THE LEARNING HEALTH SYSTEM SERIES

Artificial Intelligence in Health Care

The Hope, the Hype, the Promise, the Peril

Michael Matheny,
Sonoo Thadaney Israni, Mahnoor Ahmed,
and **Danielle Whicher,** *Editors*

NATIONAL ACADEMY OF MEDICINE

Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril

Michael Matheny, Sonoo Thadaney Israni, Mahnoor Ahmed,
and Danielle Whicher, *Editors*

WASHINGTON, DC
NAM.EDU

NATIONAL ACADEMY OF MEDICINE • 500 Fifth Street, NW • WASHINGTON, DC 20001

Library of Congress Cataloging-in-Publication Data to Come

Printed in the United States of America.

Suggested citation: Matheny, M., S. Thadaney Israni, M. Ahmed, and D. Whicher, Editors. 2019. *Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril.* NAM Special Publication. Washington, DC: National Academy of Medicine.

"Knowing is not enough; we must apply.
Willing is not enough; we must do."
--GOETHE

LEADERSHIP
INNOVATION
IMPACT

for a healthier future

NATIONAL
ACADEMY
of MEDICINE

ABOUT THE NATIONAL ACADEMY OF MEDICINE

The **National Academy of Medicine** is one of three Academies constituting the National Academies of Sciences, Engineering, and Medicine (the National Academies). The National Academies provide independent, objective analysis and advice to the nation and conduct other activities to solve complex problems and inform public policy decisions. The National Academies also encourage education and research, recognize outstanding contributions to knowledge, and increase public understanding in matters of science, engineering, and medicine.

The **National Academy of Sciences** was established in 1863 by an Act of Congress, signed by President Lincoln, as a private, nongovernmental institution to advise the nation on issues related to science and technology. Members are elected by their peers for outstanding contributions to research. Dr. Marcia McNutt is president.

The **National Academy of Engineering** was established in 1964 under the charter of the National Academy of Sciences to bring the practices of engineering to advising the nation. Members are elected by their peers for extraordinary contributions to engineering. Dr. John L. Anderson is president.

The **National Academy of Medicine** (formerly the Institute of Medicine) was established in 1970 under the charter of the National Academy of Sciences to advise the nation on issues of health, health care, and biomedical science and technology. Members are elected by their peers for distinguished contributions to medicine and health. Dr. Victor J. Dzau is president.

Learn more about the National Academy of Medicine at NAM.edu.

AUTHORS

MICHAEL MATHENY *(Co-Chair),* Vanderbilt University Medical Center and the Department of Veterans Affairs
SONOO THADANEY ISRANI *(Co-Chair)***,** Stanford University
ANDREW AUERBACH, University of California, San Francisco
ANDREW BEAM, Harvard University
PAUL BLEICHER, OptumLabs
WENDY CHAPMAN, University of Melbourne
JONATHAN CHEN, Stanford University
GUILHERME DEL FIOL, University of Utah
HOSSEIN ESTIRI, Harvard Medical School
JAMES FACKLER, Johns Hopkins School of Medicine
STEPHAN FIHN, University of Washington
ANNA GOLDENBERG, University of Toronto
SETH HAIN, Epic
JAIMEE HEFFNER, Fred Hutchinson Cancer Research Center
EDMUND JACKSON, Hospital Corporation of America
JEFFREY KLANN, Harvard Medical School and Massachusetts General Hospital
RITA KUKAFKA, Columbia University
HONGFANG LIU, Mayo Clinic
DOUGLAS MCNAIR, Bill & Melinda Gates Foundation
ENEIDA MENDONÇA, Regenstrief Institute
JONI PIERCE, University of Utah
W. NICHOLSON PRICE II, University of Michigan
JOACHIM ROSKI, Booz Allen Hamilton
SUCHI SARIA, Johns Hopkins University
NIGAM SHAH, Stanford University
RANAK TRIVEDI, Stanford University
JENNA WIENS, University of Michigan

NAM Staff

Development of this publication was facilitated by contributions of the following NAM staff, under the guidance of **J. Michael McGinnis,** Leonard D. Schaeffer Executive Officer and Executive Director of the Leadership Consortium for a Value & Science-Driven Health System:

DANIELLE WHICHER, Senior Program Officer (until September 2019)
MAHNOOR AHMED, Associate Program Officer
JESSICA BROWN, Executive Assistant to the Executive Officer (until September 2019)
FASIKA GEBRU, Senior Program Assistant
JENNA OGILVIE, Deputy Director of Communications

REVIEWERS

This special publication was reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise, in accordance with review procedures established by the National Academy of Medicine. We wish to thank the following individuals for their contributions:

SANJAY BASU, Harvard Medical School
PAT BAIRD, Philips
SARA BURROUS, Washington University, St. Louis
KEVIN JOHNSON, Vanderbilt University Medical Center
SALLY OKUN, PatientsLikeMe
J. MARC OVERHAGE, Cerner Corporation
JACK RESNECK, JR., American Medical Association
SARA ROSENBAUM, George Washington University
PAUL TANG, IBM Watson Health

TIMOTHY M. PERSONS, Chief Scientist and Managing Director, Science, Technology Assessment, and Analytics, United States Government Accountability Office
(NOTE: Dr. Persons only provided editorial comments and technical advice on the description of the artificial intelligence technology described in the publication. Dr. Persons did not comment on any policy related recommendations and did not review or comment on any of the legal content in the publication.)

The reviewers listed above provided many constructive comments and suggestions, but they were not asked to endorse the content of the publication, and did not see the final draft before it was published. Review of this publication was overseen by **Danielle Whicher,** Senior Program Officer, NAM; **Mahnoor Ahmed,** Associate Program Officer, and **J. Michael McGinnis,** Leonard D. Schaeffer Executive Officer, NAM. Responsibility for the final content of this publication rests with the editors and the NAM.

FOREWORD

In 2006, the National Academy of Medicine established the Roundtable on Evidence-Based Medicine for the purpose of providing a trusted venue for national leaders in health and health care to work cooperatively toward their common commitment to effective, innovative care that consistently generates value for patients and society. The goal of advancing a "Learning Health System" quickly emerged and was defined as "a system in which science, informatics, incentives, and culture are aligned for continuous improvement and innovation, with best practices seamlessly embedded in the delivery process and new knowledge captured as an integral by-product of the delivery experience"[1].

To advance this goal, and in recognition of the increasingly essential role that digital health innovations in data and analytics contribute to achieving this goal, the *Digital Health Learning Collaborative* was established. Over the life of the collaborative, the extraordinary preventive and clinical medical care implications of rapid innovations in artificial intelligence (AI) and machine learning emerged as essential considerations for the consortium. The publication you are now reading responds to the need for physicians, nurses and other clinicians, data scientists, health care administrators, public health officials, policy makers, regulators, purchasers of health care services, and patients to understand the basic concepts, current state of the art, and future implications of the revolution in AI and machine learning. We believe that this publication will be relevant to those seeking practical, relevant, understandable and useful information about key definitions, concepts, applicability, pitfalls, rate-limiting steps, and future trends in this increasingly important area.

Michael Matheny, MD, MS, MPH and Sonoo Thadaney Israni, MBA have assembled a stellar team of contributors, all of whom enjoy wide respect in their fields. Together, in this well-edited volume that has benefitted from the thorough review process ingrained in the National Academy of Medicine's culture, they present expert, understandable, comprehensive, and practical insights on topic areas that include the historical development of the field; lessons learned from other industries; how massive amounts of data from a variety of sources can be appropriately analyzed and integrated into clinical care; how innovations can be used to facilitate population health models and social determinants of health interventions; the opportunities to equitably and inclusively advance precision medicine; the applicability for health care organizations and businesses to reduce the cost of care delivery; opportunities to enhance

1 https://nam.edu/wp-content/uploads/2015/07/LearningHealthSystem_28jul15.pdf

interactions between health care professionals and patients, families, and caregivers; and the role of legal statutes that inform the uptake of AI in health care.

As the co-chairs of the *Digital Health Learning Collaborative*, we are excited by the progress being demonstrated in realizing a virtuous cycle in which the data inevitably produced by every patient encounter might be captured into a "collective memory" of health services to be used to inform and improve the subsequent care of the individual patient and the health system more generally. Enormous datasets are increasingly generated, not only in the formal health care setting, but also emanating from data streams from medical and consumer devices, wearables, patient-reported outcomes, as well as environmental, community and public health sources. They include structured (or mathematically operable) data as well as text, images and sounds. The landscape also includes data "mash-ups" from commercial, legal, and online social records.

AI has been the tool envisioned to offer the most promise in harvesting knowledge from that collective memory, and as this volume demonstrates, some of that promise is being realized. Among the most important of these promises in the near term is the opportunity to assuage the frustration of health care providers who have been clicking away on electronic records with modest benefit beyond increased data transportability and legibility. Our hope is that AI will be the "payback" for the investment in both the implementation of electronic health records and the cumbersomeness of their use by facilitating tasks that every clinician, patient, and family would want, but are impossible without electronic assistance—such as monitoring a patient for emergent sepsis 24 × 7 × 365 and providing timelier therapy for a condition in which diagnostic delay correlates with increased risk of death.

However, we also appreciate that AI alone cannot cure health care's ills and that new technologies bring novel and potentially under-appreciated challenges. For example, if a machine learning algorithm is trained with data containing a systematic bias, then that bias may be interpreted as normative, exacerbating rather than resolving disparities and inequities in care. Similarly, association of data does not prove causality, and it may not even be explanatory, suggesting that a simultaneous revolution in research methods is also necessary. Finally, the mere existence of substantial and sensitive data assets raises concerns about privacy and security. Aspiring to the promise of AI requires both continuing innovation and attention to the potential perils.

In our opinion, this publication presents a sober and balanced celebration of accomplishments, possibilities, and pitfalls. We commend Drs. Michael McGinnis and Danielle Whicher for their thoughtful sponsorship of the NAM Consortium and *Digital Health Learning Collaborative*, Dr. Matheny and Mrs. Thadaney Israni for their leadership in producing this volume, and to all the contributors who have produced an exceptional resource with practical relevance to a wide array of key stakeholders.

Jonathan B. Perlin, MD, PhD, MACP
Reed V. Tuckson, MD, FACP

Co-Chairs, Digital Learning Collaborative,
Consortium on Value and Science-Driven Health Care,
National Academy of Medicine

SUMMARY

The emergence of artificial intelligence (AI) as a tool for better health care offers unprecedented opportunities to improve patient and clinical team outcomes, reduce costs, and impact population health. Examples include but are not limited to automation; providing patient, "fRamily" (friends and family unpaid caregivers), and health care professionals' information synthesis; and recommendations and visualization of information for shared decision making.

While there have been a number of promising examples of AI applications in health care, we believe it is imperative to proceed with caution, else we may end up with user disillusionment and another AI winter, and/or further exacerbate existing health and technology driven disparities. The National Academy of Medicine's Special Publication: *Artificial Intelligence in Health Care: The Hope, The Hype, The Promise, The Peril* synthesizes current knowledge to offer a reference document for relevant health care stakeholders such as: AI model developers, clinical implementers, clinicians and patients, regulators, and policy makers, to name a few. It outlines the current and near-term AI solutions; highlights the challenges, limitations, and best practices for AI development, adoption, and maintenance; offers an overview of the legal and regulatory landscape for AI tools designed for health care application; prioritizes the need for equity, inclusion, and a human rights lens for this work; and outlines key considerations for moving forward. The major theses are summarized in the section below.

POPULATION-REPRESENTATIVE DATA ACCESSIBILITY, STANDARDIZATION, QUALITY IS VITAL

AI algorithms must be trained on population-representative data to achieve performance levels necessary for scalable "success." Trends such as the cost for storing and managing data, data collection via electronic health records, and exponential consumer health data generation have created a data-rich health care ecosystem. However, this growth in health care data struggles with the lack of efficient mechanisms for integrating and merging these data beyond their current silos. While there are multiple frameworks and standards in place to help aggregate and achieve sufficient data volume for AI use of data at rest (such as mature health care common data models) and data in motion (such as

HL7 FHIR), they need wider adoption to support AI tool development, deployment, and maintenance. There continue to be issues of interoperability and scale of data transfers due to cultural, social, and regulatory reasons. Solutions for them will require the engagement of all relevant stakeholders. Thus, the wider health care community should continue to advocate for policy, regulatory, and legislative mechanisms seeking to improve equitable, inclusive data collection and aggregation, and transparency around how patient health data may be best utilized to balance financial incentives and the public good.

ETHICAL HEALTH CARE, EQUITY, AND INCLUSIVITY SHOULD BE PRIORITIZED

Fulfilling this aspiration will require ensuring population-representative datasets and giving particular priority to what might be termed a new Quintuple Aim of Equity and Inclusion for health and health care (see *Figure S-1*). Else, the scaling possible with AI might further exacerbate the considerable existing inequities in health outcomes at a monumental scale. A single biased human or organizational impact is far less than that of global or national AI.

Prioritizing equity and inclusion should be a clearly stated goal when developing and deploying AI in health care. There are many high-profile examples of biased AI tools that have damaged the public's trust in these systems. It is judicious for developers and implementers to evaluate the suitability of the data used to develop AI tools and unpack the underlying biases in the data, to consider how the tool should be deployed, and to question whether various deployment environments could adversely impact equity and inclusivity. There are widely recognized inequities in health outcomes due to the variety of social determinants of health and perverse incentives in the existing health care system. Unfortunately, consumer-facing technologies have often worsened historical inequities in other fields and are at risk of doing so in health care as well.

FIGURE S-1 | Advancing to the Quintuple Aim
SOURCE: Developed by publication editors

THE DIALOGUE AROUND TRANSPARENCY AND TRUST SHOULD CHANGE TO BE DOMAIN- AND USE-CASE DIFFERENTIAL

Transparency is key to building this much needed trust among users and stakeholders, but there are distinct domains with differential needs of transparency. There should be full transparency on the composition, semantics, provenance, and quality of data used to develop AI tools. There also needs to be full transparency and adequate assessment of relevant performance components of AI. But, algorithmic transparency may not be required for all cases. AI developers, implementers, users, and regulators should collaboratively define guidelines for clarifying the level of transparency needed across a spectrum. These are key issues for regulatory agencies and clinical users, and requirements for performance are differential based on risk and intended use. Most importantly, we suggest clear separation of data, algorithmic, and performance reporting in AI dialogue, and the development of guidance in each of these spaces.

NEAR-TERM FOCUS SHOULD BE ON AUGMENTED INTELLIGENCE RATHER THAN FULL AUTOMATION

Some of the AI opportunities include supporting clinicians undertaking tasks currently limited to specialists; filtering out normal or low acuity clinical cases so specialists can work at the top of their licensure; helping humans address inattention, micro-aggressions, and fatigue; and business processes automation. Ensuring human-centered AI tools includes accepting that human override is important for developing user-trust because the public has an understandably low tolerance for machine error and that AI tools are being implemented in an environment of inadequate regulation and legislation. The near-term dialogue around AI in health care should focus on promoting, developing, and evaluating tools that support humans rather than replace it with full automation.

DEVELOP AND DEPLOY APPROPRIATE TRAINING AND EDUCATIONAL PROGRAMS TO SUPPORT HEALTH CARE AI

In order to benefit from, sustain and nurture AI tools in health care we need a thoughtful, sweeping, comprehensive expansion of relevant training and educational programs. Given the scale at which health care AI systems could change the landscape of the medical domain, the educational expansion must be multidisciplinary and engage AI developers, implementers, health care system leadership, frontline clinical teams, ethicists, humanists, and patients and patient caregivers, because each brings a core set of much needed requirements and expertise. Health care professional training programs should incorporate core curricula focused on teaching how to appropriately use data science and AI products and services. The needs of practicing health care professionals can be fulfilled via their required continuing education, empowering them to be more informed consumers. Additionally, retraining programs to

address a shift in desired skill sets due to increasing levels of AI deployment and the resulting skill and knowledge mismatches will be needed. Last, but not least, consumer health educational programs, at a range of educational levels, to help inform consumers on health care application selection and use are vital.

LEVERAGE EXISTING FRAMEWORKS AND BEST PRACTICES WITHIN THE LEARNING HEALTH CARE SYSTEM, HUMAN FACTORS, AND IMPLEMENTATION SCIENCE

The challenges in operationalizing AI technologies into the health care systems are countless in spite of the fact that this is one of the strongest growth areas in biomedical research and impact. The AI community must develop an integrated best practice framework for implementation and maintenance by incorporating existing best practices of ethical inclusivity, software development, implementation science, and human–computer interaction. This framework should be developed within the context of the learning health care system and be tied to targets and objectives. The cost and burden of implementing AI tools should be weighed against use case needs. AI tools should be pursued where other low- or no-technology solutions will not do as well. Successful AI implementation will need the committed engagement of health care stakeholders—leaders, AI developers, AI implementers, regulators, humanists, patients, and families. Health delivery systems should have a robust and mature underlying information technology (IT) governance strategy in place prior to them embarking on substantial AI deployment and integration. Lastly, national efforts should be deployed to provide capacity for AI deployment in lower resource environments where IT and informatics capacities are less robust. Linked to the prior considerations, this would help lower the entry barrier for adoption of these technologies and help promote greater health care equity. Health care AI could also go beyond the current limited biology-focused research to address patient and communal needs, expanding to meaningful and usable access of social determinants of health and psychosocial risk factors. AI has the potential (with appropriate consent) to link personal and public data for truly personalized health care.

BALANCING DEGREES OF REGULATION AND LEGISLATION OF AI TO PROMOTE INNOVATION, SAFETY, AND TRUST

AI applications have enormous ability to improve patient outcomes, but they could also pose significant risks in terms of inappropriate patient risk assessment, diagnostic error, treatment recommendations, privacy breaches, and other harms. Regulators should remain flexible, but the potential for lagging legal responses will remain a challenge for AI developers and deployers. In alignment with recent congressional and U.S. Food and Drug Administration developments and guidance, we suggest a graduated approach to the regulation of AI based on the level of patient risk, the level of AI autonomy, and considerations for how static or dynamic certain AI are likely to be. To the extent that machine learning–based models continuously learn from new data, regulators should adopt post-market surveil-

lance mechanisms to ensure continuing (and ideally improving) high-quality performance. Liability accrued when deploying AI algorithms will continue to be an emerging area as regulators, courts, and the risk-management industries deliberate. Tackling regulation and liability among AI adopters is vital when evaluating the risks and benefits. Regulators should engage stakeholders and experts to continuously evaluate deployed clinical AI for effectiveness and safety based on real-world data. Throughout that process, transparency can help deliver better-vetted solutions. To enable both AI development and oversight, government agencies should invest in infrastructure that promotes wider, ethical data collection and access to data resources for building AI solutions, within a priority of ethical use and data protection (See *Figure S-2*).

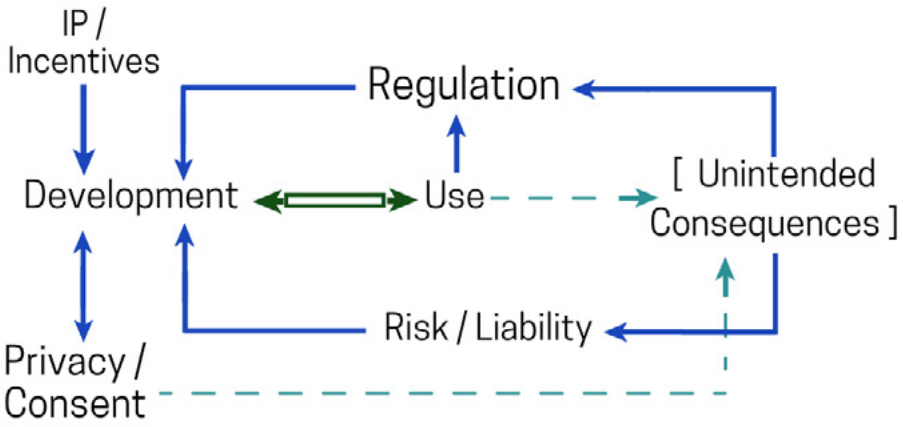

FIGURE S-2 | Appropriately regulating AI technologies will require balancing a number of important variables, including intellectual property (IP), concerns around privacy and consent, risks and liability associated with the use of the technologies, and developmental processes.
SOURCE: Developed by the authors

CONCLUSION

AI in health care is poised to make transformative and disruptive advances in health care. It is prudent to balance the need for thoughtful, inclusive health care AI that plans for and actively manages and reduces potential unintended consequences, while not yielding to marketing hype and profit motives. The wisest guidance for AI is to start with real problems in health care, explore the best solutions by engaging relevant stakeholders, frontline users, patients and their families—including AI and non-AI options—and implement and scale the ones that meet our Quintuple Aim: better health, improved care experience, clinician well-being, lower cost, and health equity throughout.