Council on Long-Range Planning and Development. *A Primer on Artificial and Augmented Intelligence*. Memorandum to the Board of Trustees, February 2018.
12. Tison HG, Singh AC, Chashi DA, et al. Cardiovascular risk stratification using off-the-shelf wearables and a multi-task deep learning algorithm. Presentation to the AHA Science Session, Anaheim, California. November 13, 2017

13. McCann E. UPMC's big data approach slashes readmissions. *Healthcare IT News* 2015;October 1. Available at <http://www.healthcareitnews.com/news/upmcs-big-data-approach-slashes-readmissions>. Accessed February 6, 2018.
14. Chen JH, Asch SM. Machine learning and prediction in medicine—beyond the peak of inflated expectations. *N Engl J Med* 2017;376:2507–2509.
15. Bogost I. The cathedral of computation. *Atlantic* 2015; January 15.
16. Burrell J. How the machine “thinks”: understanding opacity in machine learning algorithms. *Big Data & Society* 2016;January-June:1–12.
17. Hardt M. How big data is unfair. *Medium* 2014; September: 26.
18. Executive Office of the President. *Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights*. 2016.
19. Barocas S, Selbst AD. Big data's disparate impact. *California Law Review* 2016;104:671–732.
20. Osoba O, Welser IV, W. *An Intelligence in Our Image: The Risks of Bias and Errors in Artificial Intelligence*. 2017. Available at https://www.rand.org/pubs/research_reports/RR1744.html. Accessed February 19, 2018.
21. AI Now Institute. *The Social and Economic Implications of Artificial Intelligence Technologies in the Near-Term*. 2016. Available at <https://ainowinstitute.org/reports.html>. Accessed January 26, 2018.
22. Wiens J, Gutttag J, Horvitz E. Patient risk stratification with time-varying parameters: a multitask learning approach. *J Machine Learning Research* 2016;17:1–23.
23. Luo W, Phung D, Tran T, et al. Guidelines for developing and reporting machine learning predictive models in biomedical research: a multidisciplinary review. *J Med Internet Research* 2016;18(12):e323.
24. Stodden V. Enhancing reproducibility for computational methods. Presentation, National Academies of Science, Engineering, and Medicine, July 20, 2017.
25. American Medical Association. *Improving Care: Priorities to Improve Electronic Health Records*. Available at <https://www.ama-assn.org/sites/default/files/media-browser/member/about-ama/ehr-priorities.pdf>. Accessed January 30, 2018.
26. MedStar Health. *Why We Created the EHR User-Centered Design Evaluation Framework*. Available at [https://mi2.medstarhealth.org/human-factors/research/ehr-user-centered-design-evaluation-framework/methodology/#q={}](https://mi2.medstarhealth.org/human-factors/research/ehr-user-centered-design-evaluation-framework/methodology/#q={). Accessed January 29, 2018.
27. Beaulieu-Jones BK, Wu ZS, Williams C, Green CS. Privacy-preserving generative deep neural networks support clinical data sharing. bioRxiv preprint first posted online Jul. 5, 2017; doi: <http://dx.doi.org/10.1101/159756>. Accessed January 29, 2018.
28. American Medical Association. *Digital Health Study*. 2016. Available at <https://www.ama-assn.org/sites/default/files/media-browser/specialty%20group/washington/ama-digital-health-report923.pdf>. Accessed February 19, 2018.
29. Freidberg MW, Chen PG, Van Busum KR, et al. Factors affecting physician professional satisfaction and their implications for patient care, health systems, and health policy. *RAND Health Q* 2014;3(4):1. Available at <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5051918/?report=printable>. Accessed February 12, 2018.
30. American Medical Association. *Code of Medical Ethics Opinion E-1.2.12, Ethical Practice in Telemedicine*.
31. American Medical Association. *Code of Medical Ethics Opinion E-11.2.1, Professionalism in Health Care Systems*.
32. Indiana University School of Medicine. *Regenstrief EHR Clinical Learning Platform*. Accessed February 19, 2018.
33. Kuang C. Can AI be taught to explain itself? *NYT Magazine* 2017; November 21.
34. Holzinger A, Biemann C, Pattichis CS, Kell DB. What do we need to build explainable AI systems for the medical domain? *arXiv:1712.09923v1* 2017; December 28.
35. Pecci AW. “Explainable AI” could reduce readmission (and win clinicians’ trust). *Health Leaders Media* 2018; January 16.
36. Kvedar JC. *The New Mobile Age: How Technology will Extend the Healthspan and Optimize the Lifespan*. 2017.