



How Artificial Intelligence Is Changing Health Care Delivery

Citation

Sanders, Samantha. 2020. How Artificial Intelligence Is Changing Health Care Delivery. Doctoral dissertation, Harvard Medical School.

Permanent link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37364994>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Scholarly Report submitted in partial fulfillment of the MD Degree at Harvard Medical School

Date: 05 January 2020

Student Name: Samantha Sanders, MD/MBA candidate

Scholarly Report Title: How Artificial Intelligence is Changing Health Care Delivery

Mentor Name and Affiliations: Ariel D. Stern, PhD, Poronui Associate Professor of Business Administration, Harvard Business School; Associate Faculty Member, Ariadne Labs, Brigham and Women's Hospital and Harvard Chan School of Public Health; Affiliated Faculty Member, Harvard-MIT Center for Regulatory Science

Collaborator Names and Affiliations: Mats Terwiesch, BA, Research Associate, Harvard Business School; William J. Gordon, MD, Medical Director, Health Innovation Platform, Partners Personalized Medicine; Instructor of Medicine, Harvard Medical School; Associate Physician, Hospital Medicine Unit, Division of General Medicine and Primary Care, Brigham & Women's Hospital

Link to and Citation for Publication: Sanders, S.F., Terwiesch, M., Gordon, W.J., Stern, A.D. How Artificial Intelligence Is Changing Health Care Delivery. NEJM Catalyst. 2019, Oct. Available from: <https://catalyst.nejm.org/health-care-ai-systems-changing-delivery/>.

Abstract

Title: How Artificial Intelligence is Changing Health Care Delivery

Purpose: The development of intelligent machines and artificial intelligence holds great promise for making health care delivery more accurate, efficient, and accessible. This article aims to summarize these potential applications of artificial intelligence, as well as challenges that remain for incorporating this technology into clinical and administrative settings.

Methods: We performed an extensive literature review and identified frameworks used to classify different categories of artificial intelligence applications as well as different stages of development applicable to these applications.

Results: We identified three primary categories of artificial intelligence applications related to health care delivery: administrative, diagnostic, and treatment-related. In each of these categories, applications varied with regards to stages of development. For example, deep learning algorithms are advanced with regards to pathologic and radiologic diagnosis; however, models that incorporate clinical data to predict diagnoses are less developed.

Conclusions: Artificial intelligence can transform health care delivery. Nonetheless, several questions remain, including how humans will interact with artificial intelligence interfaces, how tools will be reimbursed and regulated, and how algorithms can be tailored for deployment in a variety of settings.

Table of Contents

<i>Glossary and Description</i>	<i>4</i>
<i>Appendix - Paper Text.....</i>	<i>5</i>
<i>Appendix - Figure 1.....</i>	<i>12</i>
<i>Appendix - Figure 2.....</i>	<i>15</i>

Glossary and Description

Glossary

AI – Artificial Intelligence

FDA – U.S. Food and Drug Administration

EMR – Electronic Medical Record

Description of Scholarly Project and Contribution

This project addressed the question of how readily artificial intelligence is being incorporated into a variety of health care delivery settings, as well as what challenges remain with regards to the implementation of these tools.

With regards to my contribution to the project, I came up with the idea for the project and led the research and writing process. My mentor and I worked together to define the analytical frameworks we used to present different categories of artificial intelligence applications and to define the stage of development of these applications. I worked with Mats (one of the contributors) to research and identify examples to present throughout the text of the article and in Figure 1. Will Gordon (another one of the contributors) primarily contributed during the editing and revisions process.

Link to and Citation for Publication

Sanders, S.F., Terwiesch, M., Gordon, W.J., Stern, A.D. How Artificial Intelligence Is Changing Health Care Delivery. NEJM Catalyst. 2019, Oct. Available from: <https://catalyst.nejm.org/health-care-ai-systems-changing-delivery/>.

Appendix - Paper Text

A nurse avatar named “Molly” who regularly talks with patients about their symptoms and medical needs. Voice-recognition software that helps physicians document clinical encounters. A prescription drug-monitoring platform that can detect patients’ opioid misuse. Systems that analyze millions of medical images to help physicians diagnose and predict diseases. Robots that extend the reach of surgeons.

These innovations are all powered by artificial intelligence (AI), a burgeoning field of computer science that is already reshaping many aspects of health care by harnessing vast amounts of data to improve diagnosis and treatment, save time and costs, and expand access to care worldwide.

AI is broadly defined as the development of intelligent machines – computer systems – that perform functions such as thinking and problem solving that normally require human intelligence. It encompasses natural language processing (where devices like the iPhone’s Siri decipher and respond to human language), machine processing (where technology like self-driving cars process visual data – e.g., identify other cars and pedestrians on the road), and machine learning (where computers learn by experience, such as repeatedly performing and perfecting a skill like playing chess).

Health care AI includes a growing collection of algorithms that drive hardware and software systems to analyze health care data. These systems have the potential to do everything from detecting insurance fraud to improving clinical trial recruitment to sharpening diagnostic images. Significant advances in AI methods have propelled well-known technology firms to invest in health care AI, alongside a plethora of specialized startups that have emerged in the field. AI is expected to continue to improve patient care and meaningfully change the activities undertaken by clinicians, health care provider organizations, payers, pharmaceutical firms, and medical technology companies.

In this article, we focus on applications in the health care *delivery* setting – that is, for clinicians and provider organizations. AI offers [compelling opportunities](#) to improve efficiency, reduce

errors, and incorporate increased evidence-based decision support. However, challenges abound in areas such as data security, patient privacy, legal liability, and the challenges of [applying AI tools in new contexts](#).

We describe AI's existing and potential role in reducing workloads, lowering costs, and bettering outcomes across three key domains of health care delivery: administrative work, diagnosis, and treatment. Within each domain, we provide examples of AI applications (Table 1). We also apply a framework for categorizing the stages of AI technology development and diffusion into research and practice (Table 2). We developed this tool based on [other frameworks](#) that have outlined key stages of innovation and new product development. We conclude by discussing several universal obstacles to applying AI to health care delivery.

Administrative Work

Administration represents a significant cost to the health care delivery system. For every office-based physician in the U.S., there are [2.2 administrative workers](#), exceeding the number of nurses, clinical assistants, and technical staff combined. Studies suggest that the “paper work of medicine” is a large burden to physicians and hospitals worldwide, and in the U.S. in particular, due to the variety and complexity of insurance and reimbursement systems. Applying AI to administrative tasks presents opportunities to improve the quality and efficiency of health care delivery for both patients and providers.

For example, AI-driven voice-recognition software can be used to document a broad range of clinical encounters. Indeed, this technology is already at a mature stage of development (diffusion and implementation), with the [Dragon Medical One](#) platform available in hospitals around the world to help providers [dictate patient notes](#). Speech-recognition software could also be used in scheduling follow-up appointments and generating emails, orders, and prescriptions. These tools hold great promise for improving provider efficiency by reducing time spent on manual data entry, although it remains uncertain whether they will live up to their potential, given the [steep learning curve](#) for voice-recognition software.

Patient triage represents another promising application of AI in health care delivery. For instance, a virtual nurse “chatbot” – a computer program that simulates conversation – regularly monitors patients and can direct them to providers as needed. In another case, an AI-powered therapist has been used to triage mental health patients, and research shows that many veterans with PTSD [feel more comfortable](#) speaking with a virtual chatbot than a human provider. This [therapist avatar](#) uses facial analysis and natural language processing (similar to the technology in Amazon’s virtual assistant Alexa) to detect patients’ moods, identify whether they have PTSD, and direct them to appropriate care. This type of technology is at an intermediate stage of development (prototype available), but in the future, AI-driven tools could play a significant role in patient triage in both virtual and live settings ranging from the patient’s home to urgent care.

Several challenges exist in deploying and using AI for administrative health care work. For example, developers may need to adapt software tools to meet provider preferences and site-specific norms, such as how a particular clinic handles triaging, clinical note-writing conventions, or scheduling. Moreover, patient privacy considerations will continue to emerge. These are likely to be especially important when AI generates new forms of data or handles sensitive information like mental health or genetic risk data. Patient consent may be needed before clinical encounters in some cases. Additionally, much of this technology will need to comply with the [HIPAA Privacy Rule](#), which establishes standards for safeguarding protected health information. The recent announcement that Amazon’s Alexa is now HIPAA compliant represents an important step toward applying AI to health care data in a responsible way.

Diagnosis

Artificial intelligence supports several diagnostic activities and processes – for example, reviewing imaging, visual findings, and pathology results. AI is currently being used to develop tools to support pathologists, radiologists, dermatologists, ophthalmologists, and other physicians who conduct visual diagnoses. These systems typically use deep learning, a type of machine learning in which human brain-like algorithms learn to solve problems by repeatedly performing tasks (e.g., reviewing different types of tumor scans or pathology slides), without the need for additional human involvement.

Take cancer diagnosis. In a [challenge competition](#) that simulated the reading of pathology slides, [seven deep learning algorithms](#) outperformed a panel of 11 pathologists in detecting lymph node metastases in tissue sections from women with breast cancer. Additionally, [Paige.AI](#), a startup developing AI tools to improve cancer pathology diagnosis, was recently granted “Breakthrough Device” designation by the U.S. Food and Drug Administration (FDA). The company is building a dataset of de-identified digital pathology slides, with the goal of helping pathologists become faster and more accurate in their cancer diagnoses and treatment recommendations.

Diagnostic applications are already at a late stage of development (commercialization). In the short term, the most feasible implementation model is one in which providers oversee the (preliminary) diagnostic work of AI systems to increase their own efficiency. A radiologist, for instance, could save time by supervising the work of an algorithm programmed to analyze images, find abnormalities, and provide preliminary analyses. A [hybrid model](#) like this that combines AI-driven smart machines and human experts is likely to be most successful in improving outcomes, as algorithms are more likely to identify false positives (Type I error), while clinicians may be more likely to find false negatives (Type II error), as seen in the [pathology challenge](#) mentioned above.

AI can also augment the efficiency of non-visual diagnostic methods. For example, a [study](#) showed that machine-learning algorithms helped investigators streamline the process for accurately diagnosing patients with autism spectrum disorders; they did so by significantly reducing the number of screening items used in the traditional evaluation instrument. The results offer hope for creating mobile tools that could speed the pace of autism diagnosis and reach a larger at-risk population.

Another application of AI involves novel classifications of patients and disease. Using deep learning, for example, researchers recently described and validated [four new phenotypes of sepsis](#) based on clinical data from nearly 64,000 hospitalized patients, expanding clinicians’ understanding of this heterogeneous syndrome. In another study, investigators used deep learning to create a novel representation of a patient from electronic health record data, with excellent prediction performance for a [variety of diseases](#). This work is currently at an early

stage of development (academic research). Eventually, such algorithms could harness data from traditional sources, such as labs and imaging, and non-traditional sources like wearable devices to support patient diagnosis.

The key challenges related to incorporating AI into diagnostic processes include integrating algorithms into clinical workflows, demonstrating the safety and effectiveness of algorithms to meet regulatory standards, and consistently updating algorithms with new and more representative data. Moreover, diagnosing some rare diseases may prove difficult when there isn't enough data available to develop effective predictive algorithms.

Treatment

AI-powered systems can support and improve patient treatment in a range of health care delivery settings, from large hospitals to small clinics, and through a range of services. These might include administering the treatment itself (e.g., assisting in surgeries) to developing and modifying personalized treatment plans for patients.

In the operating room, AI already aids physicians in robotic-assisted procedures by providing a suggested roadmap and warnings throughout the process. Robotic surgery is used for a variety of [different procedures](#), from coronary bypass to kidney removal, although not all involve AI. [One study of 379 patients](#) undergoing minimally invasive spinal fusion surgery found that the robotic-guided technique led to a five-fold reduction in surgical complications at three months and one year post-operation. Similarly, a [prospective randomized controlled trial](#) showed that using an AI system during colonoscopy resulted in an almost two-fold higher detection rate of precancerous adenomas. The development of these technologies is advanced: many have already reached the stage of diffusion and implementation.

AI algorithms have also proven useful in generating evidence-based treatment recommendations for certain conditions. For instance, Stanford academics are developing [HealthRex](#), a clinical decision support platform that mines EMR data to inform physicians how their peers [managed similar patient cases](#). While this technology remains at an early stage of development (academic

research), it has the potential to facilitate real-time treatment recommendations based on evolving health care data and evidence-oriented clinical experience.

Like the challenges faced around diagnosis, treatment-based algorithms require integrating AI into existing clinical workflows, demonstrating safety and effectiveness above and beyond current best practices, and updating algorithms regularly to reflect representative patient datasets and rapidly evolving medical practices. Additionally, such algorithms must be flexible enough to account for provider preferences and the ambiguities of clinical treatment protocols.

Conclusion

The variety of applications described here showcase how artificial intelligence systems are already transforming health care delivery as we know it. AI holds great promise for making health care delivery more accurate, efficient, and personalized, and AI-driven tools will likely soon touch virtually every aspect of administrative, diagnostic, and treatment domains.

Nonetheless, [humans still will be required](#) to sustain the core doctor-patient relationship, and several universal obstacles could hinder the pace of AI adoption in health care delivery for the foreseeable future. For example, development of AI technology requires access to data sources that accurately and equitably reflect the general patient population. However, medical AI is likely to emerge in high-resource settings, such as academic medical centers, leading to “[contextual bias](#)” when it is deployed in lower-resource settings such as community health centers or rural areas. Additionally, developers must attain physician and staff buy-in regarding using AI in clinical practice. Despite great promise, some digital health tools have [faltered](#) at the stage of clinical adoption and diffusion.

Finally, questions remain about reimbursement, liability, and regulation. In the U.S., the FDA has been establishing regulatory policies and guidelines around software and digital health. The agency has, for example, built a pre-certification program aimed at making sure consumers have access to high-quality digital health products.

A host of different stakeholders play key roles in overseeing and implementing these AI technologies, including hardware and software developers, clinicians, hospital administrators, and regulators. Each of these stakeholders is essential to the safe and secure diffusion of AI within health care delivery. Developers and clinicians must [work together](#) to carry out rigorous studies and clinical validation before using AI systems for patient care. Hospital administrators must evaluate AI in the context of developmental stages (Table 2) to select opportunities for adopting new technologies. Finally, regulators must continue to refine their role in legitimizing and approving AI-driven tools.

Despite the abundance of challenges in this space, the application of secure, well-validated AI systems holds great potential for improving the health care delivery experience, both now and into the future.

Appendix - Figure 1

Category	Sub-category	Examples
Administrative Work	Provider documentation	<ul style="list-style-type: none"> • A proof-of-concept system that captures data from clinical encounters can recognize medical language and categorize it by diseases, medications, tests, results, and symptoms. • Dragon Medical One, a speech-recognition platform for clinical documentation, has been shown to yield significant productivity benefits for hospitals around the world. • Developers of the voice-based virtual assistant Suki, which converts a physician’s speech to text for dictating notes into the electronic medical record (EMR), secured a reported \$20 million in funding in 2018.
	Order, prescription, coding entry for providers	<ul style="list-style-type: none"> • Clinical decision-support technologies exist to improve the ordering process, such as warnings in EMR software programs when prescribed medications are known to interact. • MedAware, an Israeli startup, uses machine learning analysis to detect deviations in prescription patterns and warn physicians of potential errors when ordering medications. • Texas and Maryland use RxGov to track patients’ drug prescriptions across providers and detect signs of abuse.
	Data entry	<ul style="list-style-type: none"> • athenahealth uses machine learning to categorize faxes and quickly add them to the appropriate physician’s workflow. • Several platforms for data entry automation are available to hospitals, such as DocuPhase. • Apixio, a growing analytics company with an estimated \$100 million valuation, offers an AI-based tool to help health care organizations conduct internal audits of their risk adjustment payment data.
	Scheduling	<ul style="list-style-type: none"> • In 2016, athenahealth acquired Arsenal Health and its machine learning platform for schedule optimization and predictive analytics. • Academic research has demonstrated the superiority of using a machine learning model to predict patient no-shows and allow a hospital to effectively overbook. • PetalMD’s AI-based schedule management system advertises that it can reduce the time to create and maintain schedules by up to 80%.

	Triageing	<ul style="list-style-type: none"> • Several large health organizations, including Mayo Clinic and UK’s National Health Service, use Sensey's AI-powered nurse avatar to monitor patients and direct them to appropriate care. • An FDA-approved product from Zebra Medical, which provides AI-based technology to support radiologists, generates an alert for urgent findings of pneumothorax. • A multisite study of emergency departments finds that a machine learning-based model implemented at select hospitals is more effective than relying on the standard ED triage algorithm. • Following a successful trial of Enlitic's patient triage platform involving 10,000 chest X-rays, the company will provide its deep learning solutions to businesses offering private health exams across China. • Researchers at the University of Southern California developed a federally funded therapist avatar that uses facial analysis and natural language processing to detect a patient’s mood and identify cases of PTSD.
Diagnosis	Imaging / pathology review	<ul style="list-style-type: none"> • Deep learning algorithms can: diagnose diabetic retinopathy from eye scans with over 90% accuracy; outperform pathologists in detecting metastases from axillary lymph node tissue sections; and differentiate images of malignant and benign skin lesions as well as dermatologists can. • Paige.AI, a startup developing AI tools to improve cancer pathology diagnosis, is granted “Breakthrough Device” designation by the U.S. Food and Drug Administration. • A worldwide competition produces AI solutions that replicate the accuracy of expert radiation oncologists in targeting lung tumors while performing more rapidly than physicians. • A deep learning model uses full-field mammograms to predict likely development of breast cancer years in advance.
	Diagnostic models / symptom analysis	<ul style="list-style-type: none"> • An algorithm is capable of streamlining the process for diagnosing patients with autism spectrum disorder. • Ada DX has been validated as a Diagnostic Decision Support System for diagnosing rare diseases.
	Phenotyping	<ul style="list-style-type: none"> • Researchers identified four novel sepsis phenotypes using a “k-means” clustering algorithm. • An unsupervised machine learning approach derives a general-purpose patient representation by aggregating EHR data from 700,000 patients to facilitate clinical prediction.

	Incorporating non-traditional data sources	<ul style="list-style-type: none"> • A machine learning algorithm that detects abnormal speech features in free-text samples was used to accurately predict which patients later developed psychosis.
Treatment	Surgical assistance	<ul style="list-style-type: none"> • Researchers demonstrate that robotic-guided spinal fusion surgery significantly reduces post-operation surgical complications. • Mazor Robotics claims its robotics technology is the subject of 55 peer-reviewed articles and is used in about 36,000 procedures.
	Individualized / personalized medicine	<ul style="list-style-type: none"> • Participants on Kaggle, a crowd-sourcing data science community, have shown that machine learning can be used to enable personalized medicine.
	Adherence and health coaching	<ul style="list-style-type: none"> • A neural network computer vision algorithm can identify patient and drug information, as well as confirm the ingestion of direct oral anticoagulants with smartphone technology. • Twine Health, a health-coaching and chronic disease management platform acquired by Fitbit, has demonstrated positive health outcomes in a case study at Joslin Diabetes Center. • The FDA approves Abilify MyCite, a sensor that allows physicians to digitally monitor whether a patient has ingested medication to treat mental health disorders.
	Generating treatment recommendations	<ul style="list-style-type: none"> • Stanford academics are developing HealthRex, a clinical decision-support platform that informs physicians how their peers managed similar patient cases. • A randomized clinical trial demonstrates success in using an algorithm for guiding treatment of staph infection, compared to standard care.
	Digital therapeutics	<ul style="list-style-type: none"> • Pear Therapeutics released two FDA-approved digital therapeutics that offer cognitive behavioral therapy to support treatment of substance use disorder and opioid use disorder. • Akili Interactive has created a tablet-based digital therapeutic that resembles a video game to treat pediatric patients with ADHD. • In January 2019, Otsuka Pharmaceutical agreed to pay up to \$300 million to Click Therapeutics to develop digital therapies for major depression.

Appendix - Figure 2

Stages of Development: Overview and Examples

Stages of development	Stage 1: Idea generation	Stage 2: Academic research	Stage 3: Prototype available	Stage 4: Commercialization	Stage 5: Diffusion and implementation
Description	Initial proposals are being made	Academic or commercial research is being conducted	Prototype is available for use	Product has been brought to market	Product is widely used
Example	<ul style="list-style-type: none"> In January 2019, Otsuka Pharmaceutical agreed to pay up to \$300 million to Click Therapeutics to develop digital therapies for major depression. 	<ul style="list-style-type: none"> Stanford academics are developing HealthRex, a clinical decision-support platform that informs physicians how their peers managed similar patient cases. 	<ul style="list-style-type: none"> Researchers at the University of Southern California developed a federally funded therapist avatar that uses facial analysis and natural language processing to detect a patient's mood and identify cases of PTSD. 	<ul style="list-style-type: none"> MedAware, an Israeli startup, uses machine learning analysis to detect deviations in prescription patterns and warn physicians of potential errors when ordering medications. 	<ul style="list-style-type: none"> Dragon Medical One, a speech-recognition platform for clinical documentation, has been shown to yield significant productivity benefits for hospitals around the world.