



Jointly published with



November 2020

United States Government Accountability Office
Science, Technology Assessment, and Analytics

Report to Congressional Requesters

TECHNOLOGY ASSESSMENT

Artificial Intelligence in Health Care

Benefits and Challenges of Technologies to
Augment Patient Care

With content from the National Academy of Medicine

Accessible Version



The cover image displays a stylized representation of data inputs from a variety of sources—including patient data from health records, research data, and textual data from scientific and medical literature—to a computer representing artificial intelligence (AI) algorithms. Those algorithms then offer recommendations or other assistance to providers that could augment their ability to care for patients.

This report is being jointly published by the Government Accountability Office (GAO) and the National Academy of Medicine (NAM). Part One presents GAO's Technology Assessment *Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care*. Part Two presents the NAM publication *Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic* discussing the use and challenges associated with AI technology in the delivery of health care services outside of settings where health care providers are employed, including the home. Although GAO and NAM staff consulted with and assisted each other throughout this work, reviews were conducted by GAO and NAM separately and independently, and authorship of the text of Part One and Part Two of the report lies solely with GAO and NAM, respectively.

With the exception of Part Two of this joint publication, this is a work of the U.S. government and is not subject to copyright protection in the United States. All but Part Two may be reproduced and distributed in its entirety without further permission from GAO. However, because the joint publication may contain copyrighted images or other material, permission from the copyright holder may be necessary if you wish to reproduce this material separately.

The National Academy of Medicine is the author of Part Two and waives its copyright rights for that material.

Foreword

The U.S. health care system is at an important crossroads as it faces major demographic shifts and burgeoning costs. Every day more than 10,000 Americans turn age 65, becoming eligible for Medicare. Although the growth in health care costs has moderated recently, total annual (non-pandemic) health care spending in the United States is projected to reach nearly \$6 trillion by 2027. Additionally, the COVID-19 pandemic has resulted in catastrophic loss of life and substantial damage to the global economy, stability, and security; in response, Congress has appropriated \$2.6 trillion in emergency assistance to date. These realities help illustrate the critical need to better address the effectiveness and efficiency of our nation's health care delivery systems.

Artificial intelligence (AI) represents a set of technologies that includes automated systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, and decision-making. AI has promising applications in health care, including in augmenting patient care. For example, it may have the potential to improve treatment, reduce burden on providers, and generally increase the efficiency with which health care facilities and providers use resources, resulting in potential cost savings or health gains. However, as might be expected with a tool with such broad potential use in health and health care decision-making, applying AI tools for health and health care also raises ethical, legal, economic, and social questions.

The Government Accountability Office (GAO) and the National Academy of Medicine (NAM), individually and in collaboration, have taken up the charge to explore AI in augmenting patient care both inside and outside traditional clinical settings, assess its implications, and identify key policy options available for optimizing its use. In recognition of mutual interests and obligations, and to reinforce and complement each other's work, GAO and NAM have cooperated on the development of two publications. The first is GAO's Technology Assessment: *Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care*, presented as Part One. The second is NAM's Special Publication: *Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic*, presented as Part Two. Any recommendations in Part Two are those of NAM alone.

This cooperative effort included an expert meeting, bringing diverse, interdisciplinary, and cross-sectoral perspectives to the discussions. We are grateful to the exceptionally talented staff of GAO and NAM as well as the experts, all of whom worked hard with enthusiasm, great skill, flexibility, clarity, and drive to make this joint publication possible.

Sincerely,



Karen L. Howard, PhD
Director,
Science, Technology Assessment, and Analytics
U.S. Government Accountability Office



J. Michael McGinnis, MD, MA, MPP
Leonard D. Schaeffer Executive Officer, and
Executive Director, NAM Leadership
Consortium



Executive Summary

This report is being jointly published by the Government Accountability Office (GAO) and the National Academy of Medicine (NAM). Part One of this joint publication is the full presentation of GAO's Technology Assessment: *Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care*. Part Two is the full presentation of NAM's Special Publication: *Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic*. Although GAO and NAM staff consulted with and assisted each other throughout this work, reviews were conducted by GAO and NAM separately and independently, and authorship of the text of Part One and Part Two of this Executive Summary and the following report lies solely with GAO and NAM, respectively.

OVERVIEW OF PART ONE – GAO Technology Assessment: Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care

The GAO report *Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care* is the second in a series of technology assessments that GAO is conducting at the request of Congress on the use of AI technologies in health care.¹ This report discusses three topics: (1) current and emerging AI tools available for augmenting patient care and their potential benefits, (2) challenges to the development and adoption of these tools, and (3) policy options to maximize benefits and mitigate challenges to the use of AI tools to augment patient care.

One of the report's high-level findings is that AI-enabled tools have shown promise in both clinical and administrative applications. Use of these tools could improve patient treatment, reduce burden on providers, and increase resource efficiency at health care facilities, among other potential benefits. Developers have demonstrated these tools in a number of clinical applications, such as supporting clinical decision-making. These tools are at varying stages of maturity and adoption, but with the exception of population health management tools, many we describe have not achieved widespread use. Use of AI tools for administrative applications could also affect patient care, including by reducing provider burden, and are also at varying stages of maturity and adoption, ranging from emerging to widespread.

¹Part One of this Joint Publication presents the GAO Technology Assessment: *Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care*. Although NAM staff and leadership provided assistance and advice in the identification of issues and experts consulted during the development process (identified in app. II), responsibility for the text, findings, and options lies solely with GAO.

We identified five categories of clinical applications where AI tools have shown promise to augment patient care: predicting health trajectories, recommending treatments, guiding surgical care, monitoring patients, and supporting population health management. These are at various stages of implementation. For example, developers are scaling machine learning-based clinical decision support software that can predict sepsis from individual to multiple health systems, while AI-enabled surgical robots are in the early stages of development, with one expert estimating another 20-30 years of development being necessary before some level of automated AI surgeon might be available.²

We also identified three categories of administrative applications where AI tools have shown promise to reduce provider burden and increase the efficiency of patient care: recording digital clinical notes, optimizing operational processes, and automating laborious tasks. These tools are also at various stages of technological maturity. For example, speech recognition software is in widespread use, while AI-powered robots have only been minimally deployed to hospitals.

This technology assessment also identifies challenges that hinder the adoption and impact of AI tools to augment patient care, according to stakeholders, experts, and the literature. Difficulties accessing sufficient high-quality data may hamper innovation in this space. Further, some available data may be biased, which can reduce the effectiveness and accuracy of the tools for some people. Addressing bias can be difficult because the electronic health data do not currently represent the general population. It can also be challenging to scale tools up to multiple locations and integrate them into new settings because of differences in institutions and the patient populations they serve. The limited transparency of AI tools used in health care can make it difficult for providers, regulators, and others to determine whether an AI tool is safe and effective. A greater dispersion of data across providers and institutions can make securing patient data difficult. Finally, one expert described how existing case law does not specifically address AI tools, which can make providers and patients reticent to adopt them. Some of these challenges are similar to those identified previously by GAO in its first publication in this series, such as the lack of high-quality, structured data, and others are more specific to patient care, such as liability concerns.³

²We refer to experts specifically as those people who attended our expert meeting, and refer to anyone else we interviewed as either a stakeholder or by their entity affiliation throughout this summary and the remainder of the report.

³GAO, *Artificial Intelligence in Health Care: Benefits and Challenges of Machine Learning in Drug Development*, GAO-20-215SP (Washington, D.C.: Jan. 31, 2020).

GAO described six options for policymakers—which GAO defines broadly to include Congress, elected officials, federal agencies, state and local governments, academic and research institutions, and industry, among others—to use in addressing these challenges:

- **Collaboration.** Policymakers could encourage interdisciplinary collaboration between developers and health care providers. This could result in AI tools that are easier to implement and use within an existing workflow.
- **Data Access.** Policymakers could develop or expand high-quality data access mechanisms. This could help developers address bias concerns by ensuring data are representative, transparent, and equitable.
- **Best Practices.** Policymakers could encourage relevant stakeholders and experts to establish best practices (such as standards) for development, implementation, and use of AI technologies. This could help with deployment and scalability of AI tools by providing guidance on data, interoperability, bias, and formatting issues.
- **Interdisciplinary Education.** Policymakers could create opportunities for more workers to develop interdisciplinary skills. This could allow providers to use AI tools more effectively, and could be accomplished through a variety of methods, including changing medical curricula or grants.
- **Oversight Clarity.** Policymakers could collaborate with relevant stakeholders to clarify appropriate oversight mechanisms. Predictable oversight could help ensure that AI tools remain safe and effective after deployment and throughout their lifecycle.
- **Status Quo.** Policymakers could allow current efforts to proceed without intervention.

OVERVIEW OF PART TWO – NAM: Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic

Artificial intelligence (AI) is a promising and rapidly developing field, and AI-driven technologies are becoming increasingly prevalent in health care settings. Current AI algorithms support diagnostic and prognostic assessment in many medical specialties, including radiology and dermatology. While these applications exist in typical hospital and clinic settings, AI presents extraordinary opportunities for health monitoring, intervention, and promoting overall well-being outside the hospital and clinic. The authors of this paper focus on health-related applications of AI specifically in these environments, including the home, work, and community settings, and refer to them as “health settings outside the hospital and clinic,” abbreviated HSOHC. This paper aims to provide an analysis of: 1) current technologies and future applications of AI in HSOHC, 2) the logistical steps and challenges involved in integrating AI-HSOHC applications into existing provider workflows, and 3) the ethical and legal considerations

of such AI tools, followed by a brief proposal of potential key initiatives to guide the development and adoption of AI in HSOHC.

Numerous AI-powered health applications designed for personal use have been shown to improve patient outcomes, building predictions based on large volumes of granular, real-time, and individualized behavioral and medical data. For instance, some forms of telehealth, a technology that has been critical during the COVID-19 pandemic, benefit considerably from AI software focused on natural language processing, which enables efficient triaging of patients based on urgency and type of illness. Beyond patient-provider communication, AI algorithms relevant to diabetic and cardiac care have demonstrated remarkable efficacy in helping patients manage their blood glucose levels in their day-to-day lives and in detecting cases of atrial fibrillation. AI tools that monitor and synthesize longitudinal patient behaviors are also particularly useful in psychiatric care, where of the exact timing of interventions is often critical. For example, smartphone-embedded sensors that track location and proximity of individuals can alert clinicians of possible substance use, prompting immediate intervention. On the population health level, these individual indicators of activity and health can be combined with environmental- and system-level data to generate predictive insight into local and global health trends. The most salient example of this may be the earliest warnings of the COVID-19 outbreak, issued in December 2019 by two private AI technology firms.

Successful implementation and widespread adoption of AI applications in HSOHC requires careful consideration of several key issues related to personal data, algorithm development, and health care insurance and payment. Chief among them are data interoperability, standardization, privacy, ameliorating systemic biases in algorithms, reimbursement of AI-assisted services, quality improvement, and integration of AI tools into provider workflows. Overcoming these challenges and optimizing the impact of AI tools on clinical outcomes will involve engaging diverse stakeholders, deliberately designing AI tools and interfaces, rigorously evaluating clinical and economic utility, and diffusing and scaling algorithms across different health settings. In addition to these potential logistical and technical hurdles, it is imperative to consider the legal and ethical issues surrounding AI, particularly as it relates to the fair and humanistic deployment of AI applications in HSOHC. Important legal considerations include the appropriate designation of accountability and liability of medical errors resulting from AI-assisted decisions for ensuring the safety of patients and consumers. Key ethical challenges include upholding the privacy of patients and their data—particularly with regard to non-HIPAA covered entities involved in the development of algorithms—building unbiased AI algorithms based on high-quality data from representative populations, and ensuring equitable access to AI technologies across diverse communities.

The authors of this paper believe that AI technologies will lead to an inevitable and fundamental shift in how health care is delivered in the U.S. This shift, however, requires transitioning from a model that is hospital- and clinic-centric to one that is decidedly more patient-focused and benefits from the richness of health data collected in the environments where individuals actually live, work, and play. Indeed, AI applications in HSOHC show exceptional promise in


ultimately improving the quality of life of individual patients, but ensuring efficient implementation, broad adoption, and fair access to these tools will be contingent upon close public-private collaboration and insurance reform, along with major governmental and policy initiatives.

Table of Contents

Foreword.....	i
Executive Summary.....	iii
OVERVIEW OF PART ONE – GAO Technology Assessment: Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care.....	iii
OVERVIEW OF PART TWO – NAM: Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic	v
GAO Highlights	13
Why GAO did this study.....	13
What GAO found	13
Introduction	1
Part One – (GAO) Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care	3
1 Background	3
1.1 AI Systems.....	3
1.2 Development of AI Systems.....	3
2 Promise and Status of AI Tools to Augment Patient Care	7
2.1 Promise of AI Tools to Augment Patient Care	7
2.2 Clinical AI Tools to Augment Patient Care	8
2.3 Administrative AI tools to Augment Patient Care	17
3 Challenges Surrounding AI Tools to Augment Patient Care	23
3.1 Difficulties Accessing High-Quality Data.....	24
3.2 Potential Bias in Data.....	27
3.3 Difficulties in Scaling	28
3.4 Limited Transparency of AI Tools	29
3.5 Difficulty Protecting Patient Privacy	30
3.6 Uncertainty about Liability for AI Tools.....	33
4 Policy Options to Enhance Benefits or Address Challenges of AI Tools to Augment Patient Care.....	34
Policy Option: Collaboration.....	35
Policy Option: Data Access	36
Policy Option: Best Practices	37

Policy Option: Interdisciplinary Education	38
Policy Option: Oversight Clarity.....	39
Policy Option: Status Quo.....	40
5 Agency and expert comments	41
Part Two—(NAM) Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic.....	45
Introduction and Scope.....	46
1 Surveying Key Examples of AI Outside the Hospital and Clinic Setting: Evaluating Current and Emerging Technologies	49
1.1 Implementing AI on the Individual Level for Better Personal Health.....	49
2 Leveraging AI and Patient-Level Data from Remote Monitoring Tools to Gather Population-Level Insight	54
2.1 Integrating AI into Population Health Strategies	54
2.2 Improving Medication Adherence with AI Tools	55
2.3 AI Efforts in Public and Environmental Health	55
2.4 Combating COVID-19 with AI-Assisted Technologies	56
3 Development and Integration of Health-Related AI Tools: Overarching Logistical Challenges and Considerations.....	58
3.1 Data Interoperability and Standardization	58
3.2 Data Handling and Privacy Protection.....	58
3.3 Systematic Biases in AI Algorithm Development.....	60
3.4 Insurance and Health Care Payment Reform	61
3.5 Quality Improvement and Algorithm Updates	62
3.6 AI Tool Integration into Provider Workflows Outside of the Hospital or Clinic	63
3.7 Practical Steps for Integration of New AI Tools into the U.S. Health System	64
4 Equitable and Humanistic Deployment of Health- Related AI Tools: Legal and Ethical Considerations	66
4.1 Safety Oversight.....	66
4.2 Accountability and Liability Issues	67
4.3 Data Privacy and Individual Consent	68
4.4 Ensuring Equitable Use of AI in Health.....	69

5 Setting the Stage for Impactful AI Tools in HSOCH: Calls to Action..... 71
References 72
Appendix I: Objectives, Scope, and Methodology..... 89
Appendix II: Expert Participation 93
Appendix III: GAO Contacts and Staff Acknowledgments 95



PART ONE

Technology Assessment: Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care

U. S. Government Accountability Office
(GAO)

Part One of this Joint Publication presents the GAO Technology Assessment: *Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care*. Although NAM staff and leadership provided assistance and advice in the identification of issues and expertise consulted during the development process (identified in Appendix 2), responsibility for the text, findings, and options lies solely with GAO.

GAO Highlights

Highlights of [GAO-21-7SP](#), a report to congressional requesters

November 2020

Why GAO did this study

The U.S. health care system is under pressure from an aging population, rising disease prevalence, and increasing costs. New technologies, such as AI, could augment patient care in health care facilities, including outpatient and inpatient care, emergency services, and preventative care. However, the use of AI tools in health care raises a variety of ethical, legal, economic, and social concerns.

GAO was asked to conduct a technology assessment on the use of AI technologies to improve patient care, with an emphasis on foresight and policy implications. This report discusses (1) current and emerging AI tools available for augmenting patient care and their potential benefits, (2) challenges surrounding the use of these tools, and (3) policy options to address challenges or enhance benefits of the use of these tools.

GAO assessed AI tools developed for or used in health care facilities; interviewed a range of stakeholder groups including government, health care, industry, academia, and a consumer group; convened a meeting of experts in collaboration with the National Academy of Medicine; and reviewed key reports and scientific literature. GAO is identifying policy options in this report.

View [GAO-21-7SP](#). For more information, contact Karen L. Howard at (202) 512-6888 or howardk@gao.gov.

TECHNOLOGY ASSESSMENT


ARTIFICIAL INTELLIGENCE IN HEALTH CARE

Benefits and Challenges of Technologies to Augment Patient Care

What GAO found

Artificial Intelligence (AI) tools have shown promise for augmenting patient care in the following two areas:

- **Clinical AI tools** have shown promise in predicting health trajectories of patients, recommending treatments, guiding surgical care, monitoring patients, and supporting population health management (i.e., efforts to improve the health outcomes of a community). These tools are at varying stages of maturity and adoption, but many we describe, with the exception of population health management tools, have not achieved widespread use.
- **Administrative AI tools** have shown promise in reducing provider burden and increasing efficiency by recording digital notes, optimizing operational processes, and automating laborious tasks. These tools are also at varying stages of maturity and adoption, ranging from emerging to widespread.



Source: GAO. | [GAO-21-7SP](#)

GAO identified the following challenges surrounding AI tools, which may impede their widespread adoption:

- **Data access.** Developers experience difficulties obtaining the high-quality data needed to create effective AI tools.
- **Bias.** Limitations and bias in data used to develop AI tools can reduce their safety and effectiveness for different groups of patients, leading to treatment disparities.
- **Scaling and integration.** AI tools can be challenging to scale up and integrate into new settings because of differences among institutions and patient populations.
- **Lack of transparency.** AI tools sometimes lack transparency—in part because of the inherent difficulty of determining how some of them work, but also because of more controllable factors, such as the paucity of evaluations in clinical settings.
- **Privacy.** As more AI systems are developed, large quantities of data will be in the hands of more people and organizations, adding to privacy risks and concerns.
- **Uncertainty over liability.** The multiplicity of parties involved in developing, deploying, and using AI tools is one of several factors that have rendered liability associated with the use of AI tools uncertain. This may slow adoption and impede innovation.

GAO developed six policy options that could help address these challenges or enhance the benefits of AI tools. The first five policy options identify possible new actions by policymakers, which include Congress, elected officials, federal agencies, state and local governments, academic and research institutions, and industry. The last is the status quo, whereby policymakers would not intervene with current efforts. See below for details of the policy options and relevant opportunities and considerations.

Policy Options to Address Challenges or Enhance Benefits of AI to Augment Patient Care

	Opportunities	Considerations
<p>Collaboration (report page 32) Policymakers could encourage interdisciplinary collaboration between developers and health care providers.</p>	<ul style="list-style-type: none"> • Could result in AI tools that are easier to implement and use within a providers’ existing workflow. • Could help implement tools on a larger scale. • Approaches to encourage collaboration include agencies seeking input from innovators. For example, agencies have used a challenge format to encourage the public to develop innovative technologies. 	<ul style="list-style-type: none"> • May result in the creation of tools that are specific to one hospital or provider. • Providers may not have time to both collaborate and treat patients.
<p>Data Access (report page 33) Policymakers could develop or expand high-quality data access mechanisms.</p>	<ul style="list-style-type: none"> • A “data commons”—a cloud based-platform where users can store, share, access, and interact with data—could be one approach. • More high-quality data could facilitate the development and testing of AI tools. • Could help developers address bias concerns by ensuring data are representative, transparent and equitable. 	<ul style="list-style-type: none"> • Would likely require resources to successfully coordinate across different domains. • Organizations with proprietary data could be reluctant to participate. • Cybersecurity and privacy risks could increase, and threats would likely require additional precautions.
<p>Best Practices (report page 34) Policymakers could encourage relevant stakeholders and experts to establish best practices (such as standards) for development, implementation, and use of AI technologies.</p>	<ul style="list-style-type: none"> • Could help providers deploy AI tools by providing guidance on data, interoperability, bias, and implementation, among other things. Could help improve scalability of AI tools by ensuring data are formatted to be interoperable. • Could address concerns about bias by encouraging wider representation and transparency. 	<ul style="list-style-type: none"> • Could require consensus from many public- and private-sector stakeholders, which can be time- and resource-intensive. • Some best practices may not be widely applicable because of differences across institutions and patient populations.
<p>Interdisciplinary Education (report page 35) Policymakers could create opportunities for more workers to develop interdisciplinary skills.</p>	<ul style="list-style-type: none"> • Could help providers use tools effectively. • Could be implemented in a variety of ways, including through changing academic curriculums or through grants. 	<ul style="list-style-type: none"> • Employers and university leaders may have to modify their existing curriculums, potentially increasing the length of medical training.
<p>Oversight Clarity (report page 36) Policymakers could collaborate with relevant stakeholders to clarify appropriate oversight mechanisms.</p>	<ul style="list-style-type: none"> • Predictable oversight could help ensure that AI tools remain safe and effective after deployment and throughout their lifecycle. • A forum consisting of relevant stakeholders could help recommend additional mechanisms to ensure appropriate oversight of AI tools. 	<ul style="list-style-type: none"> • Soliciting input and coordinating among stakeholders, such as hospitals, professional organizations, and agencies, may be challenging. • Excess regulation could slow the pace of innovation.
<p>Status Quo (report page 37) Policymakers could maintain the status quo (i.e., allow current efforts to proceed without intervention).</p>	<ul style="list-style-type: none"> • Challenges may be resolved through current efforts. • Some hospitals and providers are already using AI to augment patient care and may not need policy-based solutions to continue expanding these efforts. • Existing efforts may prove more beneficial than new options. 	<ul style="list-style-type: none"> • The challenges described in this report may remain unresolved or be exacerbated. For example, fewer AI tools may be implemented at scale and disparities in use of AI tools may increase.

Source: GAO.

Abbreviations

AI	artificial intelligence
CMS	Centers for Medicare & Medicaid Services
CDS	clinical decision support
COVID-19	Coronavirus Disease 2019
HHS	Department of Health and Human Services
EHR	electronic health record
ER	emergency room
FDA	Food and Drug Administration
HIPAA	Health Insurance Portability and Accountability Act of 1996
HIV	human immunodeficiency virus
ICU	intensive care unit
MRI	magnetic resonance imaging
NAM	National Academy of Medicine
PPE	personal protective equipment
PHI	protected health information
SaMD	software as a medical device



Introduction

November 30, 2020

Congressional Requesters

The U.S. health care system is under pressure. People age 65 and older are projected to make up one-fifth of the U.S. population by 2030. The overall prevalence of disease is increasing, even setting aside the recent Coronavirus Disease 2019 (COVID-19) pandemic.⁴ Further, individuals, health insurers, and federal and state governments spent approximately \$3.5 trillion in 2018 on health consumption expenditures, representing 16.9 percent of the nation's gross domestic product. One potential tool for addressing concerns surrounding the quality and cost of health care is emerging from the massive volume of health data, which is increasing at an unprecedented rate. Humans alone are not capable of meaningfully analyzing this flood of data on a reasonable time scale. Artificial intelligence (AI) is a promising alternative that can rapidly process and analyze large amounts of complex data.

New AI tools have the potential to both reduce administrative burdens and improve treatment, including outpatient and inpatient care, emergency services, and preventative care. Examples of AI being used to augment patient care include systems that provide personalized treatment recommendations, software that interprets vital signs to monitor patients in intensive care units (ICU), and smart speakers that convert words spoken during a medical appointment into electronic health records (EHR) and codes for insurance billing.⁵ However, the use of AI technologies in health care raises a variety of ethical, legal, economic, and social concerns. For example, AI tools developed using historical data could perpetuate biases such as underrepresentation of certain groups based on race, socioeconomic status, or gender.

In view of the potential for AI to help improve patient care, you asked us to conduct a technology assessment in this area, with an emphasis on foresight and policy implications. This report discusses (1) current and emerging AI tools available for augmenting patient care and their potential benefits, (2) challenges surrounding the use of these tools, and (3) policy options to address challenges or enhance benefits of the use of these tools.

⁴According to the Centers for Disease Control and Prevention, COVID-19 is a respiratory disease caused by a new coronavirus. The complete clinical picture of COVID-19 is not fully known. Reported illnesses have ranged from very mild to severe, including illness resulting in death. The nature of the COVID-19 pandemic and treatment of patients is an evolving situation.

⁵An EHR is a digital medical record and can contain the medical and treatment history of a patient, diagnoses, medications, treatment plans, immunization dates, allergies, radiology images, and laboratory and test results. EHRs are widely used in the health care system to give providers access to information on their patients. They can also automate certain tasks and supply evidence-based tools for making decisions about a patient's care.

To address these objectives, we assessed available and developing AI technologies for clinical or administrative purposes that companies or health care providers may use to augment patient care as well as the challenges associated with using such technologies.⁶ We focused our review on selected technologies used at locations that employ health care providers, including but not limited to physicians, registered nurses, medical assistants, and physical therapists. We excluded technologies used in other environments, such as the home, and excluded technologies that are exclusively focused on diagnostics.⁷

We reviewed key reports and scientific literature describing current and emerging technologies and interviewed a variety of stakeholders, including agency officials, industry members, academic researchers, and a consumer group. We also collaborated with the National Academy of Medicine to convene a 2-day expert meeting on current and emerging AI technologies for augmenting patient care. The meeting included experts from academia and industry, as well as legal scholars, with expertise covering all significant areas of our review. Following the meeting, we continued to use the experts' advice to clarify and expand on what we heard. We then identified six policy options in response to the challenges identified during our work and examined potential opportunities and considerations of each. Consistent with our quality assurance framework, we provided the experts and relevant agencies with a draft of our report and solicited their feedback, which we incorporated as appropriate. See appendix I for additional information on our scope and methodology.

We conducted our work from November 2019 through November 2020 in accordance with all sections of GAO's Quality Assurance Framework that are relevant to technology assessments. The framework requires that we plan and perform the engagement to obtain sufficient and appropriate evidence to meet our stated objectives and to discuss any limitations to our work. We believe that the information and data obtained, and the analysis conducted, provide a reasonable basis for any findings and conclusions in this product.

⁶GAO did not evaluate the effectiveness or utility of any equipment—software or hardware—mentioned in this report.

⁷Diagnostic-focused AI tools will be the subject of the third report in this series.

Part One – (GAO) Artificial Intelligence in Health Care: Benefits and Challenges of Technologies to Augment Patient Care

1 Background

1.1 AI Systems

The field of AI was founded on the idea that machines could be used to simulate human intelligence. Early AI technologies were often expert or rules-based systems, whereby a computer is programmed based on expert knowledge or criteria and produces outputs consistent with its programming. Software programs that do tax preparation or logistics scheduling are examples of expert systems. More recently, a second wave of AI, known as machine learning, has led to significant new capabilities. Machine learning begins with data—generally in large amounts—and infers rules or decision procedures that aim to predict specified outcomes. This inference happens when the system is able to train itself using the data to increase the accuracy of its predictions. Increased availability of large data sets and computing power has enabled recent machine learning advances such as voice recognition by personal assistants on smart phones (an example of natural language processing) and image recognition (an example of computer vision).

Researchers use several methods to train machine learning algorithms, including:

- **Supervised machine learning.** An algorithm with labeled data or input identifies logical patterns in the data and uses those patterns to predict a specified answer to a problem. For example, an algorithm trained on many labeled images of cats and dogs could then classify new,

unlabeled images as containing either a cat or a dog.

- **Unsupervised machine learning.** An algorithm with unlabeled data that allows the algorithm to identify structure in the inputs, for example by clustering similar data, without a preconceived idea of what to expect. In this technique, an algorithm could, for example, cluster images into groups based on similar features, such as a group of cat images and a group of dog images, without being told that the images in the training set are those of cats or dogs.
- **Semisupervised learning.** An algorithm with a training set that is partially labeled uses the labeled data to determine a pattern and apply labels to the remaining data.
- **Reinforcement learning.** An algorithm performs actions and receives rewards or penalties in return. The algorithm learns by developing a strategy to maximize rewards.

1.2 Development of AI Systems

The development of AI tools to address health care challenges is a complex process that varies for each tool. One reason is that funding sources and people involved may vary over time. For example, funding sources may include health systems, venture capital, or government grants. A tool might initially be developed in a university, then licensed to another organization, and ultimately scaled

and deployed by a commercial entity. Users of the tools also vary. For example, a user may be a clinician who works directly with patients, or an administrator in a health system who makes decisions about health

management of a broader population in the system (see fig. 1).⁸

Figure 1: Potential users of artificial intelligence (AI) tools and how they might interact with the tools



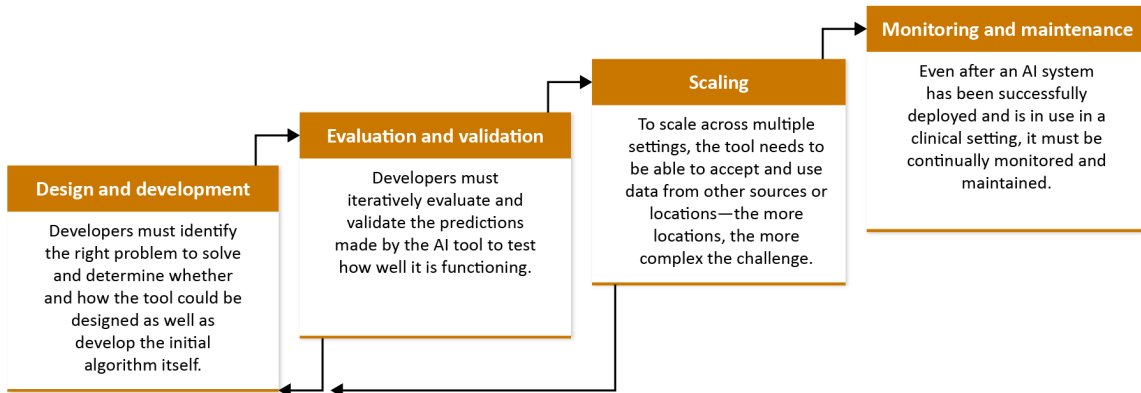
Sources (photos from top to bottom): GAO; Jacob Lund/stock.adobe.com; GAO; GAO; (illustrations): GAO. | GAO-21-7SP

⁸A health system is an organization that includes at least one hospital and at least one group of physicians who are connected through common ownership or joint management.

Development of AI tools involves a multi-stage process that is not necessarily linear and may be iterative, both within and across stages. We provide a highly simplified description of such stages for machine-learning tools below (see fig. 2).⁹

training and test sets, ensuring that there is adequate representation of different categories of interest, and identifying low-quality or irrelevant data.¹⁰ Finally, the developers select the most suitable type of algorithm based on factors such as the types and quantity of data it will

Figure 2: Artificial Intelligence (AI) product development stages



Source: GAO. | GAO-21-7SP

- Design and development.** Developers of AI tools for health care must first identify the problem they wish to solve and determine whether and how the tool could be designed to help solve the problem. Developers must also identify appropriate data sets, a key source of which is patient data from EHRs. Data are needed both to develop the model and, later in the process, to evaluate it. These data must also be curated—a sometimes labor-intensive process involving steps such as appropriately dividing data into

process, how understandable the algorithm is (interpretability), and time taken to apply what the algorithm learned to new data.

- Evaluation and validation.** Next, developers must iteratively evaluate and validate the predictions made by the AI tool to test how well it is functioning.¹¹ Evaluation is based on discrimination—the ability of the model to correctly rank or distinguish categories—and calibration, which measures how well its predicted

⁹These stages can be characterized in different ways. We include one potential way to characterize them. This characterization is adapted from Sendak et al. It applies to machine learning-based AI tools; expert systems-based AI tools may be different. M. P. Sendak, J. D’Arcy, S. Kashyap, M. Gao, M. Nichols, K. Corey, W. Ratliff, and S. Balu, “A Path for Translation of Machine Learning Products into Healthcare Delivery,” *European Medical Journal Innovations*, (2020).

¹⁰In order to know whether a model is generalizable beyond the initial data used to create it, a developer must try that model on new cases. This involves taking the full data set and splitting it into two pieces: a training set, used to initially train the model, and a test set, used to test the model trained on the training data.

¹¹Evaluation is the process by which a developer selects the best machine learning model for their needs, and validation is a set of steps taken to accomplish this.

probabilities match the actual probabilities. The initial round of evaluation and validation can be done by computer simulation using retrospective data.¹² Validation by computer simulation uses multiple models with different parameters on reduced training sets to determine which models make the most accurate predictions. However, high performance during these validation studies is not sufficient to demonstrate clinical validity or impact. Later rounds take place in production environments such as clinical settings, which can differ dramatically from the environments that store retrospective data. To measure clinical validity or impact, these tools may have to undergo additional prospective studies.¹³

- **Scaling.** Many AI systems are initially designed to solve a problem at one health

care system, based on the patient population specific to that location and problem. To scale across different settings, the tool needs to be able to accept and use data from other sources or locations—the more locations, the more complex the challenge.

- **Monitoring and maintenance.** Even after an AI system has been successfully deployed and is in use in a clinical setting, it must be continually monitored and maintained. This is necessary in order to regularly improve the AI systems when, for example: new models are developed, scientific understanding of the disease improves, and new data become available. Iteration of AI systems may also be warranted if regulatory frameworks evolve.

¹²Retrospective data are existing data that were recorded prior to the work described here.

¹³Prospective studies are those where the outcome has not occurred when the study starts, and participants are followed up with over a period of time to determine the occurrence of outcomes.

2 Promise and Status of AI Tools to Augment Patient Care


AI-enabled tools have shown promise in both clinical and administrative applications. Use of these tools could improve patient treatment, reduce burden on providers, and increase resource efficiency at health care facilities, among other potential benefits. Developers have demonstrated these tools in a number of clinical applications, such as supporting clinical decision-making. However, many of the tools we discuss have not yet been widely adopted, with the exception of a tool supporting population health management. Use of AI tools for administrative applications can also affect patient care. The tools we discuss are at varying stages of maturity and adoption, ranging from emerging to widespread technologies.

2.1 Promise of AI Tools to Augment Patient Care

AI in health care has the potential to deliver many benefits, according to the scientific literature and stakeholders, including industry representatives, academic researchers, and health care providers. In general, AI tools augment rather than replace human providers. Studies have demonstrated improved results when providers and AI tools work together rather than each working independently. Examples of key areas where AI tools show promise to augment patient care (see fig. 3) include:

- **Improving treatment.** AI tools could improve provider decision-making with more accurate predictions of a patient's health trajectory. For example, providers could use AI to predict hospital length-of-stay, readmission, or mortality. Availability of large amounts of health data and advanced analytical tools, such as those using AI, also promise significant advances in precision medicine (i.e., the tailoring of medical treatment to the individual characteristics of each patient). The goal of precision medicine is to allow providers to select patient-specific treatments that minimize harmful side effects and ensure a more successful outcome. By integrating individual patient information with lessons gleaned from large volumes of data on prior patient cases and clinical trajectories, AI tools could assist provider decision-making with a greater comprehensiveness and speed than would be possible without such tools.

Figure 3: How AI could change the patient experience of the future



Source: GAO. | GAO-21-7SP

- **Reducing burden on providers.** By taking over routine or standardized tasks not requiring human judgement, intuition, or empathy, AI tools could also reduce stress and free providers to spend time on more complex tasks. Provider burnout—a long-term stress reaction marked by emotional exhaustion, depersonalization, and a lack of a sense of personal accomplishment—has been on the rise in recent years and can threaten patient safety or quality of care. For example, AI tools that ease the process of documenting clinical notes or automate aspects of the clinical workflow could give providers more time with their patients, thereby enhancing the patient-provider relationship.
- **Increasing resource efficiency.** AI tools could also increase the efficiency with which health care facilities and providers use resources, potentially resulting in cost

savings, health gains, or both. For example, AI tools could reduce the need for costly equipment, inform staffing decisions, or help direct resources to patients in the most need of care.

2.2 Clinical AI Tools to Augment Patient Care

We identified five categories of clinical applications in which AI tools have shown promise to augment patient care: predicting health trajectories, recommending treatments, guiding surgical care, monitoring patients, and supporting population health management. The selected tools described in these categories are at varying stages of maturity and adoption.¹⁴ We identified several tools in the design and development phase; these tools have been studied in the





¹⁴The tools described in this section are selected examples intended to demonstrate the breadth of current and emerging

clinical AI tools. This list is not intended to be exhaustive or representative of all clinical AI tools.

scientific literature, but have not yet been clinically validated or tested in humans. The risk-prediction tool for supporting population health management is the only tool we describe that has reached the monitoring and maintenance stage of development. Figure 4

provides a summary of categories and selected tools, which are discussed in more detail after the figure.

Figure 4: Assessment of selected clinical AI tools to augment patient care

	Selected tool, benefits, and limitations	Stage of development
 <p>Predicting health trajectories</p>	<p>Example: Sepsis prediction clinical decision support (CDS) Can predict sepsis and septic shock before onset and has demonstrated improved outcomes. However, additional studies are needed.</p>	<ul style="list-style-type: none"> ▶ Evaluation and validation ▶ Scaling
 <p>Recommending treatments</p>	<p>Example: Mechanical ventilator CDS Could predict when to successfully wean patients from ventilators and may overcome limitations of ventilators with automated weaning modes; however, many studies did not assess effectiveness in a clinical setting.</p>	<ul style="list-style-type: none"> ▶ Design and development
 <p>Guiding surgical care</p>	<p>Example: Planning and post-operative care CDS Can assist with surgical planning or predict risk of complications, such as infections or in-hospital mortality.</p>	<ul style="list-style-type: none"> ▶ Evaluation and validation
	<p>Example: Intra-operative CDS Could use surgical video data to enhance surgical decision-making by predicting adverse events and guiding providers in real time, according to early studies, but necessary technology is still in its infancy.</p>	<ul style="list-style-type: none"> ▶ Design and development
	<p>Example: Autonomous surgical robots Could match or outperform human surgeons for certain tasks, according to early studies in animals, but has not been tested in humans.</p>	<ul style="list-style-type: none"> ▶ Design and development
 <p>Monitoring patients</p>	<p>Example: Patient fall prevention Can analyze movements in a patient's room and alert the care team when a fall is predicted.</p>	<ul style="list-style-type: none"> ▶ Scaling
 <p>Supporting population health management</p>	<p>Example: Risk-prediction Can support efforts to use population-level data to identify broad health risks and treatment opportunities for a group of individuals or a community. However, a recent study uncovered racial bias in a widely used tool.</p>	<ul style="list-style-type: none"> ▶ Monitoring and maintenance

Source: GAO. | GAO-21-7SP

Data table for Figure 4, assessment of selected clinical AI tools to augment patient care

Category	Selected tool, benefits, and limitations	Stage of development
Predicting health trajectories	<p>Example: Sepsis prediction clinical decision support (CDS)</p> <p>Can predict sepsis and septic shock before onset and has demonstrated improved outcomes. However, additional studies are needed.</p>	<p>Evaluation and validation</p> <p>Scaling</p>
Recommending treatments	<p>Example: Mechanical ventilator CDS</p> <p>Could predict when to successfully wean patients from ventilators and may overcome limitations of ventilators with automated weaning modes; however, many studies did not assess effectiveness in a clinical setting.</p>	<p>Design and development</p>
Guiding surgical care	<p>Example: Planning and post-operative care CDS</p> <p>Can assist with surgical planning or predict risk of complications, such as infections or in-hospital mortality.</p>	<p>Evaluation and validation</p>
	<p>Example: Intra-operative CDS</p> <p>Could use surgical video data to enhance surgical decision-making by predicting adverse events and guiding providers in real time, according to early studies, but necessary technology is still in its infancy.</p>	<p>Design and development</p>
	<p>Example: Autonomous surgical robots</p> <p>Could match or outperform human surgeons for certain tasks, according to early studies in animals, but has not been tested in humans.</p>	<p>Design and development</p>
Monitoring patients	<p>Example: Patient fall prevention</p> <p>Can analyze movements in a patient’s room and alert the care team when a fall is predicted.</p>	<p>Scaling</p>
Supporting population health management	<p>Example: Risk-prediction</p> <p>Can support efforts to use population-level data to identify broad health risks and treatment opportunities for a group of individuals or a community. However, a recent study uncovered racial bias in a widely used tool.</p>	<p>Monitoring and maintenance</p>

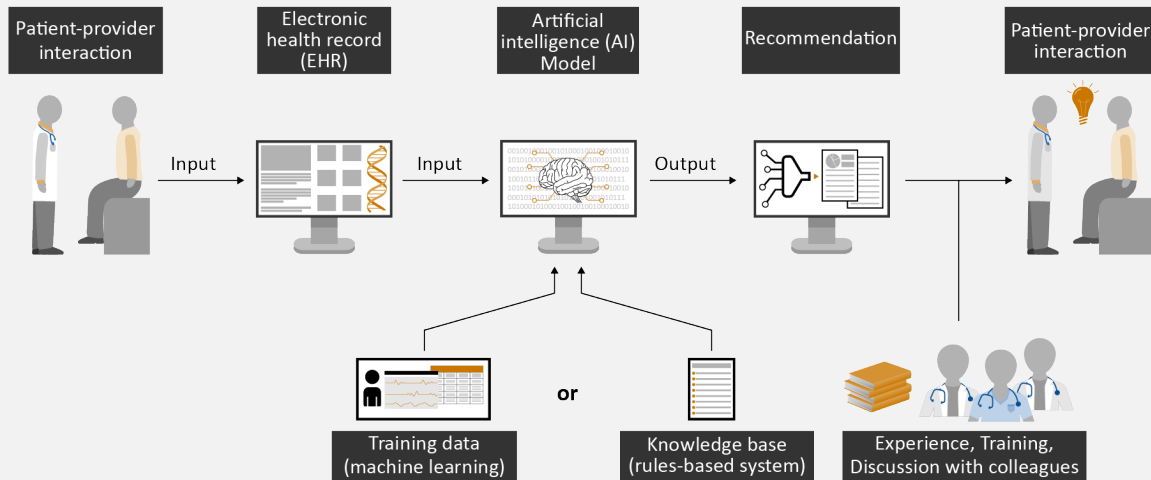
The first three categories include machine-learning-enabled CDS tools (see text box below for more information). According to the scientific literature, many of these tools

are described in academic journals but have not been fully integrated into or evaluated in a clinical setting, and evidence of the clinical or economic effect of these tools is limited.

AI-enabled clinical decision support tools

Clinical decision support (CDS), which encompasses a variety of tools, intelligently filters information or presents it at appropriate times to enhance health and health care (see fig. 5). Not all CDS tools contain AI software—researchers have been exploring the use of rules-based systems, or first wave AI, in CDS tools since the 1970s, but recent interest has focused on incorporating machine learning, or second wave AI, into these tools.

Figure 5: Sample workflow for AI-based clinical decision support system



Source: GAO. | GAO-21-7SP

While the ability of machine-learning-enabled CDS tools to make health-related predictions has shown promise in clinical applications, the development and implementation of these tools is generally still in the early stages of maturity. To date, no machine learning CDS tool using electronic health record (EHR) data has successfully scaled across the health care sector.

Source: GAO analysis of agency documentation and scientific literature. | GAO-21-7SP

2.2.1 Predicting health trajectories

Machine-learning-enabled CDS tools can help predict the likelihood that a patient’s condition will deteriorate. In one example, in 2013-2014 a large integrated health system successfully piloted a machine learning model to identify patients at risk for transfer to the intensive care unit. Other applications in this category include prediction of acute kidney injury and *Clostridioides difficile* infection.¹⁵

One of the most common health concerns targeted by recent machine learning models is sepsis.¹⁶ Machine-learning-based sepsis prediction models can predict sepsis and septic shock before onset and have demonstrated improved patient outcomes (see text box for examples). Several studies have shown that early identification and treatment of patients with severe sepsis or septic shock can reduce morbidity, mortality, and hospital length of stay. Of the 21 case studies identified in a 2020 review, six are studies of sepsis prediction models.¹⁷ Developers have begun or published prospective clinical validation of three of these models, and five have received funding from either private or public sources to begin scaling adoption.

Examples of sepsis prediction models

In a 2015 study, researchers developed a model for predicting the risk of septic shock and identifying at-risk patients hours before onset of the condition. Researchers trained the model using supervised machine learning on publicly available, de-identified intensive care unit (ICU) data. At a specificity of 67 percent and sensitivity of 85 percent, this tool was able to identify patients approximately 28 hours before the onset of septic shock.^a Additionally, the model was able to identify a majority of patients a median of around 7 hours before any sepsis-related organ dysfunction—a point after which mortality rates increase—which is an improvement over routine screening protocols. Study authors validated the model using retrospective data, and among other limitations noted the need for a prospective study to evaluate effects of the tool on therapeutic judgement. The tool is now being commercialized and implemented in two hospitals with three more planned as of 2019.

In another study, researchers conducted a small randomized clinical trial of a different machine-learning-based sepsis prediction model in two ICUs, demonstrating improved patient outcomes. The average length of stay as well as the in-hospital mortality rate decreased in the experimental group as compared to the control group. However, the study authors noted that patient outcomes may have been improved because of increased provider awareness of high-risk patients rather than the early prediction of sepsis.

^aSpecificity, or the true negative rate, describes the tool’s ability to accurately identify a negative, which in this case is a patient who will not go on to develop sepsis. Sensitivity, or the true positive rate, describes the tool’s ability to accurately identify a positive, which in this case is a patient who will go on to develop sepsis.

Source: K. E. Henry, D. N. Hager, P. J. Pronovost, and S. Saria, “A Targeted Real-Time Early Warning Score (TREWscore) for Septic Shock,” *Science Translational Medicine*, vol. 7, no. 299 (2015); and D. W. Shimabukuro, C. W. Barton, M. D. Feldman, S. J. Mataraso, and R. Das, “Effect of a machine learning-based severe sepsis prediction algorithm on patient survival and hospital length of stay: a randomised clinical trial,” *BMJ Open Respiratory Research*, vol. 4 (2017). | GAO-21-7SP.

¹⁵*Clostridioides difficile* is a bacterium that causes diarrhea and colitis (inflammation of the colon). It’s estimated to cause half a million illnesses in the United States each year, according to the Centers for Disease Control and Prevention.

¹⁶Sepsis is the body’s extreme response to an infection. In 2013, it was the most expensive condition treated in U.S. hospitals, costing payers more than \$20 billion, according to the Agency for Healthcare Research and Quality. Sepsis is life-threatening, and without timely treatment, can rapidly lead to tissue damage, organ failure, and death.

¹⁷Sendak, D’Arcy, Kashyap, Gao, Nichols, Corey, Ratliff, and Balu, “A Path for Translation of Machine Learning Products into Healthcare Delivery,” *European Medical Journal Innovations* (2020).

Despite promising results, there are limitations to sepsis prediction models. One researcher we spoke to pointed out that the retrospective data used to develop these models may not be representative of all patients who go on to develop sepsis. For example, a provider may treat a patient arriving at the emergency room (ER) with an early-stage infection with antibiotics, preventing the patient from developing sepsis or delaying the onset. The model may then learn that this patient is low-risk for sepsis, when in fact they were high-risk. Additionally, tools may not work when a new disease or other phenomenon emerges, according to an expert. For example, this expert reported having received many questions about using their sepsis prediction model and related tools during the COVID-19 pandemic. However, the protocol for ordering laboratory tests for COVID-19 patients is different than that of sepsis patients—for COVID-19, providers are ordering many different types of laboratory tests, whereas the expert’s sepsis prediction model uses one specific laboratory test as proxy for suspicion of infection. The emergence of this new virus may require changes to the way providers think about and respond to sepsis, according to this expert.

2.2.2 Recommending treatments

AI-enabled CDS tools can also recommend treatments to health care providers, potentially helping them make decisions more effective and patient-specific. Machine

learning makes it possible to process and use large-scale data from previous cases for clinical decision making in a way that would have been difficult previously. For example, these tools could help personalize treatment decisions for patients by learning from the collective experience of others to identify patients with similar conditions and the outcomes of their treatment. However, when these tools are trained on retrospective data they risk learning the prescribing habits of physicians rather than ideal practices. Examples from the scientific literature include tools used to recommend treatments for cancer, sepsis, and stroke.

Another promising area for AI-enabled CDS tools is in treating patients with mechanical ventilators. While ventilators can be life-saving, both prolonged use and premature weaning are associated with complications, increased mortality rates, and higher hospital costs.¹⁸ Deciding when to wean patients receiving ventilator treatment is an essential aspect of their care. However, there is uncertainty surrounding the best methods for conducting this process. For example, providers are using ventilators to treat COVID-19 patients, and the variation in mortality rates across ICUs suggests that different methods for ventilator management may affect outcomes.¹⁹ Many commercially available ventilators contain automated weaning modes based on rules-based systems. In scientific studies, the use of these tools or clinical guidelines has outperformed the common practice of providers deciding on

¹⁸Weaning refers to the process of taking a patient off of mechanical ventilation support and removing the endotracheal tube.

¹⁹J. J. Marini, L. Gattinoni, “Management of Covid-19 Respiratory Distress,” *Journal of the American Medical Association*, published online April 24, 2020. Accessed May 7, 2020. <https://jamanetwork.com/journals/jama/fullarticle/2765302>.

their own when to wean. However, these tools may produce suboptimal results in a clinical environment depending on what data are available and the assumptions of the rules-based system.

Recent research has explored the use of machine-learning-enabled CDS tools for this application, which may overcome some of the limitations of ventilators with automated weaning modes. Machine learning technology can help build models that incorporate the many factors affecting the respiratory system in the ICU. Supervised machine learning is a commonly used technique to train machine-learning-enabled tools for ventilator weaning management. These tools have shown promise in predicting when to successfully wean patients, according to the scientific literature, although mostly as proof-of-concept rather than assessing the effectiveness of the tools in a clinical setting.²⁰

Researchers have also explored the use of reinforcement learning for this task; however, this technology is less mature. In a recent study, researchers used reinforcement learning to develop a CDS tool that would alert providers when a patient is ready to begin weaning.²¹ Using publicly available, retrospective ICU data, researchers showed that fewer patients had to be put back on a ventilator when providers followed a protocol similar to the one recommended by the CDS tool. However, using machine learning tools trained on outcomes data, such as timing and success of weaning, to guide this process can

also be challenging. For example, successful weaning at a specified time does not preclude the possibility that the patient was ready to wean earlier.

2.2.3 Guiding surgical care

In the field of surgical care, planning and postoperative care are the most mature applications for machine-learning-enabled CDS tools. Other applications, including real-time CDS for surgery and AI-enabled surgical robots, are also active areas of research but are more nascent. One tool for surgical planning that entered clinical trials in 2018 uses computer vision on magnetic resonance imaging (MRI) scans to identify areas for prostate cancer biopsy. In postoperative care, studies have demonstrated that AI tools can predict risk of complications from surgery, such as surgical site infections and in-hospital mortality.

Other surgical applications of AI, such as using machine-learning-enabled CDS tools to enhance surgical decision-making, are also in development but are in their infancy, according to an expert. One application currently under study is the use of computer vision to analyze surgical video in real time to

²⁰M. T. Kwong, G. W. Colopy, A. M. Weber, A. Ercole, J. H. M. Bergmann, "The efficacy and effectiveness of machine learning for weaning in mechanically ventilated patients at the intensive care unit: a systematic review," *Bio-Design and Manufacturing* vol. 2 (2019): p. 31.

²¹N. Prasad, L.-F. Cheng, C. Chivers, M. Draugelis, and B. E. Engelhardt, "A Reinforcement Learning Approach to Weaning of Mechanical Ventilation in Intensive Care Units," Preprint at <https://arxiv.org/abs/1704.06300> (2017). Accessed May 6, 2020.

identify or predict adverse events and guide providers.²²

Another area of active research is AI-enabled surgical robotics; however, most current commercial surgical robots are primarily remote-controlled rather than AI-enabled. According to one expert, there are commercial surgical robots in use or coming to the U.S. market that claim to be AI-enabled, but these generally just contain certain AI-enabled features rather than actually guiding surgery. For example, these robotic tools may alert a surgeon when their surgical time is longer than average or predict how many cartridges they can use for the surgical stapler.

AI-enabled autonomous surgical robots have not yet been tested in humans, but studies in animals have demonstrated that they can match or outperform human surgeons for certain tasks. These tools could potentially increase the safety, efficiency, and access to soft tissue surgical procedures.²³ Early attempts at autonomous robotic surgery focused on automating relatively simple tasks such as suturing or knot-tying, but recently researchers have demonstrated more complex tasks. For example, hospital researchers developed a robotic catheter that can autonomously navigate a beating heart in a pig, and demonstrated that its performance

was similar to that of experienced surgeons.²⁴ Automating such tasks could free the surgeon to focus on other aspects of the procedure, potentially reducing the learning curve involved in mastering a new procedure. AI-enabled surgical robots might also have the effect of allowing less-experienced surgeons to safely perform operations and expanding the skills of providers in under-resourced areas. An expert in this area expects to see some level of automation for soft tissue surgical robots in the next 20 to 30 years. This expert said researchers will need to acquire and label surgical video data to properly develop these machines before allowing them to perform surgical tasks autonomously.

2.2.4 Monitoring patients

AI-enabled tools can use the increasing availability of health data, including data from EHRs, wearables, and other sensors, to help monitor patients in health care facilities. According to a recent review, patient monitoring is one of the areas where AI is likely to have the greatest influence.²⁵ For example, providers can use AI analysis of vital signs for cardiovascular and respiratory monitoring in the ICU.

In another example, health care facilities can use AI-enabled monitoring tools in hospitals to prevent patient falls and reduce provider

²²D. A. Hashimoto, G. Rosman, D. Rus, O. R. Meireles, "Artificial Intelligence in Surgery: Promises and Perils," *Annals of Surgery*, vol. 268, no. 1 (2018).


²³Soft tissue refers to muscle, fat, fibrous tissue, blood vessels, or other supporting tissue of the body. Unlike rigid tissue such as bone, unpredictable changes in soft tissue can pose additional challenges during surgery. According to an expert, in certain specialties, such as orthopedic surgery and neurosurgery, there are currently robots that can autonomously drill through bone given certain guidelines.

²⁴G. Fagogenis, M. Mencattelli, Z. Machaidze, B. Rosa, K. Price, F. Wu, V. Weixler, M. Saeed, J. E. Mayer, P. E. Dupont, "Autonomous robotic intracardiac catheter navigation using haptic vision," *Science Robotics*, vol. 4 (2019).

²⁵S. Reddy, J. Fox, and M. P. Purohit, "Artificial intelligence-enabled healthcare delivery," *Journal of the Royal Society of Medicine*, vol. 112, no. 1 (2019).

burden (see fig. 6). According to a 2015 report, hundreds of thousands of hospital patients fall each year.²⁶ Thirty to 50 percent of those falls result in injury, which can increase the length and cost of the hospital stay or even result in death. One commercial AI tool aiming to help providers address these issues uses computer vision, Bluetooth, and sensors to analyze movements in the patient’s room and alert the care team when a fall is predicted, according to company representatives. They stated the tool can also reduce provider burden by eliminating the need for hourly patient checks, and help predict staffing needs by identifying patients who are visited most frequently. According to company materials, this company also developed new algorithms to use the device for infection control activities in response to the COVID-19 pandemic. For example, the tool can be used to automatically detect whether staff are complying with personal protective equipment (PPE) requirements, and to conduct contact tracing within the health care facility.

Figure 6: Patient monitoring tool at the bedside



Source: Inspiren. | GAO-21-7SP

Note: The patient monitoring tool is the white and green device in the upper right-hand corner of the image.

2.2.5 Supporting population health management

Large health systems and payers widely use AI-enabled tools to support population health management activities.²⁷ Population health management involves using population-level data to identify broad health risks and treatment opportunities for a group of individuals or a community. One task where AI has the potential to improve population health management is the prioritization of clinical resources and targeting health care services to patients who are the most likely to benefit.²⁸ For example, large integrated health systems have deployed machine learning models to predict individual patients’ risk of chronic diseases and associated complications and offer specialized care programs to high-risk patients.

²⁶The Joint Commission, “Preventing falls and fall-related injuries in health care facilities,” *Sentinel Event Alert*, no. 55 (2015).

²⁷Population health management activities can serve either a clinical or administrative function depending on the application. The risk-prediction tools discussed in this section are clinical applications.

²⁸M. Matheny, S. Thadaney Israni, M. Ahmed, and D. Whicher, eds. *Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril* (Washington, DC: National Academy of Medicine, 2019).

However, one risk of using such tools is bias. For example, a commercial risk-prediction tool that large health systems and payers collectively apply to around 200 million people in the United States per year was found to have racial bias in a recent study.²⁹ (For more information on this study see the text box in section 3.2.) According to the study, correcting algorithmic bias requires expert knowledge of the relevant field, the ability to identify and extract relevant data elements, and the capacity to iterate and experiment. However, one of the study's co-authors stated that the problems in machine learning tools can at least be fixed once identified, whereas bias among humans may not be as readily remedied.

2.3 Administrative AI tools to Augment Patient Care




While administrative applications of AI do not always receive the same level of attention as clinical applications, these tools can still address essential aspects of patient care. For example, AI can take over repetitive and routine tasks such as entering patient data, giving the provider more time to spend listening to, examining, and caring for patients. We identified three categories of administrative applications in which AI tools

have shown promise to reduce provider burden and increase the efficiency of patient care: recording digital clinical notes, optimizing operational processes, and automating laborious tasks. Similar to clinical applications, the selected tools described in these categories are at varying stages of maturity and adoption.³⁰ However, speech recognition tools are widespread and have reached the monitoring and maintenance phase of development. Only one of the tools we identified, autonomous documentation support, is considered an emerging technology. Figure 7 provides a summary of categories and selected tools, which are described in more detail after the figure.

²⁹Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, "Dissecting racial bias in an algorithm used to manage the health of populations," *Science*, vol. 366, no. 6464 (2019): 447.

³⁰The tools described in this section are selected examples intended to demonstrate the breadth of current and emerging administrative AI tools. This list is not intended to be exhaustive or representative of all administrative AI tools used in health care settings.

Figure 7: Assessment of selected administrative AI tools to augment patient care

	Selected tool, benefits, and limitations	Stage of development
 Recording digital clinical notes	Example: Speech recognition Can reduce report turnaround time, but has higher error rates than human transcriptionists.	▶ Monitoring and maintenance
	Example: More autonomous documentation support Could use various AI tools to go beyond word-for-word transcription and place structured data in the EHR, but fully autonomous tools are not achievable with current technology, according to a recent review.	▶ Design and development
 Optimizing operational processes	Example: Scheduling optimization Can increase efficiency of resource use by health care facilities and improve the patient experience by decreasing delays.	▶ Evaluation and validation ▶ Scaling
 Automating laborious tasks	Example: Socially intelligent hospital robot assistant Can automate tasks that are simple but are labor-intensive, such as gathering and delivering supplies.	▶ Scaling

Source: GAO. | GAO-21-7SP

Data table for Figure 7, Assessment of selected administrative AI tools to augment patient care

Category	Selected tool, benefits and limitations	Stage of development
Recording digital clinical notes ICON	Example: Speech recognition Can reduce report turnaround time, but has higher error rates than human transcriptionists.	Monitoring and maintenance
	Example: More autonomous documentation support Could use various AI tools to go beyond word-for-word transcription and place structured data in the EHR, but fully autonomous tools are not achievable with current technology, according to a recent review.	Design and development
Optimizing operational processes ICON	Example: Scheduling optimization Can increase efficiency of resource use by health care facilities and improve the patient experience by decreasing delays.	Evaluation and validation Scaling
Automating laborious tasks ICON	Example: Socially intelligent hospital robot assistant Can automate tasks that are simple but are labor-intensive, such as gathering and delivering supplies.	Scaling

2.3.1 Recording digital clinical notes

Providers are beginning to use speech recognition and natural language processing technologies for recording digital notes into EHR systems. Generally, the current generation of EHR systems requires providers to either type or dictate notes. Adoption of EHR systems, although it can reportedly improve care coordination and decision-making, has also been associated with decreased provider satisfaction. In addition, according to a review, it has increased documentation times, reduced quality and length of interaction with patients, and created safety issues, such as corrupted files preventing entry of diagnoses or orders.³¹

To help address some of these problems, a number of products—either on the market now or in development—use speech recognition and natural language processing to transform a provider-patient conversation (and other words spoken during a clinical encounter) into digital notes. The most advanced documentation support tools currently available still rely primarily on the provider to document clinical encounters, but they make the task more simple or effective. Speech recognition tools for transcription appear to reduce turnaround time and are widely used in certain specialties, although studies show that they have higher error rates than human transcriptionists.³² In another application, EHR systems increasingly use speech recognition for providers to interact with software. These speech interaction tools

have had limited evaluation but currently appear to increase error rates and documentation times.

Support tools for documentation with greater autonomy are being tested, but are considered emerging technologies. Using speech recognition to place clinical notes into EHR systems as unstructured data—such as free text—is straightforward, according to a health care system representative we spoke with. However, transforming this information into structured data—such as medical codes—is more challenging for digital tools, as it involves understanding medical terminology and the context of the conversation. According to the representative, their health care system still relies on humans to generate these structured data—transcripts are sent to an external organization where they are transformed into structured data by a doctor. A digital tool may eventually replace humans for this task, but the health care system representative does not expect to see it in the next two to three years. Text summarization methods, using speech recognition, natural language processing, or other kinds of machine learning, could allow documentation support tools to achieve a higher level of automation by going beyond creating a word-for-word transcription of spoken language. However, fully autonomous documentation support tools are not achievable based on available AI technologies, according to a recent review.³³

³¹E. Coiera, B. Kocaballi, J. Halamka, L. Laranjo, “The digital scribe,” *NPJ Digital Medicine*, vol. 1, no. 58 (2018).

³²T. Hodgson and E. Coeira “Risks and benefits of speech recognition for clinical documentation: a systematic review,”

Journal of the American Medical Informatics Association, vol. 23, no. e1 (2016).

³³Coiera, Kocaballi, Halamka, Laranjo, “The digital scribe,” p. 58.

2.3.2 Optimizing operational processes

AI-enabled tools are being used in small-scale pilot projects to support operations management by, for example, optimizing scheduling, staffing, or resource allocation. While health care facilities are increasingly recognizing the importance of analytics to optimize these and related tasks, most of the published work on this topic has not been implemented in the field or was done at the pilot scale, according to one review.³⁴

Use of machine-learning-enabled tools to optimize scheduling has shown promise for increasing the efficiency with which health care facilities use their resources and improving the patient experience. One hospital that is active in this area uses machine learning to optimize scheduling of its 41 operating rooms and to facilitate the flow of people and equipment within the rooms. The machine learning model suggests schedule modifications that could improve efficiency and save costs. It also predicts—and identifies strategies for minimizing—negative effects that schedule changes could have on patient care. The same hospital also uses a machine learning model to identify and send electronic reminders to patients most likely to skip appointments.

In another example, a large cancer treatment center was able to see more patients and decrease their wait times by deploying a data-driven scheduling tool. The center developed

and deployed this tool in collaboration with academic researchers. Cancer patients often receive chemotherapy as well as related or supportive treatments through intravenous infusion or pills administered at outpatient facilities. These appointments can last from 30 minutes to 12 hours and are generally scheduled up to 6 weeks in advance. In 2013, the center was experiencing severe overcrowding during peak times, leading to long waits, frustration among providers, and concerns about safety. The treatment center deployed the scheduling tool in 2017 and has seen several benefits, including an increase in patient throughput of 9.5 percent and a 30 minute decrease in the average time patients spend at the center—signaling fewer operational delays.

2.3.3 Automating laborious tasks

Although many of the clinical and administrative applications discussed in this chapter are decision support tools, AI can also take over certain tasks, such as gathering supplies. According to the scientific literature, AI can automate some tasks that are simple but labor-intensive, allowing providers more time to spend with patients. According to reports from NAM, gathering supplies, documentation, and similar activities consume a significant amount of nurses' time, taking them away from direct patient care.³⁵ One study found that hospital nurses spent the majority of their time walking between patient rooms and the nursing station.

³⁴M. S. Copenhaver, M. Hu, R. Levi, K. Safavi, A. C. Z. Langle, "Health System Innovation: Analytics in Action," *Tutorials in Operations Research* (2019). According to an expert, however, many health care facilities have adopted emergency department triaging tools that improve operational flow.

³⁵These reports were originally published by the Institute of Medicine which has since been renamed the National Academy of Medicine. Institute of Medicine, *Keeping Patients Safe: Transforming the Work Environment of Nurses* (Washington, D.C.: National Academies Press, 2004); Institute of Medicine, *The Future of Nursing: Leading Change, Advancing Health* (Washington, D.C.: National Academies Press, 2011).

Another study found that surgical nurses walked an additional mile while on duty to obtain supplies and equipment. Distractions and interruptions from non-nursing activities such as gathering supplies present a risk to patient safety, according to NAM.

Many hospitals already employ robots for delivering supplies, among other activities, but they have limitations. For example, many of the currently deployed supply delivery robots do not have an arm and therefore still require a human for tasks such as opening doors and picking up items. These robots also do not significantly interact with clinical staff.

To further automate the gathering and delivering of supplies, companies are developing more sophisticated robots with the ability to execute complex manipulation tasks. The use of AI-enabled robots in a clinical care setting is fairly new and dynamic, and the effects on nursing and patient care have not yet been well studied. According to one company, AI provides their hospital robot assistant with social intelligence, mobility, and the ability to learn from humans. The ability of AI technology to safely navigate indoors and around people has been a key advance in the field to enable the use of robots in hospitals and other settings. As of late 2019, this hospital robot assistant had exited beta testing and was to be officially deployed in a hospital, according to trade press (see fig. 8). A provider organization told us that providers have had a positive response to this robot, finding that it increases efficiency and treating it like part of the team. However, they noted that in emergency situations, nurses preferred to retrieve supplies themselves because they did not want to wait for the robot. Health care facilities are also deploying robot assistants for a number of

tasks during the coronavirus pandemic in an attempt to avoid exposing workers.

Figure 8: Hospital robot assistant delivering supplies to a patient's room



Source: Diligent Robotics. | GAO-21-7SP

3 Challenges Surrounding AI Tools to Augment Patient Care

Drawing on information from experts, stakeholders, and the scientific literature, we identified several challenges affecting the use of AI for patient care (see fig. 9). These challenges affect technology developers, health care providers, and patients across the tools' life cycles and may slow the adoption of

such tools. We highlight the following challenges below: difficulties accessing high-quality data, potential bias in data, difficulties in scaling, limited transparency of AI tools, difficulties protecting patient privacy, and uncertainty about liability of AI tools.

Figure 9: Challenges Surrounding AI Tools to Augment Patient Care

 <p>Difficulties Accessing High-Quality Data</p>	<ul style="list-style-type: none"> ▶ Accessing sufficient high-quality data to develop AI tools is a significant challenge. ▶ As a result, innovation in AI tools for augmenting patient treatment is being hampered.
 <p>Potential Bias in Data</p>	<ul style="list-style-type: none"> ▶ Bias in data used to develop AI tools can reduce their effectiveness and accuracy. ▶ Addressing bias is difficult because electronic health data currently available do not represent the general population.
 <p>Difficulties in Scaling</p>	<ul style="list-style-type: none"> ▶ AI tools can be challenging to scale up and integrate into new settings because of differences among institutions and the patient populations they serve.
 <p>Limited Transparency of AI Tools</p>	<ul style="list-style-type: none"> ▶ Both interpretability and explainability pose challenges to explaining an AI tool's decision-making in an understandable way. ▶ This limited transparency can make it difficult or impossible for providers to understand how an AI tool came to a decision and whether and how an error occurred, as well as hampering the development of trust in the AI system.
 <p>Difficulties Protecting Patient Privacy</p>	<ul style="list-style-type: none"> ▶ As more AI systems are developed, large quantities of patient data will be in the hands of more people and organizations. This dispersion of data contributes to patient privacy risks. ▶ Patient advocacy groups and others have raised concerns, such as about the proliferation of potentially sensitive patient data, potentially without patient consent.
 <p>Uncertainty about Liability for AI Tools</p>	<ul style="list-style-type: none"> ▶ There is uncertainty about liability issues related to AI tools for augmenting patient treatment. ▶ The large number of people involved with developing and using AI tools as well as limited transparency of the tools contribute to this uncertainty.

Source: GAO. | GAO-21-7SP

Data table for Figure 9, Challenges Surrounding AI Tools to augment patient care

Difficulties Accessing High-Quality Data

- Accessing sufficient high-quality data to develop AI tools is a significant challenge.
- As a result, innovation in AI tools for augmenting patient treatment is being hampered.

Potential Bias in Data

- Bias in data used to develop AI tools can reduce their effectiveness and accuracy.
- Addressing bias is difficult because electronic health data currently available do not represent the general population.

Difficulties in Scaling

- AI tools can be challenging to scale up and integrate into new settings because of differences among institutions and the patient populations they serve.

Limited Transparency of AI Tools

- Both interpretability and explainability pose challenges to explaining an AI tool's decision-making in an understandable way.
- This limited transparency can make it difficult or impossible for providers to understand how an AI tool came to a decision and whether and how an error occurred, as well as hampering the development of trust in the AI system.

Difficulties Protecting Patient Privacy

- As more AI systems are developed, large quantities of patient data will be in the hands of more people and organizations. This dispersion of data contributes to patient privacy risks.
- Patient advocacy groups and others have raised concerns, such as about the proliferation of potentially sensitive patient data, potentially without patient consent.

Uncertainty about Liability for AI Tools

- There is uncertainty about liability issues related to AI tools for augmenting patient treatment.
- The large number of people involved with developing and using AI tools as well as limited transparency of the tools contribute to this uncertainty.

3.1 Difficulties Accessing High-Quality Data

Accessing sufficient high-quality data to develop AI tools is a significant challenge—so much so that it can be considered one of the most important factors when deciding what tools to develop. Data are integral to all phases of AI tool development and deployment. Large quantities of high quality

data are needed to train, tune, evaluate, and validate AI models. For example, machine

learning models often require millions of pieces of information for training.³⁶

Many factors make it difficult for developers to access high quality data. Data are often siloed in different systems, such as medical imaging archival systems, diagnostic systems, EHRs, electronic prescribing tools, and insurance databases. The various siloes store data in varying formats that cannot easily be reconciled and aggregated, even if the data could be shared. The more systems involved, the more difficult it can be to aggregate the data. In addition, providers and developers may find it difficult to share their data. One reason is that providers may have to take additional steps to obtain patient authorization, or to qualify for an authorization exception, in order to comply with specific use and disclosure requirements in the Health Insurance Portability and Accountability Act of 1996 (HIPAA), and its implementing regulations; see text box. In addition, the commoditization of data sometimes creates disincentives to sharing. Some companies do not share data they have collected because they view it as giving them a competitive advantage over others in developing AI tools. In other cases, companies

have purchased or otherwise acquired large data sets and obtained exclusive rights to their use, potentially precluding others from using the data.³⁷ Furthermore, it can be time consuming to gain access to data. In the case of some federal health data, gaining access can take 12 to 18 months, according to experts.

³⁶Z. Obermeyer and E. Emanuel. "Predicting the Future — Big Data, Machine Learning, and Clinical Medicine." *New England Journal of Medicine*, September 29; 375(13): 1216–1219 (2016). Mark Sendak, Joshua D'Arcy, Sehj Kashyap, Michael Gao, Marshall Nichols, Kristin Corey, William Ratliff, Suresh Balu, "A Path for Translation of Machine Learning Products into Healthcare Delivery," *EMJ Innovations* (Jan. 2020).

³⁷Whether health care providers or other entities that hold patients' data have legal ownership of the data depends on state law and other factors. Analysis of this complex ownership question is beyond the scope of this Technology Assessment. Additionally, patients have access rights to their data held by HIPAA-covered entities. HIPAA's implementing regulations, known as the Privacy Rule, generally require covered entities, such as providers, to provide individuals, upon request, with access to PHI about them. 45 C.F.R. § 164.524. (The HIPAA access right does not however, apply to large stores of health-relevant data held by non-HIPAA regulated entities). With regard to providers' sale of patient data, HIPAA-covered providers may not sell PHI unless the affected patients grant a HIPAA authorization. 45 C.F.R. §§ 164.502(a)(5)(ii), 164.508(a)(4). The Privacy Rule does not restrict the use or disclosure of de-identified health information as it is no longer considered PHI.

Health Insurance Portability and Accountability Act of 1996 (HIPAA)

Health data may include individually identifiable health and genetic information that, in many instances, is protected by the Health Insurance Portability and Accountability Act of 1996 (HIPAA), as amended, and its implementing regulations, known as the Privacy Rule, as well as by other federal and state laws. The Privacy Rule, when it applies, governs the use and disclosure of individuals' health and genetic information and provides individuals with privacy rights regarding their health information.

The Privacy Rule applies to certain "HIPAA-covered entities"—which include health plans, health care clearinghouses, and health care providers. They are generally prohibited from using or disclosing protected health information without the individual's authorization, unless the use fits within one of the HIPAA Privacy Rule's specific authorization exceptions. These exceptions include, for example, disclosure of protected health information that has been de-identified; use of data for review preparatory to research or for a limited data set; disclosures of data to public health authorities; or disclosures for use in research if the covered entity obtains documentation that an institutional review or a privacy board has granted waiver of the authorization requirement

Source Health Insurance Portability and Accountability Act of 1996 (HIPAA), as amended, and its implementing regulations, known as the Privacy Rule. | GAO-21-75P

Another impediment to data access is that patient records do not always include all of the data developers need. EHRs generally are not a comprehensive source of a patient's entire health care record. A key reason is lack of interoperability.³⁸ When EHR systems are interoperable, information can be sent from one provider to another and then seamlessly

³⁸In previous work, GAO described five key challenges to achieving EHR interoperability: 1) insufficiencies in health data standards, 2) variation in state privacy rules, 3) accurately matching patients' health records, 4) costs associated with interoperability, and 5) the need for governance and trust among entities, such as agreements to facilitate the sharing of information among all participants in an initiative. GAO, *Electronic Health Records: Nonfederal Efforts to Help Achieve Health Information Interoperability*, GAO-15-817 (Washington, D.C.: Sept. 16, 2015).

integrated into the receiving provider's EHR system. However, EHRs are maintained at the level of health care institutions and providers and in practice, standards used by EHR systems can vary among providers and even among EHR systems used within the same hospital. This can lead to disjointed, incomplete records.³⁹

Yet another obstacle is that some data sets might not contain data of value in large enough quantities to train algorithms. In some cases this is because data are stored in fragmented systems and the processes to make the data interoperable may be too time- and resource-intensive to be worthwhile. In other cases, the data may not exist. For example, rare diseases occur in only a small fraction of the general population, resulting in a limited amount of rare disease data.

Innovation in AI tools is being hampered by such data access issues.⁴⁰ Furthermore, some developers may be disadvantaged compared to others. For example, developers within a company or in an academic setting may still have access to data through their own organizations. However, developers who are not part of a health care system may need to first partner with other organizations to gain

³⁹In response to requirements governing interoperability and information blocking in the 21st Century Cures Act, the HHS Office of the National Coordinator for Health IT promulgated regulations that, among other things, adopt interoperability standards for data access, exchange, and constituent data elements required to support interoperability. 21st Century Cures Act: Interoperability, Information Blocking, and the ONC Health IT Certification Program. 85 Fed. Reg. 25,642 (May 1, 2020).

⁴⁰W. Nicholson Price II, "Big Data, Patents, and the Future of Medicine", *Cardozo Law Review*, vol. 37, no. 4 (2016). Mark A. Hall, "Property, Privacy, and the Pursuit of Interconnected Electronic Medical Records", *Iowa Law Review*, vol. 95, no. 2 (2010).

access to certain data sets. In addition, high-resource hospitals are more likely to have greater volumes of data and higher quality data compared to smaller, lower-resource hospitals.⁴¹ This may exacerbate disparities among such locations, as tools can be more readily developed for high-resource settings.

3.2 Potential Bias in Data

Bias in data used to develop AI tools can reduce their safety and effectiveness for patients who differ—whether genetically or in socioeconomic status, general health status, or other characteristics—from the population whose data were used to develop the tool.⁴² For example, if an AI tool is developed using data from a high-resource hospital and later applied to a community-based hospital with a different patient population, the tool may not necessarily be safe for the new population or perform as effectively in the new location. Bias can be introduced in a number of ways, including when 1) the data used come only from certain populations and do not represent all the populations for which the tool would be used, 2) certain subgroups of patients are not represented in sufficient numbers, and 3) documentation or clinical reasoning is less accurate or systematically different across sites.

Addressing bias can be difficult because AI relies on data generated by humans or collected by systems created by humans. Biases can enter AI systems and as a result, AI tools may reproduce or increase existing bias. (See text box on following page for a potential consequence.) In particular, studies have shown that electronic health data do not currently represent the general population. High-resource hospitals often were early adopters of EHR systems and thus may have larger volumes of high-quality electronic data that now can be used to develop and train AI tools.⁴³ However, the data from such hospitals may underrepresent some patient populations. Similarly, genomic data may be helpful in developing AI tools, but they may also underrepresent populations. For example, data from individuals of African and Latin American ancestry, Hispanic people, and native or indigenous peoples represented less than 4 percent of samples collected in genome-wide association studies as of 2016.⁴⁴

⁴¹In this report, we use the term “high-resource hospital” to refer to an academic medical center or state-of-the-art hospital or hospital system.

⁴²W. Nicholson Price II, “Medical AI and Contextual Bias,” *Harvard Journal of Law and Technology*, vol. 33, no. 1 (2019).

⁴³W. Nicholson Price II, “Medical AI and Contextual Bias” and Vindell Washington, Karen DeSalvo, Farzad Mostashari, and David Blumenthal, “The HITECH Era and the Path Forward,” *New England Journal of Medicine*, vol. 377, no. 10 (2017), and John D. Halamka, Kenneth D. Mandl, and Paul C. Tang, “Early Experiences with Personal Health Records,” *Journal of the American Medical Informatics Association*, vol. 15, no. 1 (2008).

⁴⁴Alice B. Popejoy and Stephanie M. Fullerton, “Genomics is failing on diversity,” *Nature*, vol. 538, no. 7624 (2016).

Bias in a commercial AI risk prediction tool

A recent study uncovered racial bias in a commercial risk-prediction tool that large health systems and payers nationwide use to screen patients for “high-risk care management programs”—programs that aim to provide additional resources to patients with complex health needs. The tool’s algorithm used a patient’s health care costs as a proxy for their health. Although these two variables are correlated—sicker patients generally need more care—disparities in health care access, among other factors, mean that on average African American patients have lower medical costs at the same level of health as Caucasian patients. As a result, the Caucasian patients who met the criteria for referral to these high-risk care programs were significantly healthier than the African American patients, and fixing the disparity could increase the percentage of African American patients receiving additional care. Correcting such biases requires expert knowledge of the relevant field, the ability to identify and extract relevant data elements, and the capacity to iterate and experiment. The authors of this study are currently collaborating with the algorithm developer to incorporate better predictors of health into future versions.

Source: GAO summary of Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, “Dissecting racial bias in an algorithm used to manage the health of populations,” *Science*, vol. 366, no. 6464 (2019): p. 447. | GAO-21-7SP

3.3 Difficulties in Scaling

AI tools can be challenging to integrate into new settings (e.g., applying a tool developed at one hospital in another) and to scale up (e.g., taking a tool developed at one hospital and applying it across an entire health care system) because of differences among institutions and the patient populations they serve. Population differences can make it difficult to scale and integrate AI tools for the same reasons that they can introduce bias: tools developed with non-representative data may not be generalizable. Similarly, institutional differences can make scaling and integration difficult because AI tools developed in one setting, such as at a high-

resource hospital, may make recommendations that are inappropriate in another, such as a low resource hospital. For example, one article reported that an AI system developed at a high-resource setting (such as an academic medical center) that uses data from that location may recommend a powerful dose of specialized medicine because trained nurses and other specialists are available to intervene and monitor negative side effects.⁴⁵ The same recommendation at a low-resource setting, such as a community or rural hospital without such resources, might not work because the AI tool might not be able to take these constraints into account. Additionally, an AI tool may make recommendations for treatments using technologies that are unavailable in a low-resource hospital.

In some cases it may be feasible to integrate a tool into a new setting by adapting the model to the different environment. AI tools developed using EHR data at one location can sometimes be integrated into a new setting by retraining on data specific to the EHR infrastructure of new site, but this can be costly. For example, one review estimated that a sepsis prediction tool would cost \$700,000 to scale across three sites, whereas an end-of-life prediction tool would cost over \$30 million to integrate into more than 25 sites.⁴⁶ However, scaling a tool across a health care system can be even more difficult because there is more variation among settings. Examples of tools that have been integrated across the health care system are rare. In fact, one review stated that there are

⁴⁵W. Nicholson Price II, “Medical AI and Contextual Bias”.

⁴⁶Mark Sendak, Joshua D’Arcy, Sejh Kashyap, Michael Gao, Marshall Nichols, Kristin Corey, William Ratliff, Suresh Balu, “A Path for Translation of Machine Learning Products into Healthcare Delivery,” *EMJ Innovations*, (Jan. 2020).

no clinical decision support AI tools using EHR data that have successfully scaled across a health care system.⁴⁷ However, as was previously mentioned, some risk-prediction tools have been scaled more widely, such as the tool described in the previous text box, although in that case with mixed success due to bias issues.

3.4 Limited Transparency of AI Tools

The limited transparency of AI tools used in health care—referred to by some experts as “black box medicine”—makes it difficult for providers, regulators, and others to determine whether an AI tool is safe and effective.⁴⁸ For one thing, the decision-making of AI algorithms can be difficult or impossible to explain; in fact, with some complex machine learning algorithms, even their developers cannot determine how they work. Furthermore, some machine learning algorithms continuously learn and adapt, so how they work may change over time.

In addition, AI developers do not always share information so that it can be subjected to broader scrutiny. Scrutiny by other experts is critical to ensuring the accuracy and integrity of scientific research. Additionally, there can be limited information available to help ensure the validity, safety, and effectiveness of some AI tools. For example:

- As is common in the machine learning community more broadly, many studies about AI tools for patient care have only been published online without being

submitted to peer-reviewed journals, and reported results may not include assessments of data quality. In addition, we found that some tools are not assessed in clinical settings.

- One review article expressed concern that even though studies often document the accuracy of AI tools, accuracy does not necessarily represent clinical efficacy.⁴⁹ That same article stated that in order for a provider to assess the relative effectiveness of two tools for use on their patient population, both tools would need to be subjected to comparison on the same independent test set that is representative of the target population, but such comparisons are generally not done.
- Developers may not share the details of their algorithms, including how they were developed and validated, in order to protect their intellectual property from competitors.

Providers are not routinely receiving training about AI systems that could better enable them to critically assess AI tools in order to use them safely and to be able to explain them to patients. Some experts also expressed concern that providers are not being adequately involved in the development of AI tools, which would give them a better understanding of the tools and have the additional benefit of helping to ensure the tools best fit the needs of providers.

⁴⁷Sendak et al 2020.

⁴⁸W. Nicholson Price II, “Regulating Black Box Medicine,” *Michigan Law Review*, vol 116, no. 3 (2017).

⁴⁹Kelly, C.J., Karthikesalingam, A., Suleyman, M. et al. “Key challenges for delivering clinical impact with artificial intelligence,” *BMC Medicine*, vol. 17, no. 195 (2019).

The limited transparency of AI tools has many consequences. It can make it difficult or impossible for providers to understand how an AI tool came to a decision and whether and how an error occurred. It is also difficult to evaluate or otherwise determine the quality of algorithms, including their effectiveness. In addition, limited information on data sets used to develop tools could make it challenging for providers to determine whether the tool was trained on data that were representative of a particular population, as well as what the accuracy and error rates are for different patient subgroups.

Finally, limited transparency also affects trust in an AI tool because providers may not adequately understand the tool. Having an appropriate level of trust in the tool is an important aspect of determining whether to take action based on a prediction, or even whether to deploy a new tool for use within a provider's office. A lack of trust can lead providers to under-rely on a tool, such as by ignoring its recommendations or not using it. Under-reliance on a tool precludes the provider from making full use of its capabilities and the patient from reaping its benefits. In contrast, providers may over-rely on a tool. For example, automation bias occurs when users rely on recommendations given by an AI tool even when they know or should know the tool is wrong.⁵⁰ Overreliance on the AI tool eliminates a human-based intervention that acts as a check on a tool's potential errors.

⁵⁰W. Nicholson Price II, "Medical AI and Contextual Bias".

⁵¹National Institute of Standards and Technology, Four Principles of Explainable Artificial Intelligence, Draft NISTIR 8312 (Gaithersburg, MD: Aug. 2020).

Two experts and one agency expressed concerns related to this challenge. First, the National Institute of Standards and Technology described a lack of standards for what constitutes a sufficient explanation of an AI tool's performance.⁵¹ This could lead to uncertainty about what is required and allows developers to share minimal information about these already difficult to understand systems. In the absence of sufficient publicly available information, regulatory oversight could help assure providers that tools are safe and effective. However, two experts expressed concern that it is unclear what tools the Food and Drug Administration (FDA) will regulate, with one expert stating that FDA may only oversee a fraction of AI tools. One expert also stated that they did not see FDA positioning itself to play an active role in premarket clearance or approval processes for these technologies. In 2019, FDA issued a proposed regulatory framework for machine learning-based software as a medical device (SaMD), but it has not yet promulgated any regulations.⁵² Without more transparency or approval from an oversight body, it may be difficult for providers to determine whether a tool is safe and effective and may hinder widespread adoption of AI tools.

3.5 Difficulty Protecting Patient Privacy

As more AI systems are developed and deployed in clinical settings, large quantities of patient data will be in the hands of more people and organizations, which can

⁵²FDA, "Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD)," Washington, D.C.: Apr. 2, 2019.

contribute to patient privacy risks. Health records are rich in valuable and sensitive information, such as social security numbers, addresses, mental health information, communicable disease diagnoses, and credit card information.

According to a 2018 report, hackers of all types (e.g., nation-state actors, cyber criminals, hacktivists) have found numerous ways to make money from illegally obtained health care data.⁵³ Examples include selling data on the black market to facilitate Medicare fraud and identity theft, and gathering foreign intelligence. Health information technology includes connected, networked systems and leverages wireless technologies, leaving such systems more vulnerable to cyber attack. Although the HIPAA Security Rule requires appropriate administrative, physical, and technical safeguards to ensure the confidentiality, integrity, and security of electronic protected

health information (PHI), studies indicate that the number of data breaches in the health industry—including from cyber attack and from insiders intentionally or unintentionally leaking information—may be among the highest of any sector.⁵⁴ Multiple reports have shown problems related to understanding of some basic cybersecurity measures among health care organization staff.

To help protect patient privacy and reduce potential access to personally identifiable information while allowing secondary use of data, the HIPAA Privacy Rule allows providers to share patient data without authorization as long as they first “de-identify” the data—remove patients’ PHI.⁵⁵ Because about half of the U.S. population can be identified with the combination of their date of birth, sex, and five digit zip code, removing PHI is an essential step to protecting patient data.⁵⁶ The Privacy Rule’s most commonly used de-identification method requires removal of

⁵³Department of Health and Human Services, “Health Industry Cybersecurity Practices: Managing Threats and Protecting Patients (HICP),” (Dec. 28, 2018).

⁵⁴The HIPAA Security Rule establishes national standards to protect individuals’ electronic PHI that is created, received, used, or maintained by a covered entity. 45 C.F.R. pt. 160 and pt. 164, subpts. A and C.

⁵⁵The Privacy Rule was designed to protect individually identifiable health information by permitting only certain uses and disclosures of PHI provided by the rule, or as authorized by the individual subject of the information. However, in recognition of the potential utility of health information even when it is not individually identifiable, see 45 C.F.R. §164.502(d), the Privacy Rule permits a covered entity or its business associate to create information that is not individually identifiable by following the de-identification standard and implementation specifications in 45 C.F.R. §164.514(a)-(b). These provisions allow the entity to use and disclose information that neither identifies nor provides a reasonable basis to identify an individual and includes two methods: 1) a formal determination by a qualified expert; or 2) the removal of specified individual identifiers as well as absence of actual knowledge by the covered entity that the remaining information could be used alone or in combination with other information to identify the individual.

⁵⁶Department of Health and Human Services, Guidance Regarding Methods for De-Identification of Protected Health Information in Accordance with the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule, accessed May 27, 2020, <https://www.hhs.gov/hipaa/for-professionals/privacy/special-topics/de-identification/index.html#standard>.

various data fields, such as names, social security numbers, account numbers, and full faced photographs.⁵⁷ This de-identification offers some protection to patients—if their data are hacked or leaked, for example, it should reduce the likelihood that recipients will be able to identify patients’ PHI. However, governmental advisory bodies acknowledge that even properly de-identified data retain some risk of re-identification. Researchers have shown that it is possible to re-identify patients using outside sources of information, although risks of re-identification vary.⁵⁸ For example, one study re-identified patients in two states by cross-referencing anonymized patient records and local newspaper articles.⁵⁹ Furthermore, there are concerns that increasingly sophisticated computer techniques may make it easier to re-identify data.

Patient advocacy groups and others have raised additional concerns about how potentially sensitive patient data are used. Some of these concerns pertain to companies and providers profiting from data that providers gathered from patients as well as a lack of transparency for patients that such uses are occurring.⁶⁰ For example, providers are forming business arrangements with technology companies and developers whereby they transfer their patients’ de-identified data for technology development

purposes. In one example, a provider shared its de-identified patient data with an AI developer in exchange for partial ownership of the developer. In another case, an artificial intelligence company obtained a large dataset and publicly described a very narrow intended use of the data, but was contractually allowed to use the data in many more ways than were made apparent. Another concern that a patient advocacy group identified is the potential for companies to unaccountably use re-identified data, such as in a discriminatory manner. There is a long history of concern about health data being used by employers to discriminate, based, for example, on mental health status or diseases such as human immunodeficiency virus (HIV). Furthermore, even when there is patient authorization, two experts told us that it may be obtained through blanket consent forms that patients may not read. As a result, even with authorization, patients’ data may be used in ways that are not clear to the patient. Because of patient concerns about how their data may be used, patients may ultimately become less willing to communicate important medical information to their providers, which could lead to a decline in care for patients who do not share critical medical information.

⁵⁷ 15 C.F.R. § 164.514(b) and (c).

⁵⁸ Sources for re-identification information can include publicly available information, data extraction, and DNA sequences. Data extraction is a collection technique where software is used to search the internet for information about individuals, and extract information from websites that contain consumer information. GAO, *Information Resellers: Consumer Privacy Framework Needs to Reflect Changes in Technology and the Marketplace*, GAO-13-663, (Washington, D.C.: Sept. 25, 2013).

⁵⁹ Ji Su Yoo, Alexandra Thaler, Latanya Sweeney, Jinyan Zang, “Risks to Patient Privacy: A Re-identification of Patients in Maine and Vermont Statewide Hospital Data”, *Technology Science*, 2018100901. October 09, 2018

⁶⁰ 45 C.F.R. § 164.528 provides individuals with a right to receive an accounting of disclosures of PHI made by a covered entity in the six years prior to the date on which the accounting is requested, subject to certain exceptions. Health information that is de-identified in accordance with 45 C.F.R. § 164.514 is not PHI.

3.6 Uncertainty about Liability for AI Tools

Multiple experts told us that there is uncertainty about liability issues related to AI tools for augmenting patient care. There are many reasons that contribute to this. For example:

- Because many parties are involved in developing and using these tools, it may be difficult to determine who is responsible—the developer, provider, or someone else—if a problem arises.
- According to some legal experts, AI is new to clinical practice, and there is no case law on liability involving medical AI. As a result, it is unclear what types of claims courts will recognize. The absence of

litigated cases also creates uncertainty about the legal doctrines courts may apply when resolving such claims. Furthermore, FDA’s regulatory treatment of these technologies is still evolving.

- Because of limited understanding of how some AI tools perform, it may be difficult and, in some cases potentially not possible, to identify the cause of a medical error. If providers are concerned that using a tool could leave them liable for its errors or shortcomings, this uncertainty could slow the adoption of AI tools. According to one expert, this may stifle innovation because developers would take a more risk averse approach to product development.

4 Policy Options to Enhance Benefits or Address Challenges of AI Tools to Augment Patient Care

We developed six policy options that policymakers—Congress, elected officials, federal agencies, state and local governments, academic research institutions, and industry, among others—could take to enhance the benefits of AI technologies to augment care in health care facilities or to respond to the challenges discussed in the previous chapter. The first five options would aim to do so by encouraging collaboration, improving data sharing, encouraging the development of best practices, fostering interdisciplinary education, and clarifying oversight mechanisms. The sixth option is to maintain the status quo. For each option, we analyze potential opportunities and considerations.

While we present options to address the major challenges we identified, the options are not intended to be exhaustive. We intend policy options to provide policymakers with a broader base of information for decision-making. The options are neither recommendations to federal agencies nor matters for congressional consideration. We did not rank the options in any way. We are not suggesting that they be done individually or combined in any particular fashion. Additionally, depending on the options selected, additional steps might need to be taken on potential design and legal issues. We did not conduct work to assess how effective the options may be, and express no view regarding the extent to which legal changes would be needed to implement them.

Policy Option: Collaboration

Policymakers could encourage interdisciplinary collaboration between developers and health care providers.

Potential Opportunities

- Early and consistent collaboration could help developers design AI tools that are easier to implement and use within providers' existing workflow and associated constraints. According to one provider organization, providers are only seen as the end-user of the product. However, they can also contribute to product design because they have useful information on how the products may affect their workflow and the patient experience, as well as insight on how to best design the tools to be easily implementable.

Such collaboration could allow the technology to be implemented on a larger scale. For example, front-line providers are more likely to adopt AI tools that they help to create, because they have the institutional knowledge to build and refine tools that affect their daily work. Collaboration prior to deployment could ensure that tools are usable for providers. Additionally, one researcher we interviewed suggested that a major barrier to scalable implementation of AI is lack of integration into existing workflow processes and the current infrastructure of health systems. If developers and providers collaborate to design products that complement provider workflow, health systems may not need to expend resources to create workarounds to enable successful integration. For example, one company fostered collaboration by having health care providers and developers work side-by-side to create a tool designed to integrate smoothly into nurses' workflow and existing call bell system.

- Policymakers could use innovative approaches to encourage such collaboration. According to one company representative, an approach that could promote collaboration is to hold hackathons where computer engineers, other technology experts, and providers collaborate to find solutions that use AI technology. As we previously reported, another innovative approach is a challenge format, in which organizations create a challenge to encourage the public to find solutions.⁶¹ For example, the Centers for Medicare & Medicaid Services (CMS) launched an AI Health Outcomes Challenge in 2019, seeking input from innovators inside and outside the health care sector to harness AI solutions for predicting health outcomes for potential use at the agency.⁶²

Potential Considerations

- Collaboration may create tools that are specific to one hospital or provider, further contributing to scalability challenges we identified earlier. For example, AI tools that were designed around site-specific variations, such as size and patient population, may enable the AI tool to work more effectively in one location than another.
- Health care providers may not have time to both collaborate with developers and treat patients. However, organizations can provide protected time for employees to engage in innovation activities such as collaboration, according to a guide developed by providers.⁶³

Source: GAO. | GAO-21-75P

⁶¹GAO, *Open Innovation: Practices to Engage Citizens and Effectively Implement Federal Initiatives*, GAO-17-14 (Washington, D.C.: Oct. 13, 2016).

⁶²Centers for Medicare & Medicaid Services, *Artificial Intelligence (AI) Health Outcomes Challenge*, accessed May 12, 2020, <https://innovation.cms.gov/innovation-models/artificial-intelligence-health-outcomes-challenge>.

⁶³Rosina Cianelli, Bonnie Clipper, Rebecca Freeman, Jill Goldstein, Tami H, Wyatt, *The Innovation Road Map. A Guide for Nurse Leaders*, accessed May 18, 2020, <https://www.nursingworld.org/globalassets/ana/innovations-roadmap-english.pdf>.

Policy Option: Data Access

Policymakers could develop or expand high-quality data access mechanisms.

Potential Opportunities

- Increasing the availability of high-quality data could facilitate the development and testing of AI tools. One expert told us that increased data access could expand the number of developers and researchers who create and test tools. As stated in chapter 3, there are constraints related to collecting high-quality, representative data for developing AI systems. Increased data access could eliminate some of these barriers and allow developers to create AI projects that are developed and tested on more representative datasets.
- Policymakers could consider increasing data access by creating a type of mechanism known as a data commons—a cloud-based platform where users can store, share, access, and interact with data and other digital objects. For example, the Stanford Institute for Human-Centered Artificial Intelligence proposed a National Research Cloud, which would be a partnership between academia, government, and industry to provide access to resources, potentially including a large-scale, government-held data set in a secure cloud environment to develop and train AI.

Additionally, a 2019 executive order asked federal agencies to identify opportunities to increase access to and use of federal data and models.⁶⁴ In response to the executive order as well as other internal plans, the Department of Health and Human Services (HHS) released a plan to change the way the department internally shares and leverages data across HHS agencies using an interagency data hub.⁶⁵ The hub is intended to give HHS staff quicker and easier access to data from other HHS agencies.

- Increasing access to high-quality data could help developers address bias concerns by ensuring data are representative, transparent, and equitable. A common platform would allow people to test and validate their algorithms across multiple health systems or data sets. The replication of outputs in multiple situations could prevent the introduction of bias into the algorithm as it is being tested and validated. Enhanced data sharing can also mitigate bias by ensuring open access to the data so developers and providers can assess how the AI was trained and tested.

Potential Considerations

- Agencies and other stakeholders would likely need to expend resources to successfully coordinate across their respective domains. One expert expressed concern that a large amount of existing data may not be useful because the data are not interoperable across domains. Agencies and stakeholders that collect and store data could take steps to make the data interoperable but this would likely require a large component of skilled labor and investment, according to the same expert. For example, we previously described how the Department of Veteran's Affairs and the Department of Defense abandoned a joint EHR integration project because of concerns about the program's cost, schedule, and inability to meet deadlines, and replaced the project with a goal of developing separate systems that could still be made interoperable.⁶⁶ The interagency program office reported spending about \$564 million on the project between 2011 and 2013.
- Health care providers that own patient data often consider those data proprietary and valuable, and often sell such data to other firms. These companies may be reluctant to share such information freely with their competitors.
- Cybersecurity and privacy risks could increase as more data are accessed and shared. If PHI is accessed inappropriately, patients can face discrimination or financial threats, according to a think tank report.⁶⁷ According to a Department of Defense Office of the Inspector General report, health care breaches of PHI affected over 46 million patients from April 2018 through April 2020 and were caused by cyber attacks, theft, improper disposal of data, and unauthorized access.⁶⁸ To try to avoid risks associated with the use and disclosure of PHI, as an alternative, de-identified data may be shared in some circumstances. For example, the National Institutes of Health's *All of Us* Research Program aims to gather detailed health information from over one million Americans while taking precautions such as removing identifying information from data before its use, enacting strict internal policies and procedures to prevent misuse of data, and requiring researchers to agree not to re-identify data.

Source: GAO. | GAO-21-7SP

⁶⁴Maintaining American Leadership in Artificial Intelligence, Exec. Order No. 13859, § 5a, 84 Fed. Reg. 3967, 3969 (Feb. 14, 2019).

⁶⁵United States Department of Health and Human Services, "Leveraging Data for the Nation's Health," December 2019.

⁶⁶GAO, *Electronic Health Records: VA and DOD Need to Establish Goals and Metrics for Their Interoperability Efforts*, GAO-16-184T (Washington, D.C.: Oct. 27, 2015).

⁶⁷With regard to providers' sale of patient data, HIPAA-covered providers may not sell PHI unless the affected patients grant a HIPAA authorization. 45 C.F.R. §§ 164.502(a)(5)(ii), 164.508(a)(4).

⁶⁸U.S. Department of Defense, Office of the Inspector General, *Special Report on Protecting Patient Health Information During the COVID-19 Pandemic*, DODIG-2020-080 (Alexandria, VA: April 2020).

Policy Option: Best Practices

Policymakers could encourage relevant stakeholders and experts to establish best practices (such as standards) for development, implementation, and use of AI technologies.

Potential Opportunities

- Best practices can help providers deploy AI tools in hospitals and health care systems by providing guidance on a number of issues, such as those relating to data, interoperability, bias, and implementation. For example, best practices could help hospitals identify the appropriate amount of testing required to help ensure it is safe to implement an AI tool in their hospital, as well as guide them through gathering data on the AI tool's performance to ensure it is working as intended.
- Best practices could improve scalability of AI tools by enhancing interoperability. As discussed above, scaling AI tools can be difficult because of challenges related to retraining models using differently formatted data, among other things. Standards such as those identified by the Interoperability Standards Advisory (ISA) may be able to help address this challenge. As we previously reported, the Office of the National Coordinator for Health Information Technology (ONC) uses the ISA process to coordinate interoperability standards and implementation specifications for clinical information technology needs.⁶⁹ These standards could help hospitals improve interoperability of their data, which could support data gathering efforts to develop AI tools.
- Best practices may also help reduce bias in AI tools. As discussed above, data used to train AI algorithms may underrepresent certain groups. Best practices for representation in training data sets may help mitigate the introduction of bias into algorithms and may help improve the performance of the AI technology for a wider demographic base. Similarly, a standard that requires or encourages transparency in data sets can help reduce bias by allowing researchers to assess sources of data, characteristics within the data, and other information that may put algorithms at risk of bias.
- Standing working groups or committees could identify the areas in which best practices would be most beneficial, develop, and periodically update best practices to help ensure they remain current and relevant. Meetings could occur with representatives from academia, patient and physician advocacy groups, industry, and the federal government, among other entities.

Potential Considerations

- The creation of best practices, such as standards, could require consensus from many public- and private-sector stakeholders, which can be time-intensive as well as a significant resource commitment. We previously reported that development of standards can take anywhere from 18 months to a decade to complete and require multiple iterations.⁷⁰
- Some best practices may not be widely applicable because institutions and patient populations differ across locations. For example, we previously reported that rural hospitals serve patients who have distinct characteristics, such as a higher percentage of residents who are elderly, have limitations in activities caused by chronic conditions, and have lower median household income.⁷¹ The resources required to manage patients in these settings may differ from those required in non-rural hospitals.

Source: GAO. | GAO-21-75P

⁶⁹GAO, *Health Information Technology: Approaches and Challenges to Electronically Matching Patients' Records across Providers*, GAO-19-197 (Washington, D.C.: Jan. 15, 2019).

⁷⁰GAO, *National Institute of Standards and Technology: Additional Review and Coordination Could Help Meet Measurement Service Needs and Strengthen Standards Activities*, GAO-18-445 (Washington, D.C.: July 26, 2018).

⁷¹GAO, *Rural Hospital Closures: Number and Characteristics of Affected Hospitals and Contributing Factors*, GAO-18-634 (Washington, D.C.: Aug. 29, 2018).

Policy Option: Interdisciplinary Education

Policymakers could create opportunities for more workers to develop interdisciplinary skills.

Potential Opportunities	Potential Considerations
<ul style="list-style-type: none">• Interdisciplinary education could help providers use tools effectively. For example, a 2019 National Academy of Medicine (NAM) report described the importance of expanding training and educational programs in AI and health care.⁷² According to NAM, education that engages health care providers, AI developers, implementers, and health care system leaders could help them benefit from and sustain the use of AI tools.• Policymakers could implement this option in a variety of ways, including changing academic curriculum or through research grants. For example, the National Institutes of Health provides funding for mentored research and career development support for certain students.⁷³ Further, the VA administers a Big Data-Scientist Training Enhancement Program, which provides opportunities for postdoctoral fellows and other researchers to gain experience working in one of the VA Medical Centers on AI/big data initiatives.⁷⁴ Policymakers could consider more efforts that help provide researchers with interdisciplinary skills and experience.• Providers might more critically evaluate recommendations from AI tools and therefore make better decisions if they better understand how such technologies work. According to one study, interdisciplinary education might help providers more effectively supervise AI tools and recognize cases where algorithms are not working as intended.⁷⁵	<ul style="list-style-type: none">• Employers and university leaders may have to modify their existing curriculum. One researcher we spoke with discussed how adding AI-related coursework without removing other requirements may increase the length of medical training. Medical training is already lengthy in some cases—in a prior report we described how medical students must complete an undergraduate degree and typically 4 years of medical school, and may require an additional 3-5 years of residency training, depending on the specialty.⁷⁶ However, another expert described how creating new career tracks that integrate AI into medical training could mitigate this concern. This would allow providers to complete AI training as part of normal coursework.

Source: GAO. | GAO-21-75P

⁷²National Academy of Medicine, *Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril* (Washington, D.C.: National Academy of Sciences, 2019).

⁷³National Institutes of Health, Office of Data Science Strategy, Stimulating Access to Research in Residency Transition Scholar (StARRTS), accessed September 15, 2020, <https://datascience.nih.gov/foa/stimulating-access-research-residency-transition-scholar-starrts-k38-clinical-trial-not-allowed>.

⁷⁴United States Department of Veterans Affairs, Office of Academic Affiliations, Big Data-Scientist Enhancement Program (BD-STEP), accessed September 15, 2020. https://www.va.gov/oaa/specialfellows/programs/sf_bdstep.asp.

⁷⁵Ketan Paranjape, Michiel Schinkel, Rishi Nannan Panday, Josip Car, Prabath Nanayakkara, “Introducing Artificial Intelligence Training in Medical Education,” *JMIR Med Educ*, vol. 5, no. 2 (2019).

⁷⁶GAO, *Graduate Medical Education: Trends in Training and Student Debt*, GAO-09-438R (Washington, D.C.: May 4, 2009).

Policy Option: Oversight Clarity

Policymakers could collaborate with relevant stakeholders to clarify appropriate oversight mechanisms.

Potential Opportunities	Potential Considerations
<ul style="list-style-type: none">Regular, predictable oversight of AI tools could help ensure that AI tools remain safe and effective after deployment in clinical settings and throughout their lifecycle, in the case of machine learning tools that evolve over time. FDA’s 2019 discussion paper on a proposed regulatory framework for AI/machine-learning-based SaMD describes a regulatory framework that aims to promote a mechanism for tool manufacturers to be continually vigilant in maintaining the safety and effectiveness of their SaMD.⁷⁷ It includes a total product lifecycle regulatory approach, which could allow devices to continually improve after they are in use, while also providing safeguards. However, the proposed framework is in draft form and is not intended to serve as guidance or regulations for developers of SaMD. Finalized guidance or regulations from FDA, including details and examples of what types of AI tools they cover, could improve predictability, so that developers better understand what will be required of them. In addition, knowing whether AI tools will be regulated as medical devices may provide insights into questions of liability.A forum consisting of relevant stakeholders could help recommend additional mechanisms to ensure appropriate oversight of AI tools. Such efforts could focus on tools that FDA does not oversee or provide additional support to FDA for those that it does oversee. For example, the International Medical Device Regulators Forum is a collaborative effort across governments to clarify appropriate oversight mechanisms. The Forum is comprised of a group of medical device regulators from around the world who voluntarily come together to harmonize the regulatory requirements for medical products that vary from country to country. A similar forum for AI tools to augment patient care could meet to pursue clarity on issues such as liability, evaluation, and others.	<ul style="list-style-type: none">Soliciting input and coordinating among stakeholders, such as hospitals, professional organizations, and agencies, may be challenging. For example, interagency groups face challenges defining outcomes, measuring performance, and establishing leadership approaches. These challenges for collaborative groups may slow the pace of oversight or lead to oversight uncertainty. For example, we previously reported on how oversight of certain federal laws applicable to private health coverage was split between states, HHS, the Department of Labor, and the Department of the Treasury.⁷⁸ This required a multitude of coordination efforts, including: formal agreements between states and the federal government, informal ongoing communication, technical assistance and outreach, and grant funding. The states also instituted surveys and the agencies held a listening session to fine tune their efforts. While these coordination efforts were useful, we discussed how stakeholders told us that more guidance and clarity would also be helpful.Experts we spoke with indicated that excess regulation could slow the pace of innovation. A 2020 draft Office of Management and Budget memorandum encourages federal agencies to avoid regulatory or non-regulatory actions that needlessly hamper AI innovation and growth.⁷⁹

Source: GAO. | GAO-21-75P

⁷⁷FDA, “Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD),” Washington, D.C.: Apr. 2, 2019.

⁷⁸GAO, Mental Health and Substance Use: State and Federal Oversight of Compliance with Parity Requirements Varies, [GAO-20-150](#) (Washington, D.C.: Dec. 13, 2019).

⁷⁹Office of Management and Budget, Guidance for Regulation of Artificial Intelligence Applications, OMB Draft Memorandum (Washington, D.C.); Request for Comments on a Draft Memorandum to the Heads of Executive Departments and Agencies, “Guidance for Regulation of Artificial Intelligence Applications”, 85 Fed. Reg. 1825 (Jan. 13, 2020).

Policy Option: Status Quo

Policymakers could maintain the status quo (i.e., allow current efforts to proceed without intervention).

Potential Opportunities

- Challenges surrounding the use of AI to augment patient care that we identified earlier in our report may be resolved through current efforts.
- As we discuss earlier in this report, some hospitals and providers are already using AI to augment patient care and may not need policy-based solutions to continue expanding the use of such technologies. For example, Duke Health has developed and is using an AI tool that helps identify patients in the early stages of sepsis.
- Existing efforts may prove more beneficial than new options. For example, some stakeholders are engaging in knowledge sharing with others in the field already, which may address some issues associated with deployment of AI tools in patient care. For example, the National Institutes of Health, the National Center for Advancing Translational Sciences, along with the National Institute of Biomedical Imaging and Bioengineering and the National Cancer Institute, held a meeting with representation from industry, academia, and federal agencies to discuss issues related to data, bias, and transparency for AI in health care. A subsequent white paper from the meeting proposed solutions for issues associated with machine learning in health care.⁸⁰

Potential Considerations

- The challenges described earlier in the report may remain unresolved or be exacerbated. For example, fewer AI tools may be implemented at scale. According to one researcher we spoke with, under the status quo there will be slow and limited progress.
- Disparities in the use of AI tools because of scalability challenges may be exacerbated. High resource health systems may be more likely to overcome implementation barriers than their low resource counterparts.

Source: GAO. | GAO-21-75P

⁸⁰Christine M. Cutillo, Karlie R. Sharma, Luca Foschini, Shinjini Kundu, Maxine Mackintosh, and Kenneth D. Mandl, "Machine intelligence in healthcare- perspectives on trustworthiness, explainability, usability, and transparency," *npj Digital Medicine*, vol. 3, no. 47 (2020).

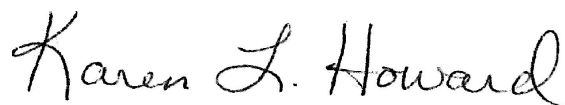
5 Agency and expert comments

We provided a draft of this report to the Department of Health and Human Services (Food and Drug Administration, National Institutes of Health, and the Centers for Medicare and Medicaid Services) and the Department of Veterans Affairs with a request for technical comments. We incorporated agency comments into this report as appropriate.

We also provided a draft of this report to 14 participants from our expert meeting, and incorporated comments received as appropriate.

We are sending copies of this report to the appropriate congressional committees, relevant federal agencies, and other interested parties. In addition, the report is available at no charge on the GAO website at <http://www.gao.gov>.

If you or your staff members have any questions about this report, please contact me at (202) 512-6888 or howardk@gao.gov. Contact points for our Offices of Congressional Relations and Public Affairs may be found on the last page of this report. GAO staff who made key contributions to this report are listed in appendix III.



Karen L. Howard, PhD
Director
Science, Technology Assessment, and Analytics

List of Requesters

The Honorable Lamar Alexander

Chairman
Committee on Health, Education, Labor, and Pensions
United States Senate

The Honorable Greg Walden


Ranking Member
Committee on Energy and Commerce
House of Representatives

The Honorable Michael C. Burgess, MD

Ranking Member
Subcommittee on Health
Committee on Energy and Commerce
House of Representatives

The Honorable Brett Guthrie

Ranking Member
Subcommittee on Oversight and Investigations
Committee on Energy and Commerce
House of Representatives



PART TWO

Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic

National Academy of Medicine

Part Two presents the NAM publication *Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic* discussing the use and challenges associated with AI technology in the delivery of health care services outside of settings where health care providers are employed, including the home. Although GAO and NAM staff consulted with and assisted each other throughout this work, reviews were conducted by GAO and NAM separately and independently, and authorship of the text of Part One and Part Two of the report lies solely with GAO and NAM, respectively.

Part Two—(NAM) Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic

Nakul Aggarwal, University of Wisconsin-Madison; **Mahnoor Ahmed**, National Academy of Medicine; **Sanjay Basu**, Harvard University; **John J. Curtin**, University of Wisconsin-Madison; **Barbara J. Evans**, University of Florida; **Michael E. Matheny**, Vanderbilt University Medical Center and Tennessee Valley Healthcare System VA; **Shantanu Nundy**, Accolade Inc.; **Carmel Shachar**, Harvard University; **Mark P. Sendak**, Duke University; **Rashmee U. Shah**, University of Utah; and **Sonoo Thadaney-Israni**, Stanford University

November 30, 2020

The health care ecosystem is witnessing a surge of artificial intelligence (AI)-driven technologies and products that can potentially augment care delivery outside of hospital and clinic settings. These tools can be used to conduct remote monitoring, support telehealth visits, or target high-risk populations for more intensive health care interventions. With much of patients' time spent outside of a hospital or a provider's office, these tools can offer invaluable benefits in facilitating patients' access to their provider teams in convenient ways, facilitating providers' understanding of their patients' daily habits, extending care access to underserved communities, and delivering personalized, real-time care in the patient's home environment. More importantly, by expanding care to novel settings (e.g., home, office), these technologies could empower patients and caregivers, as most of these tools are aimed at helping patients adapt their own behaviors or facilitating bidirectional communication between patients and clinicians for more personalized care. The authors of this manuscript refer to these such environments as "health settings outside the hospital and clinic," abbreviated and referred to as HSOHC (pronounced "h-


sock") hereafter (see Figure 1). In some instances, the capabilities of these tools are proving to be extremely timely in continuing care delivery amidst the disruptions posed by the COVID-19 pandemic.

While a number of AI applications for care delivery outside of the hospital and clinical setting in medical specialties ranging from cardiology to psychiatry are either currently available or in development, their reliability and true utility in improving patient outcomes are highly variable. In addition, fundamental logistical issues exist, including product scalability, inter-system data standardization and integration, patient and provider usability and adoption, and insurance reform that must be overcome prior to effective implementation of AI technologies. Broader adoption of AI in health care and long-term data collection must also contend with urgent ethical and equity challenges, including patient privacy, exacerbation of existing inequities and bias, and fair access, particularly in the context of the U.S.'s fragmented mix of private and public health insurance programs.

Figure 1: Artificial Intelligence in Health Settings Outside the Hospital and Clinic

AI in HSOHC

Artificial Intelligence in Health Settings Outside the Hospital and Clinic



Note: Represented in the orange third are the typical hospital and clinic settings. Represented in the blue two-thirds are the settings in which most health-related events and human experiences unfold, including the home, work, and community environments. Health-relevant data captured in these settings, for example via smartphone and wearable technology, can inform personalized and timely interventions, as well as public and environmental health assessments.

Introduction and Scope

To address the U.S. health care system's deep-seated financial and quality issues [1], several key stakeholders, including health systems, retail businesses, and technology firms, are taking steps to transform the current landscape of health care delivery. Notable among these efforts is the expansion of health care services outside the hospital and clinic settings [2,3]. These novel settings, or HSOHC, and modes of

care delivery include telehealth, retail clinics, and home and office environments. Care delivered in these environments often incorporates advanced technological applications such as wearable technology (e.g., smartwatches), remote monitoring tools, and virtual assistants.

The growing adoption of these technologies in the past decade [4] presents an

Figure 2: Advancing the Quintuple Aim



Source: Matheny, M., S. Thadaneysrani, 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.

opportunity for a paradigm shift in U.S. health care toward more precise, economical, integrated, and equitable care delivery. Coupled with advances in AI, the potential impact of such technologies expands exponentially (see Box 1 for key definitions). Machine learning (ML), a subdomain of AI, can take advantage of continuous data regarding activity patterns, peripheral physiology, and ecological momentary assessments of mood and emotion (all gathered in the home, school, community, and office settings) to predict risk for future health events and behavioral tendencies, and ultimately suggest personalized lifestyle modifications and treatment options. The increasing affordability of remote monitoring devices, decreased dependence on brick-and-mortar health care infrastructure, and real-time feedback mechanisms of these tools position AI as an indispensable factor in achieving the Quintuple Aim of health care:

better patient outcomes, better population health, lower costs, increased clinician well-being, and prioritized health equity and inclusiveness [5] (see Figure 2).

These tools, which use ML and conversational agents—another application of AI—are also particularly suitable for addressing and continuing care during the COVID-19 pandemic (see Box 1 for key definitions). In fact, the spread of COVID-19 has catalyzed many digital health and AI-related tools to augment personal and population health in the U.S. and in many other parts of the world.

While a number of AI applications for care delivery outside of the hospital and clinical setting in medical specialties ranging from cardiology to psychiatry are either currently available or in development, their reliability and true utility in improving patient outcomes are highly variable. In addition, fundamental logistical issues exist, including product scalability, inter-system data standardization and integration, patient and provider usability and adoption, and insurance reform that must be overcome prior to effective implementation of AI technologies. Broader adoption of AI in health care and long-term data collection must also contend with urgent ethical and equity challenges, including patient privacy, exacerbation of existing inequities and bias, and fair access, particularly in the context of the U.S.'s fragmented mix of private and public health insurance programs.

In this discussion paper, the authors outline and examine the opportunities AI presents to transform health care in new and evolving arenas of care, as well as the significant challenges surrounding the

sustainable application and equitable development and deployment that must be overcome to successfully incorporate these novel tools into current infrastructures. The discussion paper concludes by proposing

steps for institutional, governmental, and policy changes that may facilitate broader adoption and equitable distribution of AI-driven health care technologies and an integrated vision for a home health model.

Box 1: Key Artificial Intelligence Terminology Definitions

Artificial intelligence (AI) “refers to the capacity of computers or other machines to exhibit or simulate intelligent behavior” (Oxford English Dictionary).

Artificial intelligence writ large is comprised of several domains. Some of the critical terms utilized in this paper are defined below.

Machine learning (ML) is a family of statistical and mathematical modeling techniques that uses a variety of approaches to automatically learn and improve the prediction of a target state, without explicit programming. Machine learning can be applied for predictive analytics to uncover insights about current and future trends.

Natural language processing (NLP) enables computers to understand and organize human speech. **Conversational agents** can engage in two-way dialogue with humans using NLP to comprehend human speech and respond accordingly.

Source: Oxford English Dictionary; Witten, I. H., E. Frank, M. A. Hall, and C. J. Pal. 2016. *Data Mining: Practical Machine Learning Tools and Techniques*. Burlington, MA, USA: Morgan Kaufmann; and Manning, C. D., and H. Schütze. 1999. *Foundations of Statistical Natural Language Processing*. Cambridge, MA: MIT Press.

1 Surveying Key Examples of AI Outside the Hospital and Clinic Setting: Evaluating Current and Emerging Technologies

1.1 Implementing AI on the Individual Level for Better Personal Health

1.1.1 Telehealth and AI

Telehealth has been a long-standing element of health care delivery in the U.S. [6], but not until COVID-19 has it been considered vital to sustaining the connection between patients and providers. These electronic interactions can be materially enhanced by AI in reducing the response time for medical attention and in alleviating provider case load and administrative burden. For example, AI triaging for telehealth uses conversational agents embedded in a virtual or phone visit to stratify patients based on acuity level and direct them accordingly to the most appropriate care setting [7]. By reducing the risk of patient exposure, AI triaging platforms have been especially advantageous during COVID-19, and a number of health systems, retail clinics, and payers have implemented them to continue the facilitation of care services [8] and identify possible COVID-19 cases. On the federal level, the Centers for Disease Control and Prevention (CDC) has launched a “Coronavirus Self-Checker” based on Microsoft’s Bot technology to guide patients to care using CDC guidelines [9,10]. Outside of the urgency of the COVID-19 pandemic, natural language processing has also been used to transcribe provider-patient conversations during phone visits, which can assist providers in writing care plans after the call concludes and can be

useful to patients as a reference of what was discussed [11].

These integrations are the “tip of the iceberg” of the possibilities of AI in the telehealth domain. Given the escalating pressure amid the COVID-19 pandemic to continue regulatory and financial support for telehealth [12], one could envision a burgeoning variety of AI couplings with telehealth. The future capacity of AI might include using video- and audio- capture tools with facial or tonal interpretation for stress detection in the home or office, or the incorporation of skin lesion detection apps into real-time video for dermatological visits.

1.1.2 Using AI to Augment Primary Care Outside of the Clinical Encounter

In the last six years, there has been a significant increase in the use of consumer applications for patient self-management of chronic diseases, and to a lesser degree for patient-provider shared management through home health care delivery and remote monitoring [13].

In 2018, diabetic care witnessed the landmark approval by the U.S. Food and Drug Administration (FDA) of IDx-DR, a ML-based algorithm that detects diabetic retinopathy, as the first AI-driven medical device to not require physician interpretation [14,15]. Outside the hospital, several AI applications have been developed for diabetes self- management, including those that have shown

improvements in HgbA1c through AI analysis of photos of patient meals to assess calories and nutrients [16] and another pilot trial of fully automated coaching for prediabetics, showing decreases in weight and HgbA1c [17]. For insulin management in type 1 diabetes, multiple studies have found that using self-adaptive learning algorithms in conjunction with continuous glucose monitors and insulin pumps results in decreased rates of hypoglycemia and an increase in patients reaching their target glucose range [18]. In March 2020 in the United Kingdom, the first such tool was officially licensed and launched publicly [19]. For type 2 diabetes, a promising example is an FDA-approved diabetes management system called WellDoc that gives individualized feedback and recommendations on blood sugar management and has been shown to reduce HgbA1c levels significantly [20,21].

Other consumer tools, some with approvals by regulatory agencies, help monitor and support blood pressure control and vital sign checks. One app, Binah.ai, features a validated tool that can scan a person's face in good lighting conditions and report heart rate (HR), oxygen saturation, and respiration with high levels of accuracy [22,23]. In addition, an increasing number of AI virtual health and lifestyle coaches have been developed for weight management and smoking cessation.

1.1.3 Remote Technology Monitoring for Promoting Cardiac Health

Wearable and remote monitoring technology can assist in ushering in the next era of health care data innovation by capturing physiologic data in HSOHC [24]. In

the current clinic-based paradigm, data is captured in isolated snapshots and often at infrequent time intervals. For example, blood pressure is measured and recorded during clinic visits once or twice a year, which does not provide an accurate or longitudinal understanding of an individual's blood pressure fluctuations.

The current "Internet of Things" era has changed the landscape for wearable technology. Wearables can capture data from any location and transmit it back to a hospital or clinic, moving a significant piece of the health care enterprise to places where patients spend the bulk of their time. These measurements can then be coupled with machine learning (ML) algorithms and a user interface to turn the data into relevant information about an individual's health-related behaviors and physiological conditions.

Wearable technology has been applied to many health care domains, ranging from cardiology to mental health. Prominent examples of technologies are those that incorporate cardiac monitoring, such as HR and rhythm sensors, including the Apple Watch, iRhythm, and Huawei devices. These devices are quite popular and, in the case of the Apple Watch, have received FDA approval as a medical device to detect and alert individuals of an irregular heart rhythm, a condition called atrial fibrillation [25]. Atrial fibrillation is associated with reduced quality of life and can result in the formation of blood clots in the upper heart chambers, ultimately leading to increased risk of stroke. Theoretically, enabling diagnosis of this condition outside of the clinic could bring patients to medical

attention sooner and, in turn, considerably reduce the risk of stroke.

However, the efficacy of some of these devices in relation to improving patient outcomes (increased quality of life and longevity) through detection of abnormal rhythms remains unproven, and there have been some concerns regarding the accuracy of the ML algorithms. For example, some of the Huawei and Apple Watch studies suggest that the devices seem to work well in sinus rhythm (beating normally at rest), but underestimate HR at higher rates in atrial fibrillation or in elevated sinus rhythm (i.e., with exercise) [26,27].

Hypertension, or high blood pressure, is another example of a highly prevalent, actionable condition that merits surveillance. Hypertension affects approximately 45 percent of Americans and is associated with heart failure, stroke, acute myocardial infarction, and death [28]. Sadly, hypertension control rates are worsening in the U.S., which will have downstream effects in most likely increasing the prevalence of cardiovascular disease [29]. Unlike smartwatch devices, blood pressure cuffs have been commercially available for decades. Today, several blood pressure manufacturers, including Omron, Withings, and others, offer cuffs that collect and transmit blood pressure measurements along with data like HR to health care providers [30,31]. Collecting longitudinal, densely sampled HR and blood pressure data in these ways allows for nuanced pattern detection through ML to predict increased risk of cardiovascular events like stroke or heart failure and, in turn, triage patients for medication management or more intensive

treatment. Ultimately, such prognostic capabilities could be embedded into the device itself. However, establishing the accuracy of data capture measurements relative to traditional sphygmomanometry is challenging because of the lack of scientific assessment standards [32].

In addition to established measurement standards, remote blood pressure monitoring devices should be coupled with a system to deliver interventions based on the data. One such option includes ML-powered smartphone apps paired with remote monitoring devices. The apps should effectively provide behavioral therapy for hypertensive patients [33], assess adherence to interventions, and promote patient self-awareness [34]. In terms of patient outcomes, some studies suggest that home monitoring, when coupled with pharmacist-led medication management and lifestyle coaching, is associated with improved blood pressure control; other studies are neutral [35]. Historically, the traditional health care delivery system has been unsuccessful in blood pressure control, and moving management into the home settings shows promise [36].

1.1.4 Remote Sensing and Mobile Health (mHealth) for Behavioral and Psychiatric Care

The pursuit of precision medicine—“delivering the right treatments, at the right time, every time to the right person” [37,38]—has been a long-standing goal in medicine. In particular, for psychiatry, clinical psychology, and related disciplines, increased precision regarding the timing of interventions presents an important

opportunity for mental health care. In major depressive disorder, episodes of depression contrast with periods of relatively improved mood. In bipolar disorder, patients cycle between both manic and depressive episodes. For substance use disorders, patients may alternate between periods of use and disuse. At an even more granular level, risk for returning to substance use can be instigated at times by discrete stressors but at other times in the presence of substances or peers using substances. Poor sleep and other issues that affect self-regulation may exacerbate this risk some days but not others. In each of these examples, different interventions are better suited to each of these specific moments in time to improve mental health.

The synthesis of AI with “personal sensing” provides a powerful framework to develop, evaluate, and eventually implement more precise mental health interventions that can be matched to characteristics of the patient, their context, and the specific moment in time [39]. Today, sensors relevant to medical care are ubiquitous. Smartphones log personal communications by voice calls and text messages. Facebook posts, Instagram photos, tweets, and other social media activities are also recorded. Smartphone-embedded sensors know our location (via GPS) and activity level (via accelerometer), and can detect other people in our immediate environment (via Bluetooth). Smartwatches that can monitor our physiology and many other raw signals are increasing in popularity.

Personal sensing involves collecting these many raw data signals and combining them with ML algorithms to predict thoughts,

behaviors, and emotions, as well as clinical states and disorders. This synthesis of ML and personal sensing can revolutionize the delivery of mental health care beyond the one-size-fits-all diagnoses and treatments to personalized interventions based on vast amounts of data collected not only in health care settings but in situ.

To be clear, the field of personal sensing (or digital phenotyping) is nascent and rapidly evolving [39]. However, emerging evidence already demonstrates the potential of its signals to characterize relevant mental health states at any moment in time. For example, GPS, cellular communication logs, and patterns of social media activity have all been used to classify psychiatric disorders and prognosis over time [40,41]. Natural language processing of what people write on social media can also be used to sense cognitive or motivational states (depressed mood, hopelessness, suicidal ideation) that may be more difficult to monitor with nonverbal sensors [42]. Moreover, many of these promising signals are collected passively by people’s smartphones, such that they can be measured without burden. This allows for long-term, densely sampled, longitudinal monitoring of patients that will be necessary to provide precisely timed interventions (e.g., just-in-time adaptive interventions [43]) for psychiatric disorders that are often chronic and/or cyclical.

mHealth apps are also well positioned to deliver AI-assisted precision mental health care. Mobile apps without AI have been already developed and deployed for post-traumatic stress disorder, depression, substance use disorders, and suicidal ideation, among others [44,45], many of

which have been pioneered by the U.S. Department of Veterans Affairs. These applications can screen for psychiatric disorders, track changes over time, and deliver evidence-based treatment or post-treatment support. They often include a variety of tools and services for patients including bibliotherapy, cognitive behavioral interventions, peer-to-peer or other social support, guided relaxation and mindfulness meditation, and appointment and medication reminders. In fact, many studies have demonstrated that patients are more expressive and more willing to report mental health symptoms to virtual human interviewers [46,47]. Moreover,

because smartphones are nearly always both on and available, mobile mental health care apps can provide immediate intervention while a patient is waiting for a higher level of care. Active efforts are underway to augment these systems with personal sensing AI to improve their ability to detect psychiatric risk in the moment and to recommend specific interventions among their available tools and services based on the characteristics of the patient and the moment in time [43,48].

2 Leveraging AI and Patient-Level Data from Remote Monitoring Tools to Gather Population-Level Insight

2.1 Integrating AI into Population Health Strategies

Since population health takes a holistic philosophy about caring for a large group of patients' health throughout their lives and all their activities, health management at this level necessarily goes far outside the bounds of a traditional medical encounter and into the daily lives of patients. A variety of integrated care delivery mechanisms have been used to improve population-level wellness and health, in many cases through novel partnerships and collaborations [49]. With the ongoing development of increasingly refined AI applications for individual use, next-generation population health strategies include analysis of aggregate patient-level data geared toward identification of broader population health trends and habits. Furthermore, these large-scale datasets set the stage for population-level AI algorithms for the purposes of epidemiological prediction, fueling a synergistic and powerful feedback loop of personal and population health innovation.

In the U.S., much of population health is managed and prioritized by insurance companies, employers, and disease management companies and increasingly by accountable care organizations and risk-bearing health care delivery organizations, whose primary aim is to decrease wasteful spending and improve health care quality by proactively engaging with and intervening for patients. The question becomes how to precisely identify these

patients at the right time in their care journey, so as to not engage them too late—after the health care decision is made and costs are no longer avoidable—or to engage them too early, and therefore waste administrative resources in engaging them. This need has given rise to the field of predictive analytics, which increasingly leverages AI to improve the effectiveness and efficiency of these programs. In the health care industry, these analytics typically rely on medical and pharmacy claims data, but are increasingly integrating a more diverse set of data, including health risk assessment data, electronic health record data, social determinants of health data—and even more recently, data from connected health devices and from transcribed call and messaging data between patients and these managed care organizations.

There has been tremendous interest and investment in deploying sensors, monitors, and automated tracking tools that, when combined with AI, can be used for population health management [50]. These tools and systems have been applied with varying degrees of sophistication to a wide variety of acute and chronic diseases, such as for diabetes and hypertension (described in previous sections), monitoring patients in rehabilitation [51], ongoing cardiovascular care [52], mental health care, falls [53], or dementia and elder care [54]. This category of potential applications distinguishes itself from self-management-related AI through the primary users of the systems. In this domain, the users are health care

professionals seeking to manage population health through information synthesis and recommendations.

Just like on an individual level, these algorithms remotely and passively detect physical and physiologic indicators of health and pathology, integrate them with patient-level environmental or health care system data, and generate insights, recommendations, and risks for many conditions. The challenges in this domain are melding disparate data—some from sensing information, some from image tracking, some from voice and audio analysis, and some from inertial or positional data—with more traditional medical data to improve outcomes and care.

2.2 Improving Medication Adherence with AI Tools

Another key challenge that population health faces is a lack of medication adherence. In some disease treatments, up to 40 percent of patients misunderstand, forget, or ignore health care advice [55]. Promotion of adherence to medical therapy is a complex interaction between patient preferences and autonomy; health communication and literacy; trust between patients/caregivers and the clinical enterprise; social determinants of health; cultural alignment between patients, caregivers, and health care professionals; home environment; management of polypharmacy; and misunderstandings about the disease being treated [56]. Numerous examples of adherence challenges abound, from treatments of chronic obstructive pulmonary disease [57],

asthma [58], diabetes [59], and heart failure [60].

There is tremendous opportunity for AI to identify and mitigate patient adherence challenges. One example of how AI might assist in improving adherence is in the case of direct oral anticoagulants in which an AI system embedded in smartphones was used to directly observe patients taking the medications. The AI incorporated imaging systems such as facial recognition and medication identification as well as analytics to identify those at high risk or to confirm delay in administration. Those who were identified as being at high risk were routed to a study team for in-person outreach as needed [61]. In a 12-week randomized controlled trial format, the AI arm had 100 percent adherence and the control arm had 50 percent adherence by plasma drug concentration level assessment. There are other notable examples in this area of medication adherence, such as with tuberculosis [62] and schizophrenia treatments [63].

2.3 AI Efforts in Public and Environmental Health

There is a strong need and opportunity for the use of AI technologies in public health, with opportunities that include information synthesis, outbreak detection, and responsible, appropriately governed, ethical, secure, and judicious syndromic surveillance. Public health has been incorporating and leveraging AI technologies for a number of years, and many countries have syndromic surveillance systems in place, such as RAMMIE in the U.K. [64]. As a subdomain of public health, environmental health has applied ML

techniques to tremendously benefit from the wide integration of publicly available data sources. One example is the need to assess toxicity *in silico* among chemicals used in commercial products, with over 140,000 monoconstituent chemicals in use and safety studies in less than 10 percent of them, not counting the vast number of chemical admixtures and metabolites [65,66,67]. There are important implications for environmental impacts in overall determinants of health along with genetic and chronic disease data, and AI will be critical in allowing the effective analysis of these types of data.

Another key area is the estimation of exposure histories and magnitude of patients over time, which requires diverse data ranging from location history, environmental conditions in areas of exposure, and subsequent evaluation and integration of said data into overall disease risk and clinical management strategies [68]. This also requires complex capacities in geospatial analysis and transformation [69]. In addition, the emphasis on geography and location mapping to assess potential outbreaks and environmental exposures is important for air pollution modeling [70]. AI-driven air pollution modeling uses a combination of satellite data, fixed monitoring, and professional and personal mobile monitoring devices to conduct complex assessments [71]. However, sensors such as PurpleAir require individuals to pay and install them in their homes and communities [72]. Thus, access is limited to those who have the privilege of disposable income. There have also been novel applications in assessing and informing public health policy with regard to neighborhood physical activity and

assessment of greenspace access, as well as access to healthy food outlets and grocery stores [73].

2.4 Combating COVID-19 with AI-Assisted Technologies

AI interpretation and human review of incoming data for syndromic surveillance provided early warning of the recent COVID-19 pandemic. The first early warning alert of a potential outbreak was issued on December 30, 2019, by the HealthMap system at Boston Children's Hospital, while four hours earlier a team at the Program for Monitoring Emerging Diseases had mobilized a team to start looking into the data and issued a more detailed report 30 minutes after the HealthMap alert [74]. BlueDot also issued an advisory on December 31, 2019, to all its customers [75]. These systems are interconnected and share data using a complex system of machine learning and natural language processing to analyze social media, news articles, government reports, airline travel patterns, and in some cases emergency room symptoms and reports [76,77,78]. Another set of ML algorithms consumes these processed data to make predictions about possible outbreaks [79].

In addition, wearable devices could serve an important role in the surveillance of high prevalence conditions, for which COVID-19 provides an immediate and important application. Fever alone provides inadequate screening for COVID-19 infection [80], but combining temperature with HR, respiratory rate, and oxygen saturation—all of which can be captured via wearable devices—could aid in triage and diagnosis. Prior research related to

influenza, in which investigators found that Fitbit data among 47,249 users could reliably predict prevalence rates estimated by the CDC, supports the role of wearables in infectious disease surveillance [81]. Indeed, randomized trials to test this hypothesis in relation to COVID-19 are underway [82], while others are using

wearables for COVID-19 tracking outside of the research enterprise [83]. Furthermore, many wearables can provide location data when linked to a smartphone, opening the door for geographic outbreak monitoring.

3 Development and Integration of Health-Related AI Tools: Overarching Logistical Challenges and Considerations

AI development and integration, especially of those devices deployed in HSOHC, face several logistical challenges in the health care marketplace. The authors of this discussion paper focus on six major categories of challenges that have been carefully documented in the literature and in practice: data interoperability and standardization, data handling and privacy protection, systemic biases in AI algorithm development, insurance and health care payment reform, quality improvement and algorithm updates, and AI tool integration into provider workflows.

3.1 Data Interoperability and Standardization

Logistical challenges to technology development and integration with virtual care systems include the challenges inherent to health care data collection, aggregation, analysis, and communication. In particular, AI-based programs must contend with data interoperability standards that have been created to ensure that data can be reliably transferred between scheduling, billing (including electronic health records), laboratory, registry, and insurer entities, as well as third party health data administrators, and ultimately be actionable to end users. Common data interoperability standards for health care data (e.g., the Health Level Seven standards [84] and its Fast Healthcare Interoperability Resources specification [85]) have helped to enhance communication among AI developer teams, data analysts, and engineers working on

other health care platforms, such as electronic health records. Nevertheless, considerable time can be spent by AI developers on extraction, transformation, and loading of data into different formats to both input and output data from AI platforms to health care data systems.

Often a major hurdle to AI development has been the personnel effort and time needed for data organization and cleaning, including the development of a strategy to address unclear data definitions and missing data [86]. The Observational Health Data Sciences and Informatics (or OHDSI, pronounced “Odyssey”) program involves an interdisciplinary collaboration to help address these issues for data analytics, and has introduced a common data model that many AI developers are now using to help translate and back-translate their health care data into a standard structure that aids communication with other health data management systems [87].

3.2 Data Handling and Privacy Protection

AI developer teams may also be subject to state and federal privacy regulations that affect sharing, use, and access to data for use in training and operating AI health care tools. As the major federal medical privacy statute, the Health Insurance Portability and Accountability Act (HIPAA) applies to “HIPAA-covered entities,” including health care providers such as clinics and hospitals, health care payers, and health care clearing houses that process billing information.

HIPAA-covered entities are subject to the HIPAA Privacy Rule, the federal medical and genetic privacy regulation promulgated pursuant to HIPAA. However, many entities that handle health-related information are not HIPAA-covered. Such entities can include many medical device and wearable/home monitoring manufacturers and medical software developers, unless they enter into “Business Associate Agreements” with organizations that do qualify as HIPAA-covered entities. Overall, because HIPAA is targeted to traditional health care providers, it often does not cover health AI companies that do not intersect or work closely with more traditional organizations.

Because the HIPAA Privacy Rule is directed at private sector players in health care, Medicare data and other health data in governmental databases are governed by a different statute, the Federal Privacy Act. State privacy laws add a layer of privacy protections, because the HIPAA Privacy Rule does not preempt more stringent provisions of state law. Several states, such as California, have state privacy laws that may cover commercial entities that are not subject to HIPAA, and which may provide more stringent privacy provisions in some instances. This means that companies that operate across multiple states may face different privacy regulatory requirements depending on where patients/clients are located.

When AI software is developed by a HIPAA-covered entity, such as at an academic medical center or teaching hospital that provides health care services, data must be maintained on HIPAA-compliant servers (even during model training) and not used

or distributed to others without first complying with the HIPAA Privacy Rule’s requirements. These requirements include that HIPAA-covered entities must obtain individual authorizations before disclosing or using people’s health information, but there are many exceptions allowing data to be used or shared for use in AI systems without individual authorization. An important exception allows sharing and use of data that have been de-identified, or had key elements removed, according to HIPAA’s standards [88]. Also, individual authorization is not required (even if data are identifiable) for use in treatment, payment, and health care operations (such as quality improvement studies) [88]. This treatment exception is particularly broad, and the Office for Civil Rights in the U.S. Department of Health and Human Services, which administers the HIPAA Privacy Rule, has construed it as allowing the sharing of one person’s data for treatment of other people [89]. This would allow sharing and use of data for AI tools that aim to improve treatment of patients.

Data also can be shared and used with public health authorities and their contractors, which could support data flows for public health AI systems [90]. Data can be shared for use in AI research without consent (including in identifiable form) pursuant to a waiver of authorization approved by an institutional review board or privacy board [90]. Such bodies sometimes balk at approving research uses of identifiable data, but the HIPAA Privacy Rule legally allows it, subject to HIPAA’s “minimum necessary” standard, which requires a determination that the identifiers are genuinely necessary to accomplish the purpose of the research [91,92]. These and

various other exceptions, in theory, allow HIPAA-covered care providers to use and share data for development of AI tools. However, all of HIPAA's authorization exceptions are permissive, in that they allow HIPAA-covered entities to share data but do not require them to do so.

Another concern is that much of today's health-relevant data, such as those from fitness trackers and wearable health devices, exist outside the HIPAA-regulated environment. This is because, as discussed above, HIPAA regulates the behavior of HIPAA-covered entities and their business associates only, leaving out many other organizations that develop AI. This has two implications: (1) the lack of privacy protection is of concern to consumers, and (2) it can be hard to access these data, and to know how to do it ethically, absent HIPAA's framework of authorization exceptions. Ethical standards for accessing data for responsible use in AI research and AI health tools are essential. Otherwise, public trust will be undermined.

There are an increasing number of publicly available and de-identified datasets that will allow for model comparisons, catalogued in the PhysioNet repository for biomedical data science and including the Medical Information Mart for Intensive Care dataset that involves intensive care unit data [93,94]. As most of these data are from research or hospital contexts, they highlight the need for more public, de-identified data from outpatient settings including telemedicine and patient-driven home monitoring devices.

3.3 Systematic Biases in AI Algorithm Development

Beyond data standards and regulations, a major challenge for AI developers in the U.S. health care environment is the risk that AI technologies will incorporate racial, social, or economic biases into prediction or classification models. Moreover, even if training datasets are perfectly reflective of the U.S. general population, an AI system could still be biased if it is applied in a setting where patients differ from the U.S. population at large. Many biases do, however, reflect broader historical racism and societal injustices that further perpetuate health care inequalities. Once these biased data are incorporated into ML algorithms, the biases cannot easily be interrogated and addressed. For example, while de-identified health care data from payers is increasingly available to predict which patients are higher or lower cost, Black patients in the U.S. are disproportionately at risk for lower health care access, and thus lower cost relative to their illness level (because of inadequate utilization). This artificially lower cost occurs in spite of this population's higher burden of social ills that increase the risk of poor health outcomes, such as social stressors related to hypertension or poor food security that often worsens diabetes outcomes. Researchers have observed that AI prediction models that seek to determine which people need more outreach for home-based or community-based care were developed from cost data, without a correction for differential access, and thus biased predictions against predicting care needs for Black individuals [95]. Outside of the hospital and clinic settings, historically marginalized communities may face similar

barriers to access to technologies, algorithms, and devices. Indeed, recent surveys indicate that use of smartwatches and fitness trackers correlates with household income, but ethnicity-based differences are less pronounced, with Black and Latinx Americans reporting usage rates equivalent or higher than those of white Americans [96].

Developing AI tools is a process of active discovery and simultaneously subject to counterintuitive complexities, making the work of removing bias from health care data extremely complex. For example, observing equal treatment among groups may actually be indicative of a highly inequitable AI model [97]. Some groups may be properly deserving of higher attention because of disproportionate risk for a health care event, and therefore treating them equally would be an error [98]. Bias in the data itself is also paired with bias in outcomes, in that AI models can predict risk of an event such as health care utilization, but can also make suggestions for appropriate health care treatments. If the treatment recommendations are also biased, then disadvantaged groups may get erroneous advice more often, or simply not receive AI-aided advice while their counterparts who are better represented in the data receive the advantages of the AI-aided decision making [99]. To reiterate, using AI systems and tools that utilize biased data or biased processes will further entrench and exacerbate existing inequities and must be addressed before a system or tool is deployed.

3.4 Insurance and Health Care Payment Reform

A logistical challenge for AI use outside of the hospital and clinic setting that also challenges AI development and integration is the U.S. health care payment landscape. Many, albeit an increasingly smaller percentage, of health care payments from commercial insurers or government payers (e.g., Medicare Part B) to health care delivery entities are in the form of fee-for-service payments for in-person visits or procedures. Health care delivery entities generate revenues by billing payers with attached billing codes that reference negotiated payments for different services, from routine office visits to a primary care provider to surgeries, and programs outside of the hospital and clinic setting are incentivized to fit into the fee-for-service model if they are to be paid for by traditional payment mechanisms. While telemedicine visits (video and phone) are now covered by most payment entities, and, in the initial months of the COVID-19 pandemic, were reimbursed at an equal rate to in-person visits [100], the tools used to deliver such services are traditionally not reimbursed. For example, a physician could use many AI tools, remote sensing tools included, to help improve the quality or precision of diagnosis or therapeutic recommendations. Such tool use could be costly in personnel and computational time, and as discussed earlier, these tools can have questionable validity. However, the use of these tools would not necessarily be paid for, as its use would be considered implicit in conducting a medical visit, even though various diagnostic procedures with their own personnel and equipment costs (e.g., radiology) have their own payment

rates. A change in such policies to help pay for AI tools is one step toward helping AI applications both inside and outside of the hospital or clinic be paid for; the first billing code for an AI tool is one that helps to detect diabetic retinopathy. Still, it is unclear at this point how much this new code will pave the way for payers to accept the code and pay for AI services within a fee-for-service payment structure [101].

While much billing in the U.S. health care landscape remains a fee-for-service billing approach in which services are rendered and reimbursed according to negotiated rates for services that a payer covers, alternative payment models exist that may alter the AI payment landscape. Such alternative models include capitation payments (per-patient, per-month payments) that payers could increase for practices or providers that incorporate high-quality, externally validated AI tools into their practice, and value-based payments (bonuses for providers who show that their use of AI tools has improved outcomes). Capitation payments have now been increasingly adopted for routine health care delivery in many managed-care environments, but as of yet there are no adjustments for the use of AI tools. Value-based payments have been more experimental and limited to isolated pilot programs or novel health care organization arrangements to date [102,103], and such value-based payments may incentivize use of AI tools outside of the hospital or clinic if they improve clinical outcomes, whether or not such tools require intervention within a medical visit. Both capitation and value-based payments could be adjusted to explicitly reward the use of AI tools for better outcomes

3.5 Quality Improvement and Algorithm Updates

To further aid in the adoption and implementation of AI tools into clinical practice, particularly into telemedicine and virtual care environments, it is important to solidify the practice of quality improvement and to responsibly navigate the challenges of ownership, responsibility, decision making, and liability. As telehealth and virtual care platforms continue to improve their user experiences, it becomes critical that the AI tools they rely on—from symptom checkers that direct providers toward considering particular diagnoses, to scheduling and billing tools that aid patients, to personalized recommendation systems that help remind patients of routine cancer screening and available health coaching—must have a built-in feedback process. There are numerous examples of complex chronic diseases that require detailed self-management, such as blood glucose monitoring and adjustments in daily calorie intake or insulin administration for diabetes [104,105] or management of diet, salt, exercise, and medication dosing after heart failure [106]. The key challenges in this subdomain are those of appropriate data collection through patient-facing technologies—whether linked glucose monitors, blood pressure monitoring, calories and types of food eaten, steps taken, and other features—and integrating AI algorithms and tools safely into cautions and recommendations along with information synthesis to patients.

AI-driven personal sensing algorithms will likely have limited shelf lives for a variety of reasons [39]. Given the rapid pace of

development, there is considerable churn in both the software and hardware that are used to measure these signals. As sensor software and hardware are updated, raw data signals will change. There will also be shifts in how patients interact with these software and hardware and where these digital interactions happen that necessitate changes in the devices and signals that are monitored. Additionally, as modes of data collection become more precise (e.g., more advanced HR and glucose monitors), algorithms can be regularly retrained with these more reliable data to harness greater predictive accuracy.

For example, smartphone use has changed dramatically in recent years. Communications have shifted from voice to SMS, and SMS itself has moved from native smartphone apps to separate applications like Facebook Messenger, Snapchat, and WhatsApp. Video conferencing has also been rapidly adopted during COVID-19.

More fundamentally, the meaning of the raw signals may change over time as well. Language usage and even specific words that indicate clinically relevant effects or stressors have a temporal context that may change rapidly based on sociopolitical or other current events (e.g., the COVID-19 pandemic, the Black Lives Matter movement and associated protests, political election cycles). Patterns of movement and their implications can change as well (e.g., time spent at home or in the office). This limited shelf life for personal sensing algorithms must be explicitly acknowledged and processes must be developed to monitor and update the performance of the algorithms over time to keep them current and accurate.

3.6 AI Tool Integration into Provider Workflows Outside of the Hospital or Clinic

Most health care systems today have training and execution of quality improvement programs that identify important problems such as medical errors and undergo cycles of planning, piloting, studying, and modifying workflows to reduce such problems, often using a Lean framework for improvement [107,108]. AI tools outside of the hospital or clinic can be integrated into that workflow to improve their effectiveness, efficiency, and utilization. Such tools may be vital for quality improvement of services outside of hospital or clinic settings, as well as to scale and diffuse such technologies among teams that may be initially skeptical about their value. Issues of usability have significant implications for provider adoption. The increasing volume of data collected through wearable technology can overwhelm providers who are already experiencing high rates of alert fatigue and clinician burnout. Ensuring usability entails developing an accessible user interface and presenting information in a clear and actionable way.

Inherent to the implementation and improvement process is the dilemma of how to ensure that the business models underlying AI tool innovation are tailored to their users. It is often assumed that AI tools will have a single user: a provider or a patient. Typically, however, AI tools are used in a mixed manner because of availability and access of the tools in shared environments or in the transition of settings from the home to clinical visits where providers use and show the results or

visualizations from a tool. Therefore, communications to mixed groups of users are important to consider [109].

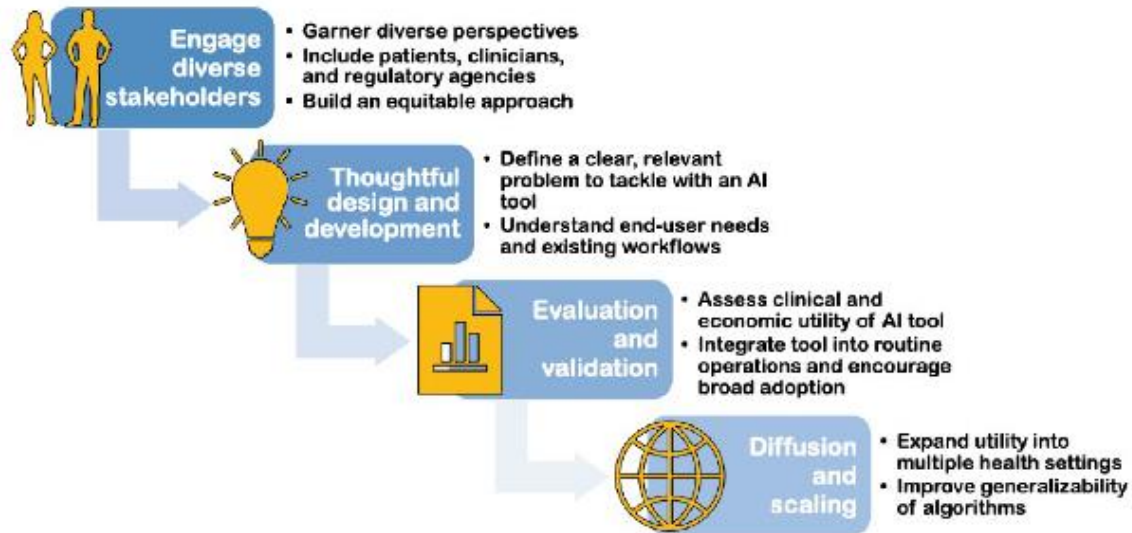
3.7 Practical Steps for Integration of New AI Tools into the U.S. Health System

To help overcome the challenges of AI tool development and deployment, the authors of this discussion paper suggest considering a series of steps for taking a model from design to health system integration and highlight challenges specific to each step (see Figure 3).

clinicians across the care continuum, and that the approach to building the technology does not “automate inequality” [110] or build “weapons of math destruction” [111].

The second step should be careful and thoughtful model design and development. During the model design step, AI developers often curate a dataset, secure initial research funding to develop a model, and build an interdisciplinary team with technical and clinical experts. A critical challenge during this stage is to develop a product that solves a real, relevant problem for end users. AI developers looking to

Figure 3: Translational Path for AI into Clinical Care



Note: Depicted here are the key steps towards successful implementation of AI applications in HSOHC into clinical workflows, including engagement of diverse stakeholders, thoughtful application design and development, evaluation and validation, and diffusion and scaling of technologies.

The first step on the translational path for AI into clinical care is to engage a wide range of stakeholders to ensure that the tools developed account for a wide range of perspectives, including patients and

translate their technologies into practice need to approach the technical task of training a model as part of a product development process. As described by Clayton Christensen, “when people find themselves needing to get a job done, they essentially hire products to do that job for

them” [112]. For an AI model to be used in practice, the model must successfully complete a job for an end user, be it a patient or expert clinician. Unfortunately, this goal usually involves more than a straightforward modeling task, and models need to be conceptualized as a single component within a more complex system required to deliver value to users. Deeply understanding the “job to be done” requires close collaboration with end users and interdisciplinary collaboration [113]. In contrast, many AI and ML technologies are built without clinical collaborators and leverage readily available datasets to model a small set of outcomes [114,115]. Teams that successfully navigate the design and develop steps deeply understand user needs and have developed an AI technology potentially able to solve a problem.

The third step on the translational path is to evaluate and validate the new AI tool. During this step, AI developers often evaluate the clinical and economic utility of a model using retrospective, population-representative data. Models may then undergo temporal and external validation, and then be integrated into a care delivery setting to assess clinical and economic impact. Unfortunately, many AI models undergo *in silico* experiments using retrospective data and do not progress further [116]. These experiments can provide preliminary data on the potential utility of a model, but do not provide evidence of realized impact. Prospective implementation in clinical care requires both clinical and technical integration of the AI model into routine operations. Technical integration requires sophisticated infrastructure that automates and monitors extraction, transformation, and load

processes that ingest data from data sources and write model output into workflow systems [117]. Clinical integration requires the design and successful implementation of clinical workflows for end users. There is a rich literature on innovation adoption in health care, and adoption barriers and facilitators specific to AI are emerging [118,119]. Teams that successfully navigate the “evaluate and validate” steps are able to demonstrate the clinical and economic impact of an AI model within at least one setting.

The fourth and fifth steps on the translational path are to diffuse and scale. To date, no AI model has efficiently scaled across all health care settings. Most models have been validated within silos or single settings, and a small number of AI technologies have undergone peer-reviewed external validations [120]. Furthermore, while some AI developers externally validate the same model in multiple settings, other teams take a different approach. For example, there is ongoing research into a generalizable approach to train site-specific *Clostridium difficile* models (a model of within-hospital infection) by which each hospital has a local model [120,121]. Externally validating and scaling a model across settings also introduces data quality challenges as institutional datasets are not interoperable and significant effort is required to harmonize data across settings [122].

4 Equitable and Humanistic Deployment of Health-Related AI Tools: Legal and Ethical Considerations

Health-related AI tools designed for use outside hospitals and clinics present special legal, regulatory, and ethical challenges.

Chief among the legal challenges are:

1. the safety of patients, consumers, and other populations whose well-being may depend on these systems; and
2. concerns about accountability and liability for errors and injuries that will inevitably occur even if these tools deliver hoped-for benefits such as improving patient care and public health, reducing health disparities, and helping to control health care costs.

Major ethical challenges are:

3. ensuring privacy and other rights of persons whose data will be used or stored in these systems;
4. ensuring ethical access to high-quality and inclusive (population representative) input data sets capable of producing accurate, generalizable, and unbiased results; and
5. ensuring ethical implementation of these tools in home care and other diverse settings.

4.1 Safety Oversight

The sheer diversity of AI tools discussed in this paper implies a nonuniform and, at times, incomplete landscape of safety oversight. Policymakers and the public often look to the FDA to ensure the safety of health-related products, and the FDA is currently working to develop suitable frameworks for regulating software as a medical device [123], including AI/ML software [124]. For software intended for use outside traditional care settings, however, the FDA cannot by itself ensure safety. Involvement of state regulators, private sector institutional and professional bodies, as well as other state and federal consumer safety regulators such as the Federal Trade Commission and Consumer Product Safety Commission, will also be required. Coordination is crucial, however, and the FDA can use its informational powers to inform and engage the necessary dialogue and cooperation among concerned oversight bodies: state, federal, and nongovernmental.

The 21st Century Cures Act of 2016 delineated types of health-related software that the FDA can and cannot regulate [125]. In 2017, the agency announced its Digital Health Innovation Action Plan [126] followed by the Digital Health Software Pre-certification (Pre-Cert) Program [127] and published final or draft guidance documents covering various relevant software categories, including consumer-grade general wellness devices such as wearable fitness trackers [128], mobile medical applications [129], and clinical decision support software [130].

The software discussed in this paper raises special concerns when it comes to regulatory oversight. First, AI software intended for population and public health applications is not subject to FDA oversight, because it does not fit within Congress's definition of an FDA-regulable device intended for use in diagnosing, treating, or preventing disease of individuals in a clinical setting [131]. Second, there is a potential for software designed for one intended use to be repurposed for new uses where its risk profile is less understood. For example, consumer-grade wearables and at-home monitoring devices, when marketed as general wellness devices, lie outside the FDA's jurisdiction and do not receive the FDA's safety oversight. These devices might be repurposed for medical uses by consumers or by developers of software applications. Repurposing raises difficult questions about the FDA's capacity to detect and regulate potential misuses of these devices [132]. Consumers may not understand the limits of the FDA's regulatory jurisdiction and assume that general wellness devices are regulated as medical devices because they touch on health concerns.

Also pertinent to the home care setting, the FDA tried in 2017 to address "patient decision support" (PDS) software, where the user is a patient, family member, or other layperson (paid or unpaid care-givers) in the home care setting (as opposed to a medical professional in a clinic or hospital), but subsequently eliminated this topic from its 2019 clinical decision support draft policy [130]. The regulatory framework for PDS software remains vague.

Even when a trained medical professional uses clinical decision support software (whether in a clinic, hospital, or HSOHC), patient safety depends heavily on appropriate application of the software. This is primarily a medical practice issue, rather than a medical product safety issue that the FDA can regulate. State agencies that license physicians, nurses, and home care agencies have a crucial role to play, as do private-sector institutional and professional bodies that oversee care in HSOHC. A singular focus on the FDA's role as a potential software regulator distracts from the need for other regulatory bodies to engage with the challenge of ensuring proper oversight for health care workers applying AI/ML software inside and outside traditional clinical settings.

4.2 Accountability and Liability Issues

AI tools for public health raise accountability concerns for the agencies that rely on them, but appear less likely to generate tort liability, because of the difficulty of tracing individual injuries to the use of the software and because public health agencies often would apply such software to perform discretionary functions that enjoy sovereign immunity from tort lawsuits.

Concerns about accountability and liability are greater for private sector users, such as a health care institution applying AI population health tools or quality improvement software that recommends approaches that, while beneficial on the whole, may result in injuries to specific patients. It remains unclear what duties institutions have (either ethically or legally,

as elaborated in future tort suits) to inform patients about the objectives of population health software (for example, is the software programmed to reduce health care costs or to ration access to scarce facilities such as ICU beds) and how these objectives may affect individual patients' care.

More generally, AI tools are relatively new, so there is not yet a well-developed body of case law with which to predict the tort causes of action that courts will recognize or the doctrines courts may apply when hearing those claims. Possible claims may include malpractice claims against physicians, nurses, or home care providers who rely on AI decision-support tools; direct suits against home care agencies and institutions for lax policies and supervision in using such systems; and suits against software developers including possible product liability actions [133]. Scholars are actively engaged in exploring the liability landscape, but uncertainties will remain until courts resolve cases in this field.

4.3 Data Privacy and Individual Consent

Many types of data that would be useful for training and ongoing operation of AI health systems, as well as the data such software may generate and store about individuals who use and rely on them, may fall outside the umbrella of HIPAA privacy protections. As discussed above, HIPAA privacy protections generally apply only to HIPAA-covered entities such as providers and payers for health care services. HIPAA's coverage excludes many device and software developers, governmental public health agencies (which may be governed by

the Federal Privacy Act), and research institutions that are not affiliated with HIPAA-covered clinics or teaching hospitals. Data generated and used in HSOHC have spotty privacy protections, subject largely to a patchwork of state privacy laws. Data bearing on social determinants of health, behavioral factors, and environmental exposures are crucial in developing AI tools tailored to diverse subpopulations, yet such data often arise in non-HIPAA-covered environments with weak oversight of sharing and data uses, creating ethical challenges such as "surveillance capitalism" [134] and "automating inequality" [110]. The absence of a uniform floor of federal privacy protections for all types of health-relevant data in all settings (medical and nonmedical) is a factor that may hinder future development of promising AI technologies in the U.S. and undermine public trust in the AI tools that do managed to be developed.

The lack of uniform privacy protections in the U.S. also encourages heavy reliance on de-identified data by AI system developers, who are navigating the ethical challenges of data uses without clear regulatory protections and guidance. The reliance on de-identified data, however, is a "second-best" solution that can diminish the accuracy and generalizability of AI tools available in the U.S. As to the concerns, noted earlier, about the risks of re-identification of such data [135], there is scholarly debate about how real these risks actually are, with empirical studies indicating the risk is considerably lower than portrayed in the popular press [136]. On one hand, there is increased awareness that de-identified data sets can be combined to re-identify individuals, a

process known as data aggregation, suggesting that de-identification may not be a complete solution to privacy concerns [137]. The public is concerned about wide dissemination of their de-identified data. On the other hand, a serious—and less understood—concern relates to the quality and generalizability of AI software developed using de-identifiable data. The process of de-identifying data diminishes its usefulness and can hinder the creation of high-quality longitudinal data sets to support accurate results from AI health tools. Moreover, de-identification strips away information that may be needed to audit data sets to ensure that they are inclusive and generalizable across all population subgroups. This can increase the danger of biased data sets that fail to produce accurate results for all racial, geographical, and gender-related subgroups [138]. Rather than rely on de-identification as a weak proxy for privacy protection, the U.S. needs a strong framework of meaningful privacy protections that would allow the best available data to be used.

As AI systems move into health care settings, patients may not be aware when AI systems are being used to inform their care [139]. Whether and when informed consent is appropriate has not received adequate discussion [140]. On one hand, it is not standard practice for physicians to inform patients about every medical device or every software tool used in their care. On the other hand, some functions performed by AI software tools (such as deciding which therapy is best for the patient) may rise to a level of materiality where consent becomes appropriate. The bioethics community, health care accreditation organizations, and state medical regulators need to engage

with the challenge of defining when, and under what circumstances, informed consent may be needed. A related topic is how future consent standards, designed for traditional clinic and hospital care settings, could be applied and enforced in HSOHC.

4.4 Ensuring Equitable Use of AI in Health

mHealth apps can address many of the current disparities that result in unmet health care needs. For example, in reference to mental health, the geographic distribution of licensed clinical psychologists across the U.S. is highly uneven, with large swaths of rural America significantly underserved [141]. Office visits with psychiatrists and psychologists can be infrequent, difficult to schedule, and typically not available at moments of peak need. Mental health care is costly, and those with greatest need are often uninsured or otherwise unable to afford necessary care [142]. In contrast, access to mental health care via mHealth apps is not limited by either geographic or temporal constraints. Furthermore, the percentage of Americans who now own smartphones is 81 percent, up from just 35 percent in 2011 [143]. Equally important, Black and Latinx adults have smartphones in shares similar to whites and are more likely than whites to use smartphones to access information about health conditions [144].

Nevertheless, equity in access to mHealth apps and technology will not happen without attention and planning. For example, the majority of telehealth visits during the initial months of the COVID-19 pandemic were based on pre-existing provider/patient relationships [145]. This

emphasis on continuity of care, rather than establishing new care relationships, suggests that individuals in medically underserved communities may not be benefiting from the shift to digital home health. Likewise, individuals who are already underserved because of racial and other disparities may struggle to access mHealth apps or other AI applications. mHealth apps that are powered by AI personal sensing can address racial and other health disparities but only if they are thoughtfully designed, developed, and distributed [146] with the intention of reducing biases and the digital divide. The

development of these AI algorithms using personal sensing signals must include data from people from racial and ethnic minorities and other underserved groups to account for differences in how these signals function in different groups of people [147]. Algorithms must also be carefully designed and scrutinized to avoid reinforcing contemporary racial and other biases by instantiating them in these algorithms [111]. Thoughtful infrastructure measurement, regulation, and accountability are necessary for the distribution and oversight of these mHealth apps.

5 Setting the Stage for Impactful AI Tools in HSOHC: Calls to Action

AI-powered digital health technology is a rapidly developing sector that is poised to significantly alter the current landscape of health care delivery in the U.S., particularly as care extends beyond the walls of the hospital and clinic. As mHealth applications and personal health devices, including wearable technology, become increasingly ubiquitous, they enable large-scale collection of detailed, continuous health data that processed through AI can support individual and population health. Illustrated by the examples discussed here, these tools signal a paradigm shift in the traditional notion of clinical point-of-care to one that meets people where they are to deliver care. However, widespread adoption, secure implementation, and integration of these novel technologies into existing health care infrastructures pose major legal and ethical challenges. Concrete steps toward ensuring the success of AI health tools outside the hospital and clinic can include:

- Building broad regulatory oversight to promote patient safety by engaging organizations beyond the FDA, including other state, federal, and nongovernmental oversight bodies;
- Reconsidering the definition and implications of informed consent in the context of big data, AI algorithm

development, and patient privacy in HSOHC;

- Developing policy initiatives that push for greater data interoperability and device integration standards with hospital clinical systems, so as to enhance stress-free provider and consumer/patient usability;
- Recognizing and mitigating biases (racial, socioeconomic) in both AI algorithms and access to personal health devices by including population-representative data in AI development and increasing affordability and access (or insurance coverage) of personal health technology, respectively;
- Advocating for insurance and health care payment reform that incentivizes adoption of AI tools into physicians' workflow; and
- Establishing clarity in regard to liability for applications of health AI, with an eye to supporting rather than hindering innovation in this field.

It is important to acknowledge that many of these steps necessitate fundamental changes in governmental oversight of health care, industry-hospital communication, and the patient-provider relationship itself. However, approaching novel applications of AI in health with a critical but receptive mindset will enable the U.S. to lead in ushering in the next generation of health care innovation.

References

1. Tikkanen, R. and M. K. Abrams. 2020. *U.S. Health Care from a Global Perspective, 2019: Higher Spending, Worse Outcomes?* The Commonwealth Fund. Available at: <https://www.commonwealthfund.org/publications/issue-briefs/2020/jan/us-health-care-global-perspective-2019> (accessed September 29, 2020).
2. American Medical Association. 2019. *Medicare's major new primary care pay model: Know the facts*. Available at: <https://www.ama-assn.org/practice-management/payment-delivery-models/medicare-s-major-new-primary-care-pay-model-know-facts> (accessed October 26, 2020).
3. American Association of Family Physicians. 2020. *Chronic Care Management*. Available at: <https://www.aafp.org/family-physician/practice-and-career/getting-paid/coding/chronic-care-management.html> (accessed October 26, 2020).
4. National Science and Technology Council, Committee on Technology, Executive Office of the President. 2016. *Preparing for the Future of Artificial Intelligence*. Available at: https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf (accessed October 26, 2020).
5. Matheny, M., S. Thadaneys 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.
6. Institute of Medicine. 2012. *The Role of Telehealth in an Evolving Health Care Environment: Workshop Summary*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/13466>.
7. Kuziemsy, C., A. J. Maeder, O. John, S. B. Gogia, A. Basu, S. Meher, and M. Ito. 2019. Role of Artificial Intelligence within the Telehealth Domain. *Yearbook of Medical Informatics* 28(01): 035-040. <https://doi.org/10.1055/s-0039-1677897>
8. Wosik, J., M. Fudim, B. Cameron, Z. F. Gellad, A. Cho, D. Phinney, S. Curtis, M. Roman, E. G. Poon, J. Ferranti, J. N. Katz, and J. Tchong. 2020. Telehealth transformation: COVID-19 and the rise of virtual care. *Journal of the American Medical Informatics Association* 27(6): 957-962. <https://doi.org/10.1093/jamia/ocaa067>
9. Bitran, H. and J. Gabarra. 2020. *Delivering information and eliminating bottlenecks with CDC's COVID-19 assessment bot*. Available at: <https://blogs.microsoft.com/blog/2020/03/20/delivering-information-and-eliminating-bottlenecks-with-cdcs-covid-19-assessment-bot/> (accessed September 29, 2020).
10. Centers for Disease Control and Prevention. 2020. *COVID-19 Testing Overview*. Available at: <https://www.cdc.gov/coronavirus/2019-ncov/symptoms-testing/testing.html> (accessed October 26, 2020).

11. Menon, R. 2020. Virtual Patient Care Using AI. *Forbes*, March 26, 2020. Available at: <https://www.forbes.com/sites/forbestech-council/2020/03/26/virtual-patient-care-using-ai/#18befa3be880> (accessed September 29, 2020).
12. Haque, S. N. 2020. Telehealth Beyond COV- ID-19. *Technology in Mental Health*. <https://doi.org/10.1176/appi.ps.202000368>
13. Tran, V-T., C. Riveros, and P. Ravaud. 2019. Patients' views of wearable devices and AI in healthcare: findings from the ComPaRe e-cohort. *npj Digital Medicine* 2(53). <https://doi.org/10.1038/s41746-019-0132-y>
14. U.S. Food and Drug Administration (FDA). 2018. *FDA permits marketing of artificial intelligence-based device to detect certain diabetes-related eye problems*. Available at: <https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-artificial-intelligence-based-device-detect-certain-diabetes-related-eye> (accessed October 30, 2020).
15. Abramoff, M. D., P. T. Lavin, M. Birch, N. Shah, and J. C. Folk. 2018. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. *npj Digital Medicine* 1(39). <https://doi.org/10.1038/s41746-018-0040-6>
16. Oka, R., A. Nomura, A. Yasugi, M. Kometani, Y. Gondoh, K. Yoshimura and T. Yoneda. 2019. Study Protocol for the Effects of Artificial Intelligence (AI)-Supported Automated Nutritional Intervention on Glycemic Control in Patients with Type 2 Diabetes Mellitus. *Diabetes Therapy* 10: 1151-1161. <https://doi.org/10.1007/s13300-019-0595-5>
17. Everett, E., B. Kane, A. Yoo, A. Dobs, and N. Mathoudakis. 2018. A Novel Approach for Fully Automated, Personalized Health Coaching for Adults with Prediabetes: Pilot Clinical Trial. *Journal of Medical Internet Research* 20(2):e72. <https://doi.org/10.2196/jmir.9723>
18. Nimri, R., N. Bratina, O. Kordonouri, M. A. Stefanija, M. Fath, T. Biester, I. Muller, E. Atlas, Miller, A. Fogel, M. Phillip, T. Danne, and Battelino. 2016. MD-Logic overnight type 1 diabetes control in home settings: A multicenter, multinational, single blind randomized trial. *Diabetes, Obesity, and Metabolism*. <https://doi.org/10.1111/dom.12852>
19. Brierly, C. 2020. World's first artificial pancreas app licensed for people with type 1 diabetes in UK. *Medical Xpress*, March 16, 2020. Available at: <https://medicalxpress.com/news/2020-03-world-artificial-pancreas-app-people.html> (accessed November 5, 2020).

20. Quinn, C. C., S. S. Clough, J. M. Minor, D. Lender, M. C. Okafor, and A. Gruber-Baldini. 2008. WellDoc Mobile Diabetes Management Randomized Controlled Trial: Change in Clinical and Behavioral Outcomes and Patient and Physician Satisfaction. *Diabetes Technology and Therapeutics*. <https://doi.org/10.1089/dia.2008.0283>
21. Quinn, C. C., M. D. Shardell, M. L. Terrin, E. A. Barr, S. H. Ballew, and A. L. Gruber-Baldini. 2011. Cluster-Randomized Trial of a Mobile Phone Personalized Behavioral Intervention for Blood Glucose Control. *Diabetes Care* 34(9): 1934-1942. <https://doi.org/10.2337/dc11-0366>
22. Insurancenewsnet.com. 2020. *Binah.ai Awarded Best Digital Health Innovation by Juniper Research*. Available at: <https://insurancenewsnet.com/oarticle/binah-ai-awarded-best-digital-health-innovation-by-juniper-research#.X6SElmhKjIW> (accessed November 5, 2020).
23. Banerjee, R., A. D. Choudhury, A. Sinha, and Visvanathan. 2014. HeartSense: smart phones to estimate blood pressure from photoplethysmography. *SenSys '14: Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*. <https://doi.org/10.1145/2668332.2668378/>
24. Binkley, P. F. 2003. Predicting the potential of wearable technology. *IEEE Engineering in Medicine and Biology Magazine* 22(3). <https://doi.org/10.1109/MEMB.2003.1213623>
25. Su, J. 2018. Apple Watch 4 Is Now An FDA Class 2 Medical Device: Detects Falls, Irregular Heart Rhythm. *Forbes*, September 14, 2018. Available at: <https://www.forbes.com/sites/jeanbaptiste/2018/09/14/apple-watch-4-is-now-an-fda-class-2-medical-device-detects-falls-irregular-heart-rhythm/?sh=3ddcdf2b2071> (accessed October 30, 2020).
26. Koshy, A. N., J. K. Sajeev, N. Nerlekar, A. J. Brown, K. Rajakariar, M. Zureik, M. C. Wong, L. Roberts, M. Street, J. Cooke, and A. W. Teh. 2018. Smart watches for heart rate assessment in atrial arrhythmias. *International Journal of Cardiology* 266(1): 124-127. <https://doi.org/10.1016/j.ijcard.2018.02.073>
27. Wang, R., G. Blackburn, M. Desai, D. Phelan, L. Gillinov, P. Houghtaling, and M. Gillinov. 2017. Accuracy of Wrist-Worn Heart Rate Monitors. *JAMA Cardiology* 2(1): 104-106. <https://doi.org/10.1001/jamacardio.2016.3340>
28. Centers for Disease Control and Prevention. n.d. *Facts About Hypertension*. Available at: <https://www.cdc.gov/bloodpressure/facts.htm> (accessed October 30, 2020).
29. Muntner, P., S. T. Hardy, L. J. Fine, B. C. Jaeger, G. Wozniak, E. B. Levitan, and L. D. Colantonio. 2020. Trends in Blood Pressure Control Among US Adults With Hypertension, 1999-2000 to 2017-2018. *JAMA* 324(12): 1190-1200. <https://doi.org/10.1001/jama.2020.14545>
30. Omron. n.d. *Blood Pressure Monitors*. Available at: <https://omronhealthcare.com/blood-pressure/> (accessed November 5, 2020).

31. Withings. n.d. *Blood Pressure Monitors*. Available at: <https://www.withings.com/us/en/blood-pressure-monitors> (accessed November 5, 2020).
32. Ringrose, J. S., G. Polley, D. McLean, A. Thompson, F. Morales, and R. Padwal. 2017. An Assessment of the Accuracy of Home Blood Pressure Monitors When Used in Device Owners. *American Journal of Hypertension* 30(7): 683-689. <https://doi.org/10.1093/ajh/hpx041>
33. Guthrie, N. L., M. A. Berman, K. L. Edwards, K. J. Appelbaum, S. Dey, J. Carpenter, D. M. Eisenberg, and D. L. Katz. 2019. Achieving Rapid Blood Pressure Control With Digital Therapeutics: Retrospective Cohort and Machine Learning Study. *JMIR Cardiology* 3(1): e13030. <https://doi.org/10.2196/13030>
34. da Silva, V. Joao, V. S. Souza, R. G. da Cruz, J. M. V. M. de Lucena, N. Jazdi, and V. F. L. Junior. 2019. Commercial Devices-Based System Designed to Improve the Treatment Adherence of Hypertensive Patients. *Sensors* 19(20): 4539. <https://doi.org/10.3390/s19204539>
35. Margolis, K. L., S. E. Asche, A. R. Bergdall, S. P. Dehmer, S. E. Groen, H. M. Kadrmas, T. J. Kerby, K. J. Klotzle, M. V. Maciosek, R. D. Michels, P. J. O'Connor, R. A. Pritchard, J. L. Sekenski, J. M. Sperl-Hillen, and N. K. Trower. 2013. Effect of Home Blood Pressure Telemonitoring and Pharmacist Management on Blood Pressure Control: A Cluster Randomized Clinical Trial. *JAMA* 310(1): 46-56. <https://doi.org/10.1001/jama.2013.6549>
36. Khatib, R., J. Schwalm, S. Yusuf, R. B. Haynes, M. McKee, M. Khan, and R. Nieuwlaat. 2014. Patient and Healthcare Provider Barriers to Hypertension Awareness, Treatment and Follow Up: A Systematic Review and Meta-Analysis of Qualitative and Quantitative Studies. *PLoS ONE*. <https://doi.org/10.1371/journal.pone.0084238>
37. Kaiser, J. 2015. Obama gives East Room roll-out to Precision Medicine Initiative. *Science Magazine*. Available at: <https://doi.org/10.1126/science.aaa6436> (accessed November 5, 2020).
38. Precision Medicine Initiative Working Group. 2015. *The Precision Medicine Initiative Cohort Program – Building a Research Foundation for 21st Century Medicine*. Available at: <https://acd.od.nih.gov/documents/reports/DRAFT-PMI-WG-Report-9-11-2015-508.pdf> (accessed November 5, 2020).
39. Mohr, D. C., M. Zhang, and S. M. Schueller. 2017. Personal Sensing: Understanding Mental Health Using Ubiquitous Sensors and Machine Learning. *Annual Review of Clinical Psychology* 13: 23-47. <https://doi.org/10.1146/annurev-clinpsy-032816-044949>
40. Epstein, D. H., M. Tyburski, W. J. Kowalczyk, A. J. Burgess-Hull, K. A. Phillips, B. L. Curtis, and K. L. Preston. 2020. *npj Digital Medicine* 3(26). <https://doi.org/10.1038/s41746-020-0234-6>

41. Barnett, I., J. Torous, P. Staples, L. Sandoval, M. Keshavan, and J. Pekka Onnela. 2018. Relapse prediction in schizophrenia through digital phenotyping: a pilot study. *Neuropsychopharmacology* 43: 1660-1666. <https://doi.org/10.1038/s41386-018-0030-z>
42. de Choudhury, M., S. Counts, and E. Horvitz. 2013. Social media as a measurement tool of depression in populations. *WebSci '13: Proceedings of the 5th Annual ACM Web Science Conference*. <https://doi.org/10.1145/2464464.2464480>
43. Nahum-Shani, I., S. N. Smith, B. J. Spring, L. M. Collins, K. Witkiewitz, A. Tewari, and S. A. Murphy. 2017. Just-in-Time Adaptive Interventions (JITAs) in Mobile Health: Key Components and Design Principles for Ongoing Health Behavior Support. *Annals of Behavioral Medicine* 52(6): 446-462. <https://doi.org/10.1007/s12160-016-9830-8>
44. Gustafson, D. H., F. M. McTavish, M. Chih, A. K. Atwood, R. A. Johnson, M. G. Boyle, M. S. Levy, H. Driscoll, S. M. Chisholm, L. Dillenburg, A. Isham, and D. Shah. 2014. A Smart-phone Application to Support Recovery from Alcoholism: A Randomized Clinical Trial. *JAMA Psychiatry* 71(5): 566-572. <https://doi.org/10.1001/jamapsychiatry.2013.4642>
45. Gould, C. E., B. C. Kok, V. K. Ma, A. M. L. Zapata, J. E. Owen, and E. Kuhn. 2019. Veterans Affairs and the Department of Defense men- tal health apps: A systematic literature re- view. *Psychological Services* 16(2): 196–207. <https://doi.org/10.1037/ser0000289>
46. Lucas, G. M., J. Gratch, A. King, and L. Morency. 2014. It's only a computer: Virtual humans increase willingness to disclose. *Computers in Human Behavior* 37: 94-100. <https://doi.org/10.1016/j.chb.2014.04.043>
47. Lucas, G. M., A. Rizzo, J. Gratch, S. Scherer, G. Stratou, J. Boberg, and L. Morency. 2017. Reporting Mental Health Symptoms: Breaking Down Barriers to Care with Virtual Human Interviewers. *Frontiers in Robotics and AI*. <https://doi.org/10.3389/frobt.2017.00051>
48. Kornfield, R., P. K. Sarma, D. V. Shah, F. McTavish, G. Landucci, K. Peromashko, and D. H. Gustafson. 2018. Detecting Recovery Problems Just in Time: Application of Automated Linguistic Analysis and Supervised Machine Learning to an Online Substance Abuse Forum. *Journal of Medical Internet Research* 20(6):e10136. <https://doi.org/10.2196/10136>
49. Shahzad, M., R. Upshur, P. Donnelly, A. Bharmal, X. Wei, P. Feng, and A. D. Brown. 2019. A population-based approach to integrated healthcare delivery: a scoping review of clinical care and public health collaboration. *BMC Public Health* 19(708). <https://doi.org/10.1186/s12889-019-7002-z>

50. Baig, M. M., H. GholamHosseini, A. A. Moqem, F. Mirza, and M. Linden. 2017. A Systematic Review of Wearable Patient Monitoring Systems – Current Challenges and Opportunities for Clinical Adoption. *Journal of Medical Systems* 41(115). <https://doi.org/10.1007/s10916-017-0760-1>
51. Paradiso, R., C. Belloc, G. Loriga, and N. Taccini. 2005. Wearable healthcare systems, new frontiers of e-textile. *Studies in Health Technology and Informatics* 117: 9-16. Available at: <https://pubmed.ncbi.nlm.nih.gov/16282648/> (accessed November 6, 2020).
52. Thomas, S. S., V. Nathan, C. Zong, K. Sounda- rapandian, X. Shi, and R. Jafari. 2016. Bio- Watch: A Noninvasive Wrist-Based Blood Pressure Monitor That Incorporates Training Techniques for Posture and Subject Variability. *IEEE Journal of Biomedical and Health Informatics* 20(5). <https://doi.org/10.1109/JBHI.2015.2458779>
53. Melillo, P., A. Orrico, P. Scala, F. Crispino, and L. Pecchia. 2015. Cloud-Based Smart Health Monitoring System for Automatic Cardio- vascular and Fall Risk Assessment in Hypertensive Patients. *Journal of Medical Systems* 39(109). <https://doi.org/10.1007/s10916-015-0294-3>
54. Sapci, A. H., and H. A. Sapci. 2019. Innovative Assisted Living Tools, Remote Monitoring Technologies, Artificial Intelligence-Driven Solutions, and Robotic Systems for Aging Societies: Systematic Review. *JMIR Aging* 2(2): e5429. <https://doi.org/10.2196/15429>
55. Martin, L. R., S. L. Williams, K. B. Haskard, and M. R. Dimatteo. 2005. The challenge of patient adherence. *Therapeutics and Clinical Risk Management* 1(3): 189-199. Available at: <https://pubmed.ncbi.nlm.nih.gov/18360559/> (accessed November 5, 2020).
56. Yeam, C. T., S. Chia, H. C. C. Tan, Y. H. Kwan, W. Fong, and J. J. B. Seng. 2018. A systematic review of factors affecting medication adherence among patients with osteoporosis. *Osteoporosis International* 29: 2623-2637. <https://doi.org/10.1007/s00198-018-4759-3>
57. Restrepo, R. D., M. T. Alvarez, L. D. Wittnebel, H. Sorenson, R. Wettstein, D. L. Vines, J. Sikkema-Ortiz, D. D. Gardner, and R. L. Wilkins. 2008. Medication adherence issues in patients treated for COPD. *International Journal of Chronic Obstructive Pulmonary Disease* 3(3): 371-384. <https://doi.org/10.2147/copd.s3036>
58. Engelkes, M., H. M. Janssens, J. C. de Jongste, M. C. J. M. Sturkenboo, and K. M. C. Verhamme. 2015. Medication adherence and the risk of severe asthma exacerbations: a systematic review. *European Respiratory Journal* 45: 396-407. <https://doi.org/10.1183/09031936.00075614>
59. Grant, R. W., N. G. Devita, D. E. Singer, and J. B. Meigs. 2003. Polypharmacy and Medication Adherence in Patients with Type 2 Diabetes. *Diabetes Care* 26(5): 1408-1412. <https://doi.org/10.2337/diacare.26.5.1408>

60. Wu, J-R., D. K. Moser, T. A. Lennie, P. V. Burkhart. 2008. Medication Adherence in Patients Who Have Heart Failure: A Review of the Literature. *Nursing Clinics of North America* 43(1): 133-153. <https://doi.org/10.1016/j.cnur.2007.10.006>
61. Labovitz, D. L., L. Shafner, M. R. Gil, D. Virmani, and A. Hanina. 2017. Using Artificial Intelligence to Reduce the Risk of Nonadherence in Patients on Anticoagulation Therapy. *Stroke* 48: 1416-1419. <https://doi.org/10.1161/strokeaha.116.016281>
62. Kim, K., B. Kim, A. J. Chung, K. Kwon, E. Choi, and J. Nah. 2018. Algorithm and System for improving the medication adherence of tuberculosis patients. *2018 International Conference on Information and Communication Technology Convergence (ICTC)*. <https://doi.org/10.1109/ICTC.2018.8539402>
63. Bain, E. E., L. Shafner, D. P. Walling, A. A. Othman, C. Chuang-Stein, J. Hinkle, and A. Hanina. 2017. Use of a Novel Artificial Intelligence Platform on Mobile Devices to Assess Dosing Compliance in a Phase 2 Clinical Trial in Subjects With Schizophrenia. *JMIR Mhealth Uhealth* 5(2): e18. <https://doi.org/10.2196/mhealth.7030>
64. Lake, I. R., F. J. Colon-Gonzalez, G. C. Barker, R. A. Morbey, G. E. Smith and A. J. Elliot. 2019. Machine learning to refine decision making within a syndromic surveillance service. *BMC Public Health* 19(559). <https://doi.org/10.1186/s12889-019-6916-9>
65. Research Triangle Environmental Health Collaborative. 2018. *11th Annual Environmental Health Summit: Artificial Intelligence in Environmental Health Science and Decision Making*. Available at: <https://environmentalhealthcollaborative.org/wp-content/uploads/2019/07/RTEHC-2018-Summit-Report.pdf> (accessed November 4, 2020).
66. Spinu, N., M. T. D. Cronin, S. J. Enoch, J. C. Madden, and A. P. Worth. 2020. Quantitative adverse outcome pathway (qAOP) models for toxicity prediction. *Archives of Toxicology* 94: 1497-1510. <https://doi.org/10.1007/s00204-020-02774-7>
67. National Academies of Sciences, Engineering, and Medicine. 2019. *Leveraging Artificial Intelligence and Machine Learning to Advance Environmental Health Research and Decisions: Proceedings of a Workshop—in Brief*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/25520>
68. Weichenthal, S., M. Hatzopoulou, and M. Brauer. 2019. A picture tells a thousand... exposures: Opportunities and challenges of deep learning image analyses in exposure science and environmental epidemiology. *Environment International* 122: 3-10. <https://doi.org/10.1016/j.envint.2018.11.042>

69. VoPham, T., J. E. Hart, F. Laden, and Y. Chiang. 2018. Emerging trends in geospatial artificial intelligence (geoAI): potential applications for environmental epidemiology. *Environmental Health* 17(40). <https://doi.org/10.1186/s12940-018-0386-x>
70. Bellinger, C., M. S. M. Jabbar, O. Zaiane, and A. Osornio-Vargas. 2017. A systematic review of data mining and machine learning for air pollution epidemiology. *BMC Public Health* 17(907). <https://doi.org/10.1186/s12889-017-4914-3>
71. Apte, J. S., K. P. Messier, S. Gani, M. Brauer, T. W. Kirchstetter, M. M. Lunden, J. D. Marshall, C. J. Portier, R. C. H. Vermeulen, and S. P. Hamburg. 2017. High-Resolution Air Pollution Mapping with Google Street View Cars: Exploiting Big Data. *Environmental Science and Technology* 51(12): 6999-7008. <https://doi.org/10.1021/acs.est.7b00891>
72. Considine, E. 2020. *Calibrating Low-Cost Air Pollution Sensors Using Machine Learning*. Earth Lab, University of Colorado, Boulder. Available at: <https://www.colorado.edu/earthlab/2020/02/07/calibrating-low-cost-air-pollution-sensors-using-machine-learning> (accessed November 5, 2020).
73. Maharana, A. and E. O. Nsoesie. 2018. Use of Deep Learning to Examine the Association of the Built Environment with Prevalence of Neighborhood Adult Obesity. *JAMA Network Open* 1(4): e181535. <https://doi.org/10.1001/jamanetworkopen.2018.1535>
74. Fliesler, N. 2020. *Crowdsourcing the COVID-19 pandemic in real time*. Boston Children's Hospital. Available at: <https://discoveries.childrenshospital.org/covid-near-you-coronavirus-tracking/> (accessed November 5, 2020).
75. Bluedot. 2020. *Home*. Available at: <https://bluedot.global/> (accessed November 5, 2020).
76. Feldman, J., A. Thomas-Bachli, J. Forsyth, Z. H. Patel, and K. Khan. 2019. Development of a global infectious disease activity database using natural language processing, machine learning, and human expertise. *Journal of the American Medical Informatics Association* 26(11): 1355-1359. <https://doi.org/10.1093/jamia/ocz112>
77. Freifeld, C. C., K. D. Mandl, B. Y. Reis, and J. S. Brownstein. 2008. HealthMap: Global Infectious Disease Monitoring through Automated Classification and Visualization of Internet Media Reports. *Journal of the American Medical Informatics Association* 15(2): 150-157. <https://doi.org/10.1197/jamia.M2544>
78. Conway, M., M. Hu, and W. W. Chapman. 2019. Recent Advances in Using Natural Language Processing to Address Public Health Research Questions Using Social Media and Consumer Generated Data. *Yearbook of Medical Informatics* 28(01): 208-217. <https://doi.org/10.1055/s-0039-1677918>

79. Prosser, M. 2020. How AI Helped Predict the Coronavirus Outbreak Before It Happened. *SingularityHub*, February 5, 2020. Available at: <https://singularityhub.com/2020/02/05/how-ai-helped-predict-the-coronavirus-outbreak-before-it-happened/> (accessed November 5, 2020).
80. Quilty, B. J., S. Clifford, CMMID nCOV working group, S. Flasche, and R. M. Eggo. 2020. Effectiveness of airport screening at detecting travellers infected with novel coronavirus (2019-nCoV). *Eurosurveillance* 25(5): pii=20000080. <https://doi.org/10.2807/1560-7917.ES.2020.25.5.2000080>
81. Radin, J. M., N. E. Wineinger, E. J. Topol, and S. R. Steinhubl. 2020. Harnessing wearable device data to improve state-level real-time surveillance of influenza-like illness in the USA: a population-based study. *The Lancet Digital Health* 2(2): e85-e93. [https://doi.org/10.1016/S2589-7500\(19\)30222-5](https://doi.org/10.1016/S2589-7500(19)30222-5)
82. Wong, C. K., D. T. Y. Ho, A. R. Tam, M. Zhou, Y. M. Lau, M. O. Y. Tang, R. C. F. Tong, K. S. Rajput, G. Chen, S. C. Chan, C. W. Siu, and I. F. N. Hung. 2020. Artificial intelligence mobile health platform for early detection of COV- ID-19 in quarantine subjects using a wearable biosensor: protocol for a randomised controlled trial. *BMJ Open* 10(7). <http://dx.doi.org/10.1136/bmjopen-2020-038555>
83. Pickman, B. 2020. The Story Behind the Ring That is Key to the NBA's Restart. *Sports Illustrated*, July 1, 2020. Available at: <https://www.si.com/nba/2020/07/01/orlando-ring-nba-restart-orlando-coronavirus> (accessed November 5, 2020).
84. HL7 International. 2020. *Introduction to HL7 Standards*. Available at: <https://www.hl7.org/implement/standards/> (accessed November 5, 2020).
85. HL7 FHIR. 2020. *Introduction*. Available at: <https://www.hl7.org/fhir/overview.html> (accessed November 5, 2020).
86. Kahn, M. G., T. J. Callahan, J. Barnard, A. E. Bauck, J. Brown, B. N. Davidson, H. Estiri, C. Goerg, E. Holve, S. G. Johnson, S-T. Liaw, M. Hamilton-Lopez, D. Meeker, T. C. Ong, P. Ryan, N. Shang, N. G. Weiskopf, C. Weng, M. N. Zozus, and L. Schilling. 2016. A Harmonized Data Quality Assessment Terminology and Framework for the Secondary Use of Electronic Health Record Data. *eGEMs (Generating Evidence & Methods to improve patient outcomes)* 4(1):18. <http://doi.org/10.13063/2327-9214.1244>
87. Observational Health Data Sciences and Informatics. 2020. *Data Standardization*. Available at: <https://ohdsi.org/data-standardization/> (accessed November 5, 2020).
88. U.S. Department of Health and Human Services. n.d. *Public Health Authority, § 164.502*. Available at: <https://www.govinfo.gov/content/pkg/CFR-2010-title45-vol1/pdf/CFR-2010-title45-vol1-sec164-502.pdf> (accessed November 5, 2020).

89. HHS.gov. n.d. *Does the HIPAA Privacy Rule permit doctors, nurses, and other health care providers to share patient health information for treatment purposes without the patient's authorization?* Available at: <https://www.hhs.gov/hipaa/for-professionals/faq/481/does-hipaa-permit-doctors-to-share-patient-information-for-treatment-without-authorization/index.html> (accessed November 5, 2020).
90. U.S. Department of Health and Human Services. n.d. *45 CFR Subtitle A (10-1-09 Edition)*. Available at: <https://www.govinfo.gov/content/pkg/CFR-2009-title45-vol1/pdf/CFR-2009-title45-vol1-sec164-512.pdf> (accessed November 5, 2020).
91. Clayton, E. W., B. J. Evans, J. W. Hazel, and M.A. Rothstein. 2019. The law of genetic privacy: applications, implications, and limitations. *Journal of Law and the Biosciences* 6(1): 1-36. <https://doi.org/10.1093/jlb/lisz007>
92. National Committee on Vital and Health Statistics. 2016. *Letter on recommendation on the HIPAA Minimum Necessary Standard*. Available at: <https://ncvhs.hhs.gov/wp-content/uploads/2018/03/2016-Ltr-Privacy-Minimum-Necessary-formatted-on-ltrhead-Nov-9-FINAL-w-sig.pdf> (accessed November 5, 2020).
93. PhysioNet. n.d. *Overview*. Available at: <https://physionet.org/about/database/> (accessed November 5, 2020).
94. MIMIC. n.d. *Home*. Available at: <https://mimic.physionet.org/> (accessed November 5, 2020).
95. Obermeyer, Z., B. Powers, C. Vogeli, and S. Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366(6464): 447-453. <https://doi.org/10.1126/science.aax2342>
96. Vogels, E. A. 2020. About one-in-five Americans use a smart watch or fitness tracker. *Pew Research Center*, January 9, 2020. Available at: <https://www.pewresearch.org/fact-tank/2020/01/09/about-one-in-five-americans-use-a-smart-watch-or-fitness-tracker/> (accessed November 5, 2020).
97. Jung, J., S. Corbett-Davies, R. Shroff, and S. Goel. 2019. Omitted and Included Variable Bias in Tests for Disparate Impact. *arXiv.org*. Available at: <https://arxiv.org/abs/1809.05651> (accessed November 5, 2020).
98. Yadlowsky, S., R. A. Hayward, J. B. Sussman, R. L. McClelland, Y-I. Min, and S. Basu. 2018. Clinical Implications of Revised Pooled Cohort Equations for Estimating Atherosclerotic Cardiovascular Disease Risk. *Annals of Internal Medicine*. <https://doi.org/10.7326/M17-3011>
99. Cahan, E. M., T. Hernandez-Boussard, S. Thadaney-Israni, and D. L. Rubin. 2019. Putting the data before the algorithm in big data addressing personalized healthcare. *npj Digital Medicine* 2(78). <https://doi.org/10.1038/s41746-019-0157-2>

100. Basu, S., R. S. Phillips, R. Phillips, L. E. Peterson, and B. E. Landon. 2020. Primary Care Practice Finances In The United States Amid The COVID-19 Pandemic. *Health Affairs* 39(9). <https://doi.org/10.1377/hlthaff.2020.00794>
101. Abramoff, M. D. n.d. *Autonomous AI and Ethics: lessons from real world implementation*. Presentation. Available at: <https://www.fda.gov/media/135730/download> (accessed November 5, 2020).
102. Baicker, K., and M. E. Chernew. 2017. Alternative Alternative Payment Models. *JAMA Internal Medicine* 177(2): 222-223. <https://doi.org/10.1001/jamainternmed.2016.8280>
103. Mechanic, R. E., P. Santos, B. E. Landon, and M. E. Chernew. 2011. Medical Group Responses To Global Payment: Early Lessons From The 'Alternative Quality Contract' In Massachusetts. *Health Affairs* 30(9). <https://doi.org/10.1377/hlthaff.2011.0264>
104. Contreras, I. and J. Vehi. 2018. Artificial Intelligence for Diabetes Management and Decision Support: Literature Review. *Journal of Medical Internet Research* 20(5): e10775. <https://doi.org/10.2196/10775>
105. Dankwa-Mullan, I., M. Rivo, M. Sepulveda, Y. Park, J. Snowdon, and K. Rhee. 2019. Transforming Diabetes Care Through Artificial Intelligence: The Future Is Here. *Population Health Management* 22(3). <https://doi.org/10.1089/pop.2018.0129>
106. Barrett, M., J. Boyne, J. Brandts, H-P. B. La Rocca, L. De Maesschalck, K. De Wit, L. Dixon, C. Eurlings, D. Fitzsimons, O. Golubnitschaja, A. Hageman, F. Heemskerk, A. Hintzen, T. M. Helms, L. Hill, T. Hoedemakers, N. Marx, K. McDonald, M. Mertens, D. Müller-Wieland, A. Palant, J. Piesk, A. Pomazanskyi, J. Ramaekers, P. Ruff, K. Schütt, Y. Shekhawat, C. F. Ski, D.R. Thompson, A. Tsirkin, K. van der Mierden, Watson and B. Zippel-Schultz. 2019. Artificial intelligence supported patient self-care in chronic heart failure: a paradigm shift from reactive to predictive, preventive and personalised care. *EPMA Journal* 10: 445-464. <https://doi.org/10.1007/s13167-019-00188-9>
107. Blackmore, C. C., B. L. Williams, J. M. Ching, L. A. Chafetz, and G. S. Kaplan. 2016. Using Lean to Advance Quality Improvement Research. *Journal for Healthcare Quality* 38(5): 275-282. <https://doi.org/10.1097/01.JHQ.0000462684.78253.a1>
108. Berwick, D. 2003. Disseminating Innovations in Health Care. *JAMA* 289(15): 1969-1975. <https://doi.org/10.1001/jama.289.15.1969>
109. Hargraves, I. G., A. K. Fournier, V. M. Montori, and A. S. Bierman. 2020. Generalized shared decision making approaches and patient problems. Adapting AHRQ's SHARE Approach for Purposeful SDM. *Patient Education and Counseling* 103(10): 2192-2199. <https://doi.org/10.1016/j.pec.2020.06.022>
110. Eubanks, V. 2018. *Automating Inequality: How High-Tech tools Profile, Police, and Punish the Poor*. St. Martin's Press.

111. O’Neil, Cathy. 2017. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown.
112. Christensen, C. M., S. Cook, and T. Hall. 2005. Marketing Malpractice: The Cause and the Cure. *Harvard Business Review*. Available at: <https://hbr.org/2005/12/marketing-malpractice-the-cause-and-the-cure> (accessed November 5, 2020).
113. Christensen, C. M., T. Hall, K. Dillon, and S. Duncan. 2016. Know Your Customers’ “Jobs to be Done”. *Harvard Business Review*. Available at: <https://hbr.org/2016/09/know-your-customers-jobs-to-be-done> (accessed November 5, 2020).
114. Sendak, M., M. Gao, M. Nichols, A. Lin, and Balu. 2019. Machine Learning in Health Care: A Critical Appraisal of Challenges and Opportunities. *eGEMs (Generating Evidence & Methods to improve patient outcomes)* 7(1):1. <http://doi.org/10.5334/egems.287>
115. Beaulieu-Jones, B., S. G. Finlayson, C. Chivers, I. Chen, M. McDermott, J. Kandola, A.V. Dalca, A. Beam, M. Fiterau, T. Naumann. 2019. Trends and Focus of Machine Learning Applications for Health Research. *JAMA Network Open* 2(10): e1914051. <https://doi.org/10.1001/jamanetworkopen.2019.14051>
116. Topol, E. J. 2019. High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine* 25: 44-56. <https://doi.org/10.1038/s41591-018-0300-7>
117. Sculley, D., G. Holt, D. Golovin, E. Davydov, Phillips, D. Ebner, V. Chaudhary, and M. Young. 2014. Machine Learning: The High Interest Credit Card of Technical Debt. *SE4ML: Software Engineering for Machine Learning (NIPS 2014 Workshop)*. Available at: <https://research.google/pubs/pub43146/> (accessed November 5, 2020).
118. Wiens, J., S. Saria, M. Sendak, M. Ghassemi, V. X. Liu, F. Doshi-Velez, K. Jung, K. Heller, D. Kale, M. Saeed, P. N. Ossorio, S. Thadaneys-Israni, and A. Goldenberg. 2019. Do no harm: a roadmap for responsible machine learning for health care. *Nature Medicine* 25: 1337- 1340. <https://doi.org/10.1038/s41591-019-0548-6>
119. Sendak, M., M. C. Elish, M. Gao, J. Futoma, W. Ratliff, M. Nichols, A. Bedoya, S. Balu, and C. O’Brien. 2002. “The human body is a black box”: supporting clinical decision-making with deep learning. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*: 99–109. <https://doi.org/10.1145/3351095.3372827>
120. Sendak, M. P., J. D’Arcy, S. Kashyap, M. Gao, M. Nichols, K. Corey, W. Ratliff, and S. Balu. 2019. A Path for Translation of Machine Learning Products into Healthcare Delivery. *EMJ*. Available at: <https://www.emjreviews.com/innovations/article/a-path-for-translation-of-machine-learning-products-into-healthcare-delivery/> (accessed November 5, 2020).

121. Oh, J., M. Makar, C. Fusco, and R. McCaffery. 2018. A Generalizable, Data-Driven Approach to Predict Daily Risk of Clostridium difficile Infection at Two Large Academic Health Centers. *Infection Control & Hospital Epidemiology*: 425-433. <https://doi.org/10.1017/ice.2018.16>
122. Sendak, M. P., S. Balu, and K. A. Schulman. 2017. Barriers to Achieving Economies of Scale in Analysis of HER Data. *Applied Clinical Informatics* 08(03): 826-831. <https://doi.org/10.4338/ACI-2017-03-CR-0046>
123. Spanou, D. 2013. Software as a Medical Device (SaMD): Key Definitions. *IMDRF SaMD Working Group*. Available at: <http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-131209-samd-key-definitions-140901.pdf> (accessed November 5, 2020).
124. U.S. Food and Drug Administration. 2020. *Artificial Intelligence and Machine Learning in Software as a Medical Device*. Available at: <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device> (accessed November 5, 2020).
125. U.S. Congress. 2016. *H.R. 34 – 21st Century Cures Act*. Available at: <https://www.congress.gov/bill/114th-congress/house-bill/34> (accessed November 5, 2020).
126. U. S. Food and Drug Administration. 2018. *Digital Health Innovation Action Plan*. Available at: <https://www.fda.gov/media/106331/download> (accessed November 5, 2020).
127. U.S. Food and Drug Administration. 2020. *Digital Health Software Precertification (Pre-Cert) Program*. Available at: <https://www.fda.gov/medical-devices/digital-health-center-excellence/digital-health-software-precertification-pre-cert-program> (accessed November 5, 2020).
128. U. S. Food and Drug Administration. 2019. *General Wellness: Policy for Low Risk Devices*. Available at: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/general-wellness-policy-low-risk-devices> (accessed November 5, 2020).
129. U. S. Food and Drug Administration. 2019. *Policy for Device Software Functions and Mobile Medical Applications*. Available at: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/policy-device-software-functions-and-mobile-medical-applications> (accessed November 5, 2020).
130. U. S. Food and Drug Administration. 2019. *Clinical Decision Support Software*. Available at: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software> (accessed November 5, 2020).
131. United States Code. 2006. *United States Code, 2006 Edition, Supplement 4, Title 21 FOOD AND DRUGS*. Available at: <https://www.govinfo.gov/app/details/USCODE-2010-title21/USCODE-2010-title21-chap9-subchapII-sec321> (accessed November 5, 2020).

132. Evans, B. J. 2020. Minding the Gaps in Regulation of Do-it-Yourself Biotechnology. *DePaul Journal of Health Care Law* 21(3). Available at: <https://via.library.depaul.edu/cgi/view-content.cgi?article=1380&context=jhcl> (accessed November 5, 2020).
133. Price, W. N., S. Gerke, D-J. Univ, and I. G. Cohen. 2019. Potential Liability for Physicians Using Artificial Intelligence. *JAMA* 322(18): 1765-1766. <https://doi.org/10.1001/jama.2019.15064>
134. Zuboff, S. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. Public Affairs Books.
135. El Emam, K., E. Jonker, L. Arbuckle, and Malin. 2011. A Systematic Review of Re-Identification Attacks on Health Data. *PLOS One*. <https://doi.org/10.1371/journal.pone.0028071>
136. Barth-Jones, D. 2012. The ‘Re-Identification’ of Governor William Weld’s Medical Information: A Critical Re-Examination of Health Data Identification Risks and Privacy Protections, Then and Now. *SSRN*. <http://dx.doi.org/10.2139/ssrn.2076397>
137. Benitez, K., and B. Malin. 2010. Evaluating re-identification risks with respect to the HIPAA privacy rule. *Journal of the American Medical Informatics Association* 17(2): 169-177. <https://doi.org/10.1136/jamia.2009.000026>
138. Barth-Jones, D. 2012. The Debate Over ‘Re-Identification’ of Health Information: What Do We Risk? *Health Affairs*. <https://doi.org/10.1377/hblog20120810.021952>
139. Stat News. 2020. An Invisible Hand: Patients Aren’t Being Told About the AI Systems Advising Their Care. *Health Leaders Media*, July 15, 2020. Available at: <https://www.healthleadersmedia.com/innovation/invisible-hand-patients-arent-being-told-about-ai-systems-advising-their-care> (accessed November 5, 2020).
140. Cohen, I. G. 2020. Informed Consent and Medical Artificial Intelligence: What to Tell the Patient? *Georgetown Law Journal* 108: 1425-1469. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3529576 (accessed November 5, 2020).
141. Lin, L., J. Conroy, and P. Chrisidis. 2020. Data- point: Which states have the most licensed psychologists? *American Psychological Association* 51(1): 19. Available at: <https://www.apa.org/monitor/2020/01/datapoint-states> (accessed November 5, 2020).
142. Rowan, K., D. D. Mcalpine, and L. A. Blewett. 2013. Access and Cost Barriers to Mental Health Care, By Insurance Status, 1999-2010. *Health Affairs* 32(10). <https://doi.org/10.1377/hlthaff.2013.0133>
143. Pew Research Center. 2019. *Mobile Fact Sheet*. Available at: <https://www.pewresearch.org/internet>

[/fact-sheet/mobile/](#) (accessed November 5, 2020).

144. Perrin, A., and E. Turner. 2019. Smartphones help blacks, Hispanics bridge some but not all digital gaps with whites. *Pew Research Center*. Available at: <https://www.pewresearch.org/fact-tank/2019/08/20/smartphones-help-blacks-hispanics-bridge-some-but-not-all-digital-gaps-with-whites/> (accessed November 5, 2020).
145. The Chartis Group. n.d. *Telehealth Adoption Tracker*. Available at: https://reports.chartis.com/telehealth_trends_and_implications-aug2020/ (accessed November 5, 2020).
146. Brewer, L. C., K. L. Fortuna, C. Jones, R. Walker, S. N. Hayes, C. A. Patten, and L. A. Cooper. 2020. Back to the Future: Achieving Health Equity Through Health Informatics and Digital Health. *JMIR Mhealth Uhealth* 8(1): e14512. <https://doi.org/10.2196/14512>
147. Stoytcheva, S., D. Cohen, and C. Blake. 2015. Exploring cultural differences in language usage: The case of negation. *Proceedings of the American Society for Information Science and Technology*. Available at: <https://doi.org/10.1002/meet.2014.14505101086> (accessed November 5, 2020).

DOI

<https://doi.org/10.31478/202011f>

Suggested Citation

Aggarwal, N., M. Ahmed, S. Basu, J. J. Curtin, B. J. Evans, M. E. Matheny, S. Nundy, C. Shachar, M. P. Sendak, R. U. Shah, and S. Thadanev-Israni. 2020. Advancing Artificial Intelligence in Health Settings Outside the Hospital and Clinic. *NAM Perspectives*. Discussion Paper, National Academy of Medicine, Washington, DC.

<https://doi.org/10.31478/202011f>

Author Information

Nakul Aggarwal, BS, is an MD-PhD Candidate and NIMH T32 Predoctoral Fellow at University of Wisconsin-Madison. **Mahnoor Ahmed, MEng**, is an Associate Program Officer at the National Academy of Medicine. **Sanjay Basu, MD, PhD**, is Director of Research at the Center for Primary Care at Harvard University. **John J. Curtin, PhD**, is Professor and Director of Clinical Training at University of Wisconsin-Madison. **Barbara J. Evans, JD, PhD**, is Professor of Law, Professor of Engineering, and Stephen O'Connell Chair at University of Florida. **Michael E. Matheny, MD, MS, MPH**, is Co-Director of the Center for Improving the Public's Health through Informatics; Associate Director of the Advanced Fellowship in Medical Informatics at TVHS Veterans Affairs; and Associate Professor in the Departments of Biomedical Informatics, Medicine, and Biostatistics at Vanderbilt University. **Shantanu Nundy, MD, MBA**, is Chief Medical Officer at Accolade Inc. and Professorial Lecturer at

George Washington University Milken Institute for Public Health. **Carmel Shachar, JD, MPH**, is Executive Director of the Petrie-Floem Center for Health Law Policy, Biotechnology, and Bioethics at Harvard University. **Mark P. Sendak, MD, MPP**, is Population Health and Data Science Lead at Duke University. **Rashmee U. Shah, MD, MS**, is Assistant Professor in Cardiovascular Medicine at University of Utah. **Sonoo Thadanev-Israni, MBA**, is Executive Director of the Presence Center at Stanford University.

Nakul Aggarwal is supported by grants from the NIH (T32MH018931; T32GM008692).

John Curtin is supported by grants from the NIH (R01DA047315; R01AA024391).

Rashmee Shah is supported by a grant from the NIH/NHLBI (5R03HL148372).

Acknowledgments

The authors would like to thank **Pat Baird**, Philips; **Sally Okun**, SallyOkun360; and **W. Nicholson Price II**, University of Michigan, for their valuable contributions to this paper.

Conflict-of-Interest Disclosures

Mark Sendak discloses that his employer holds a license agreement with Cohere Med, Inc. **Sanjay Basu** discloses that he receives grants from the National Institutes of Health and the Centers for Disease Control and Prevention; that he receives personal fees from the New England Journal of Medicine, PLoS Medicine, Collective Health, and HealthRight360; and that he

has a pending patent relative to this work. **Shantanu Nundy** discloses that he is employed by Accolade, Inc., which provides personalized advocacy and population health services.

Correspondence

Questions or comments about this paper should be sent to Mahnoor Ahmed at mahmed@nas.edu.

Disclaimer

The views expressed in this paper are those of the author and not necessarily of the author's organizations, the National Academy of Medicine (NAM), or the National Academies of Sciences, Engineering, and Medicine (the National Academies). The paper is intended to help inform and stimulate discussion. It is not a report of the NAM or the National Academies. Copyright by the National Academy of Sciences. All rights reserved.

Appendix I: Objectives, Scope, and Methodology

We describe our scope and methodology for addressing the three objectives outlined below:

Objectives

1. What current and emerging artificial intelligence (AI) tools are available for augmenting patient care, and what are their benefits?
2. What challenges surround the use of AI tools to augment patient care?
3. What policy options could help maximize benefits and mitigate challenges surrounding the use of AI tools to augment patient care?

Scope and methodology

To address all three research objectives, we assessed available and developing AI tools that companies or health care providers could use to augment patient care and for administrative purposes as well as the challenges associated with using such tools. To do so, we reviewed key reports and scientific literature describing current and developing tools; interviewed a variety of stakeholders, including agency officials, industry members, academic researchers, and a consumer group; and conducted an expert meeting in conjunction with the National Academy of Medicine.

Limitations to scope

We focused our review on selected tools used at locations that employ health care providers, including but not limited to physicians, registered nurses, medical

assistants, and physical therapists. We excluded tools used in other environments, such as the home or office. Tools discussed are examples and not an exhaustive list of all AI tools used in augmenting patient care or for administrative purposes. We did not assess all available or developing tools. We selected narrative examples to demonstrate the breadth of AI tools in augmenting patient care or for administrative purposes.

Literature Search

In the course of our work we conducted four literature searches. To establish background and identify appropriate tools and their challenges, we reviewed key articles from the scientific literature. To support objective 2, we conducted two separate literature searches, one focused on economic considerations and the other on patient privacy issues. To support objective 3, we conducted a policy options literature search using a variety of databases, including Scopus, MEDLINE, Biosis Previews, and WorldCat. We used search terms such as “health care services”, “medical services”, “delivery” and “artificial intelligence”, and narrowed our search to articles published within the last five years. For these searches, results could originate from scholarly or peer reviewed material, government reports, conference papers, dissertations, working papers, books, legislative materials, trade or industry articles, and white papers, but not from general news. We selected the most relevant articles for further review based on our objectives, and reviewed the abstracts for additional search terms to refine the results.

Interviews

We interviewed key stakeholders in the field of AI to augment patient care, including:

- relevant federal agencies including the Department of Veterans Affairs, the Centers for Medicare and Medicaid Services, the National Institutes of Health, and the Food and Drug Administration;
- four industry/professional organizations and one private firm;
- three academic researchers or institutions;
- one health care system and;
- one consumer group.

Because this is a small and non-generalizable sample of the stakeholders involved in using AI to augment patient care, the results of our interviews are illustrative and represent important perspectives, but are not generalizable.

Expert Meeting

We collaborated with the National Academy of Medicine to convene a 2-day meeting of 18 experts on current and emerging AI tools for use in augmenting patient care. We worked with National Academy of Medicine staff to identify experts from a range of stakeholder groups including federal agencies, academia,

industry, and legal scholars, with expertise covering all significant areas of our review, including individuals with research or operational expertise in using AI tools in the augmentation of patient care.⁸¹ We evaluated the experts for any conflicts of interest. A conflict of interest was considered to be any current financial or other interest (such as an organizational position) that might conflict with the service of an individual because it could (1) impair objectivity or (2) create an unfair competitive advantage for any person or organization. The 18 experts were determined to be free of reported conflicts of interest, except those that were outside the scope of the forum or where the overall design of our panel and methodology was sufficient to address them, and the group as a whole was determined to not have any inappropriate biases.⁸² (See app. II for a list of these experts and their affiliations.) The comments of these experts generally represented the views of the experts themselves and not the agency, university, or company with which they were affiliated, and are not generalizable to the views of others in the field.

We divided the 2-day meeting into six moderated discussion sessions: (1) tools to assist with augmenting patient care using clinical decision support tools and administrative tools; (2) tools to assist with augmenting patient care using robotics; (3)

⁸¹This meeting of experts was planned and convened with the assistance of the National Academy of Medicine to better ensure that a breadth of expertise was brought to bear in its preparation, however all final decisions regarding meeting substance and expert participation was the responsibility of GAO. Any conclusions and recommendations in GAO reports are solely those of the GAO.

⁸²For example, one expert had equity interest in a company developing an AI technology for use in augmenting patient

care. We determined the expert's relationship did not prevent the expert from serving on the panel as the discussion was not planned to revolve around any specific technology, pharmaceutical company, or vested interest. We did not interview representatives from or otherwise mention in the report the company in which the expert had any interest, nor did we suggest policy options that we have reason to believe will improperly promote or adversely affect this or any other company.

factors around development of such tools; (4) factors around adoption of such tools; and (5) two sessions focused on policy ideas that could facilitate augmenting patient care in the United States through the use of AI tools. Each session featured an open discussion among all meeting participants based on key questions we provided. The meeting was transcribed to ensure that we accurately captured the experts' statements. After the meeting, we reviewed the transcripts to characterize their responses and to inform our understanding of all three researchable objectives. Following the meeting, we continued to seek the experts' advice to clarify and expand on what we had heard. Consistent with our quality assurance framework, we provided 14 experts with a draft of our report and solicited their feedback, which we incorporated as appropriate.

Policy Options

We intend policy options to provide policymakers with a broader base of information for decision-making.⁸³ The options are neither recommendations to federal agencies nor matters for congressional consideration. They are also not listed in any specific rank or order. We are not suggesting that they be done individually or combined in any particular fashion. Additionally, we did not conduct work to assess how effective the options may be, and express no view regarding the extent to which legal changes would be needed to implement them.

Based on our requesters' interest in U.S. competitiveness and the use of AI tools in augmenting patient care, we decided on an objective designed to identify options that could help maximize benefits and mitigate challenges to the use of AI to augment patient care and for administrative purposes. We limited the policy options included in this report to those that met the policy objective and fell within the report scope. We present six policy options in response to the challenges identified during our work and discuss potential opportunities and considerations of each. While we present options to address the major challenges we identified, the options are not intended to be inclusive of all potential policy options.

To develop the policy options, we prepared a list of potential policy ideas (107 in total, including the status quo) based on our literature search, stakeholder interviews, and expert meeting. We removed ideas that were not likely to achieve the policy objective or did not fit into the overall scope of our work. We grouped the remaining ideas based on themes (e.g. data access). We combined those that (1) were duplicative, (2) could be subsumed into a higher-level policy option, or (3) were examples of how to implement a policy option rather than the option itself.

We conducted our work from November 2019 through November 2020 in accordance with all sections of GAO's Quality Assurance Framework that are relevant to technology assessments. The framework requires that we plan and perform the engagement to obtain sufficient and appropriate evidence to meet

⁸³Policymakers is a broad term including, for example, Congress, elected officials, federal agencies, state and local governments, academic and research institutions, and industry.

our stated objectives and to discuss any limitations to our work. We believe that the information and data obtained, and the analysis conducted, provide a reasonable basis for any findings and conclusions in this product.

Appendix II: Expert Participation

We collaborated with the National Academy of Medicine to convene a two-day meeting of experts to inform our work on artificial intelligence in patient care; the meeting was held virtually on March 31-April 1, 2020. The experts who participated in this meeting are listed below. Many of these experts gave us additional assistance throughout our work, including 2 who provided additional assistance during our study by sending material for our review or answering technical questions; and 14 who reviewed our draft report for accuracy and provided technical comment.

Jason Borenstein, PhD

Director, Graduate Research Ethics Programs;
and Associate Director, Center for Ethics
and Technology
Georgia Institute of Technology

Vivian Chu, PhD

Co-Founder and Chief Technology Officer
Diligent Robotics

Shahram Ebadollahi, PhD, MBA

Global Head of Data Science and AI
Novartis

Barbara J. Evans, PhD, JD, LLM

Professor of Law and Stephen C. O'Connell
Chair
University of Florida Levin College of Law
Professor of Engineering
University of Florida Wertheim College of
Engineering

Marzyeh Ghassemi, PhD

Faculty Member, Vector Institute
Assistant Professor, Computer Science and
Medicine
University of Toronto

Avi Goldfarb, PhD

Professor of Marketing and Rotman Chair in
Artificial Intelligence and Healthcare
Rotman School of Management, University of
Toronto

Elizabeth Kaziunas, PhD

Postdoctoral Researcher
AI Now Institute, New York University

Michael E. Matheny, MD, MS, MPH

Co-Director, Center for Improving the Publics'
Health Through Informatics
Associate Professor of Biomedical Informatics,
Biostatistics, and Medicine
Vanderbilt University Medical Center
Clinical Research Scientist, Tennessee Valley
Healthcare System VA Vanderbilt
University Medical Center

Ozanan Meireles, MD

Assistant Professor of Surgery
Harvard Medical School
Director, Surgical Artificial Intelligence and
Innovation Laboratory (SAIIL)
Massachusetts General Hospital

Nicholson Price, PhD, JD

Professor of Law
University of Michigan

Maulik P. Purohit, MD, MPH

Associate Chief Medical Information Officer
University Hospitals

Ronen Rozenblum, PhD, MPH

Director, Unit for Innovative Healthcare
Practice & Technology
Brigham and Women's Hospital
Assistant Professor, Harvard Medical School

Mark Sendak, MD, MPP

Population Health & Data Science Lead
Duke Institute for Health Innovation

Martín-José Sepúlveda MD, ScD

IBM Fellow
Chief Executive Officer and Principal
CLARALUZ LLC

Christina Silcox, PhD

Managing Associate
Duke-Margolis Center for Health Policy

William W. Stead, MD, FACMI

Chief Strategy Officer
Vanderbilt University Medical Center

Peter Szolovits, PhD

Professor, Computer Science and Engineering
Massachusetts Institute of Technology

Aleksandar Vakanski, PhD, PE

Assistant Professor, Industrial Technology
University of Idaho

Appendix III: GAO Contacts and Staff Acknowledgments

GAO contact

Karen L. Howard, PhD, (202) 512-6888 or howardk@gao.gov

Staff acknowledgments

In addition to the contact named above, Laura Holliday (Assistant Director), Jon D. Menaster (Analyst-in-Charge), Nora Adkins, Virginia Chanley, Leia Dickerson, Thomas Lombardi, Anika McMillon, Yesook Merrill, Rebecca Parkhurst, Monica Perez-Nelson, Daniel Setlow, Ben Shouse, and Britney Tsao made key contributions to this report. Darnita Akers, Rodney Bacigalupo, and Jane Eyre also contributed to this report.

(103919)

GAO's Mission

The Government Accountability Office, the audit, evaluation, and investigative arm of Congress, exists to support Congress in meeting its constitutional responsibilities and to help improve the performance and accountability of the federal government for the American people. GAO examines the use of public funds; evaluates federal programs and policies; and provides analyses, recommendations, and other assistance to help Congress make informed oversight, policy, and funding decisions. GAO's commitment to good government is reflected in its core values of accountability, integrity, and reliability.

Obtaining Copies of GAO Reports and Testimony

The fastest and easiest way to obtain copies of GAO documents at no cost is through GAO's website (<https://www.gao.gov>). Each weekday afternoon, GAO posts on its website newly released reports, testimony, and correspondence. To have GAO e-mail you a list of newly posted products, go to <https://www.gao.gov> and select "E-mail Updates."

Order by Phone

The price of each GAO publication reflects GAO's actual cost of production and distribution and depends on the number of pages in the publication and whether the publication is printed in color or black and white. Pricing and ordering information is posted on GAO's website, <https://www.gao.gov/ordering.htm>.

Place orders by calling (202) 512-6000, toll free (866) 801-7077, or TDD (202) 512-2537.

Orders may be paid for using American Express, Discover Card, MasterCard, Visa, check, or money order. Call for additional information.

Connect with GAO

Connect with GAO on [Facebook](#), [Flickr](#), [Twitter](#), and [YouTube](#).

Subscribe to our [RSS Feeds](#) or [E-mail Updates](#).

Listen to our [Podcasts](#) and read [The Watchblog](#).

Visit GAO on the web at <https://www.gao.gov>.

To Report Fraud, Waste, and Abuse in Federal Programs

Contact: Website: <https://www.gao.gov/fraudnet/fraudnet.htm>

Automated answering system: (800) 424-5454 or (202) 512-7470

Congressional Relations

Orice Williams Brown, Managing Director, WilliamsO@gao.gov, (202) 512-4400,
U.S. Government Accountability Office, 441 G Street NW, Room 7125, Washington, DC 20548

Public Affairs

Chuck Young, Managing Director, YoungC1@gao.gov, (202) 512-4800
U.S. Government Accountability Office, 441 G Street NW, Room 7149, Washington, DC 20548

Strategic Planning and External Liaison

Stephen Sanford, Acting Managing Director, spel@gao.gov, (202) 512-9715
U.S. Government Accountability Office, 441 G Street NW, Room 7B37N, Washington, DC 20548