

SPECIAL COMMUNICATION

A Clinician's Guide to Artificial Intelligence (AI): Why and How Primary Care Should Lead the Health Care AI Revolution

Steven Lin, MD

Artificial intelligence (AI) in health care is the future that is already here. Despite its potential as a transformational force for primary care, most primary care providers (PCPs) do not know what it is, how it will impact them and their patients, and what its key limitations and ethical pitfalls are. This article is a beginner's guide to health care AI, written for the frontline PCP. Primary care—as the dominant force at the base of the health care pyramid, with its unrivaled interconnectedness to every part of the health system and its deep relationship with patients and communities—is the most uniquely suited specialty to lead the health care AI revolution. PCPs can advance health care AI by partnering with technologists to ensure that AI use cases are relevant and human-centered, applying quality improvement methods to health care AI implementations, and advocating for inclusive and ethical AI that combats, rather than worsens, health inequities. (J Am Board Fam Med 2022;35:175–184.)

Keywords: Artificial Intelligence, Deep Learning, Delivery of Health Care, Health Equity, Information Technology, Machine Learning, Primary Health Care, Quality Improvement, Social Justice, Technology

Introduction

Artificial intelligence (AI) is poised as a transformational force in health care, and primary care is where the power, opportunity, and future of AI are most likely to be realized in the broadest and most ambitious scale. This article will (1) define AI, machine learning, and deep learning; (2) describe 10 ways AI is transforming health care; (3) explore the key limitations and ethical pitfalls of AI; (4) explain 3 reasons why primary care should lead health care AI; and (5) discuss 3 ways how primary care will lead health care AI.

A Primer on AI for Primary Care Providers

There are many definitions of AI, but the simplest is just a process where a computer is trained to do a task in a way that mimics human behavior (Figure 1).¹ In other words, AI is a machine programmed to do something that a human might do; eg, reminding a clinician that magnetic resonance imaging might be inappropriate for uncomplicated low back pain when it is ordered in the electronic health record (EHR). Many EHR alerts and clinical decision support tools fall under this category of “classic” AI.

In the modern era when AI is discussed, what most people are actually referring to is machine learning (ML)—algorithms that allow computers to

This article was externally peer reviewed.

Submitted 27 May 2021; revised 9 September 2021 and 16 September 2021; accepted 17 September 2021.

From Stanford Healthcare AI Applied Research Team, Division of Primary Care and Population Health, Department of Medicine, Stanford University School of Medicine, Stanford, California.

Funding: This work was supported by the Division of Primary Care and Population Health, Department of Medicine, Stanford University School of Medicine.

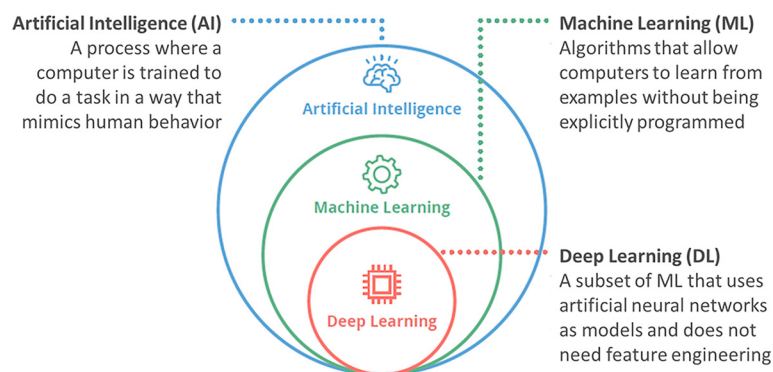
Conflict of interest: The author is a principal investigator working with companies and nonprofit organizations through grants and sponsored research agreements administered by Stanford University. Current and previous collaborators include Amazon, American Academy of Family Physicians, American Board of Artificial Intelligence in Medicine, American Board of Family Physicians, Center for Professionalism and Value in Health

Care, Codex Health, DeepScribe, Eko Health, Google Health, Quadrant Technologies, Soap Health, Society of Teachers of Family Medicine, University of California, San Francisco, and Verily. Neither the author, nor members of his immediate family, has any financial interest in these organizations.

Prior presentation: This work was presented at the Society of Teachers of Family Medicine's (STFM) Virtual Conference on Medical Student Education (February 2, 2021), STFM Virtual Annual Spring Conference (May 3, 2021), and STFM Virtual Conference on Practice and Quality Improvement (September 14, 2021).

Corresponding author: Steven Lin, MD, 211 Quarry Road, Suite 405, MC 5985, Palo Alto, CA (E-mail: stevenlin@stanford.edu).

Figure 1. Definitions of artificial intelligence, machine learning, and deep learning.



learn from examples, without being specifically programmed.¹ In other words, a machine is programmed to do something that a human might do, and every time it encounters similar examples, it improves on its ability to perform that task. Many prediction models fall under this category of ML; eg, computers that can predict which patients are at risk for unplanned intensive-care unit transfers, preventable hospital-acquired conditions, or unnecessary hospitalizations and emergency department (ED) visits.

There is an even more powerful subset of ML called deep learning (DL), which uses algorithms designed to mimic the human brain (also known as artificial neural networks) to vastly increase the computer's learning potential—even allowing it to teach itself, without the need for explicit programming.¹ Many diagnostic models fall under this category of DL; eg, computers that can analyze radiology images, pathology slides, dermatology photographs, or even a patient's physical movements and provide differential diagnoses.

AI has already permeated every industry and every facet of our daily lives. Every time a shopping website gives a recommendation for what to buy, a streaming service gives a suggestion for what to watch, a voice assistant gives directions, a search engine gives results based on the user's search history and location, a social media platform gives a personalized newsfeed, or a jobs app matches the user with a potential employer—that is AI.

The reality is, the speed of technology adoption has accelerated so much that it has become barely perceptible. It took electricity 46 years until it was used by 25% of all Americans; the TV, 26 years; but only 3 years for the smartphone to get into the hands of 1 in every 4 Americans.²

It is not just the speed of adoption that is accelerating, it is also the raw power of the technologies.

Before 2012, most technologies—including AI—followed Moore's Law, the observation that computing power doubles every 2 years.³ Moore's Law explains why there's more computing power in a modern smartphone than in the Apollo spacecraft that sent humans to the moon. But in the last decade, especially with the advancement of ML and DL, everything has changed: AI is now outpacing Moore's Law by doubling every 3 to 4 months.³

To put that power into context, consider chess. Chess has been used as a Rosetta Stone of both human and machine cognition for over a century.⁴ Ever since Garry Kasparov, the former world champion, was famously defeated by the computer Deep Blue in 1996, the best human chess players in the world have been unable to keep up with the best computer chess engines. In 2016, the most powerful chess engine was Stockfish, programmed with all the chess knowledge in human history. But Stockfish was a product of classic AI; so how did it fare against modern AI, powered by ML/DL?

In 2017, scientists introduced AlphaZero, which was a DL AI that taught itself how to play chess with no prior built-in domain knowledge, just the basic rules of the game. In only 4 hours of training, AlphaZero was playing chess at a higher rating than Stockfish. In 9 hours, when scientists put them head to head in a time-controlled 100-game tournament, AlphaZero crushed Stockfish: 28 wins, zero losses, 72 draws. This was published in the journal *Science* and the headlines were jaw-dropping: in just a few hours, AlphaZero mastered all of the chess knowledge in human history.⁵

Although the practice of medicine is substantially more complex and carries higher stakes than chess, imagine if it was possible to harness even a fraction of that power and direct it toward health

care. AI has the potential to fundamentally change the way society thinks about medicine, the way medicine is practiced, and the way medicine is taught. It is no exaggeration to say that today humanity stands at an inflection point in history, where even those who may have been disappointed in the past by health care technologies—with all the hype, the overpromise, the underdelivery—have no choice but to contend with the coming era of innovation and prepare for a future that is already here.

Ten Ways AI Is Transforming Health Care

This section will explore what is happening in health care AI today (Figure 2).

Risk Prediction and Intervention

In the United States, hospital costs for potentially preventable conditions account for 1 in every 10

dollars of total expenditures.⁶ With AI, it is possible to take EHR data and calculate risk to patients based on demographics, comorbidities, medications, labs, imaging, and social determinants of health.⁷ Understanding risk at the individual level allows health systems to devise team-based interventions in the primary care setting to engage patients at highest risk and reduce preventable ED visits, hospitalizations, and death.

Population Health Management

AI can assist with identifying and closing care gaps that would otherwise take primary care providers (PCPs) 7 to 8 hours to do every working day.⁸ Future population health platforms will go beyond simple dashboards and have semiautomated systems that can reach out to patients when their care gaps are due and have shared decision-making encounters for cancer screenings and other preventive services.

Figure 2. Ten ways artificial intelligence is transforming health care today. Abbreviations: AI, artificial intelligence; ED, emergency department; EHR, electronic health record; PCP, primary care provider.

Application Area	Description and Comment
Risk prediction and intervention	AI can take the health records of entire populations and calculate risk to patients based on EHR data. Understanding risk at the individual level allows health systems to devise team-based interventions in the primary care setting to engage patients at highest risk and reduce preventable ED visits, hospitalizations, and death.
Population health management	AI can assist with identifying and closing care gaps using semi-automated systems that can reach out to patients when their care gaps are due, and have shared decision making encounters for cancer screenings and other preventive services.
Medical advice and triage	AI can triage patient complaints and free up primary care access for more appropriate care. Chatbots can provide health advice directly to patients with common symptoms.
Risk-adjusted paneling and resourcing	AI can ensure that clinicians have adequate time to address the needs of each patient by adjusting panel sizes based on complexity.
Remote patient monitoring	AI can organize large amounts of data coming from smart devices to provide actionable clinical insights – often paired with digital health coaching.
Digital health coaching	AI-powered tools that can help patients self-manage some of the costliest chronic diseases, such as diabetes, obesity, hypertension, and depression.
Chart review and documentation	AI can listen in on patient-clinician conversations and generate a note, like a human scribe. AI can read notes and extract insights from unstructured data.
Diagnostics	AI-powered algorithms for diagnosing disease will broaden the services PCPs can provide to patients, reduce the need for unnecessary referrals, and expand access to care in regions with lack of specialty care.
Clinical decision making	AI assistants, built into EHR platforms, can give physicians evidence-based clinical suggestions at the point of care.
Practice management	AI can automate repetitive administrative tasks that are suffocating practices like billing, coding, and prior authorizations. It can also automate certain aspects of inter-visit care planning and management.

Medical Advice and Triage

Chatbots are already being used to provide health advice directly to patients with common symptoms. A recent study showed that in a dataset of 150,000 patient interactions with an AI triage tool, the urgency of patients' intended level of care decreased in more than one quarter of the cases.⁹ This shows the feasibility of using AI to triage patient complaints and free up primary care access for more appropriate care.

Risk-Adjusted Paneling and Resourcing

AI can ensure that clinicians have adequate time to address the needs of each patient by adjusting panel sizes based on complexity.¹⁰ That same data in the hands of health systems may be used to ensure optimal staffing resources based on the intensity of care being provided.

Remote Patient Monitoring and Digital Health Coaching

One in 5 Americans now has a device that tracks vital signs or other health measures.¹¹ Many companies have learned how to leverage that data by pairing them with AI-powered coaches that can help patients self-manage some of the costliest chronic diseases, such as diabetes, obesity, hypertension, and depression. Some systems have been able to demonstrate outcomes that are comparable or superior to usual care.¹²

Chart Review and Documentation

For every 1 hour clinicians spend in front of patients, they spend another 2 hours behind the computer, mostly on chart review and documentation.¹³ AI can listen in on patient-clinician conversations and generate a note, similar to a human scribe.¹⁴ This technology has the potential to unshackle clinicians from the EHR and allow them to pay more attention to their patients.

Diagnostics

AI-powered algorithms for diagnosing disease will broaden the services PCPs can provide to patients, reduce the need for unnecessary referrals, and expand access to care in regions with lack of specialty care. Here are some examples: (1) a computer that can scan facial expressions, analyze speech and affect, to screen for depression and anxiety;¹⁵ (2) a camera that can screen for diabetic retinopathy in the PCP's office;¹⁶ and (3) a system that takes

photographs of skin lesions and uses AI to provide diagnostic and management recommendations.¹⁷

Clinical Decision-Making

AI assistants, built into EHR platforms and combining some of the aforementioned technologies, can help PCPs with clinical decision-making at the point of care. Imagine a clinician speaking with a patient; there is an AI listening in and generating a note; that AI is also predicting the clinical decisions that the clinician might need to make based on the live conversation and providing clinical insights and recommendations in real time.¹⁸

Practice Management

AI can automate repetitive clerical tasks that are suffocating practices like billing, coding, and prior authorizations.¹⁹ It can also automate certain aspects of intervisit care planning to make actual visits more efficient and rewarding for both patients and clinicians alike.²⁰ Some have proposed new models of primary care, powered by humans and augmentable by AI, focused on care between office visits, recognizing that health care is moving out of the 4 walls of the examination room and into the virtual space, especially in the post-COVID-19 era.²¹

Key Limitations and Ethical Pitfalls of AI

Despite AI's vast potential to improve health care, it is important to recognize its key limitations and ethical pitfalls. Though a complete discussion falls outside the scope of this article, here are some of the critical challenges.

Coded Bias

Like any data-powered tool, AI is vulnerable to bias and abuse. There is no such thing as machine neutrality. Automated systems are not inherently neutral—they reflect the priorities, preferences, and prejudices of the humans who created them. In other words, if the data are biased, so is the AI.

The phenomenon known as coded bias was proven in a 2018 study that found that commercially available facial recognition tools performed better on males than females and better on lighter subjects than darker subjects.²² This work was a milestone in AI bias research, and in June 2020, after the killing of George Floyd, IBM, Microsoft, and Amazon

all announced that they would stop selling their facial recognition software to law enforcement.²³

Coded bias exists broadly and in many areas of AI, including health care. A recent study showed that some prediction models are more accurate for affluent White male patients—because they were trained on data from that demographic—than they are for Black, female, or low-income patients.²⁴ The risk of insufficient or unequal representation of data for underserved populations is inherent in algorithms based on EHR data. Unequal access to care among racial and ethnic minorities often leads models to underpredict their risk.²⁵ Adding social determinants of health data—including neighborhood, environment, lifestyle behaviors, habits, language, transportation, income, financial strain, social support, and education—to prediction models can improve model accuracy for hospitalization, death, and costs of care.²⁶

Racial justice requires algorithmic justice, because the potential harms from biased algorithms making decisions are real.²⁷ As Dr. Joy Buolamwini, founder of the Algorithmic Justice League, remarked: “In today’s world, AI systems are used to decide who gets hired, the quality of medical treatment we receive, and whether we become a suspect in a police investigation. While these tools show great promise, they can also harm vulnerable and marginalized people, and threaten civil rights. Unchecked, unregulated, and at times, unwanted AI systems can amplify racism, sexism, ableism, and other forms of discrimination.”²⁸


Nongeneralizability

AI models often score well on statistical tests of accuracy but perform surprisingly poorly in real-world clinical settings. A recent study showed that a commonly used early warning system for sepsis missed two thirds of cases, rarely found cases the medical staff did not notice, and frequently issued false alarms.²⁹ Some models work well in 1 geographic region but not in others.³⁰ Many AI models have a tendency to become less accurate over time, a problem known as “calibration drift.”³¹ Addressing these fundamental data science challenges will require new methods of training and retraining AI models on very large datasets while also protecting data privacy, security, and access rights.

Profit-Driven Design

Some technology companies prioritize profit over the common good, and the algorithms that they built

Figure 3. The window of humane technology.



have created major unintended consequences. This was perhaps best described in the documentary *The Social Dilemma*, which has been viewed by over 100 million people in 190 countries.³² In it the filmmakers describe how under intense pressure to grow, social media platforms have made engagement the highest priority, intentionally creating algorithms to drive clicks and views. This design has been linked to serious harms, including propagation of misinformation, conspiracy theories, social media addiction, degradation of interpersonal relationships, and even the upheaval of politics and elections that have led to systemic oppression and violence.³³

Consider what happened at the US Capitol on January 6, 2021,³⁴ or what is happening today in the fight against the COVID-19 pandemic—both complicated by the fact that fake news spreads 6 times faster than real news on social media.^{35,36} Weaponized misinformation is an example of what can happen when AI, in the hands of motivated actors, overwhelms human capabilities to discern truth from lies. In the case of health care AI, society has a vital responsibility to ensure that models cannot be co-opted by nefarious interests to block access, deny coverage, drive up prices, sell unproven products, or marginalize subgroups.

There is a concept known as the window of humane technology (Figure 3), where one wants AI to overcome human limitations so that it can augment human capabilities but never overwhelm them. Striking that right balance—hitting that window—is extraordinarily difficult, especially considering the raw power of AI, the rate at which it is progressing, and the motto of some in the technology industry to “move fast and break things.”³⁷ The path to humane technology will require society to demand that technology serve everybody, not the privileged few—that it serves the common good, not profit.

Why Primary Care Should Lead Health Care AI

This section will consider why primary care is the most uniquely suited specialty to lead the health care AI revolution.

Primary Care Is the Dominant Force at the Base of the Health Care Pyramid

With more than 500 million visits per year in the United States—55% of all clinician office visits, more than all other specialties combined—primary care is where the power, opportunity, and future of AI are most likely to be realized in the broadest and most ambitious scale.³⁸ Patients interact with PCPs often and in a variety of contexts (eg, acute, chronic, preventive care) that generate the broadest array of both clinical and biopsychosocial data. PCPs represent the single largest group of AI end users among health care professionals; therefore, PCPs have the most to gain, as well as the most to lose, in the coming era of innovation. Model developers looking to make the biggest impact on the greatest number of patients and clinicians should look no further than primary care for whom to consult, who to design for, which problems to solve, and where to deploy their solutions.

The 4 Cs of Primary Care

First Contact, Comprehensive, Coordinated, and Continuous—these are the pillars described by Dr.

Barbara Starfield, and they are as relevant today as they were when she proposed them.³⁹ In a tribute article at the time of her passing, Dr. Kevin Grumbach wrote, “In this era of dynamic primary care transformation and redesign, Starfield’s 4 C’s retain an enduring integrity and relevance. What is innovative these days is the means to deliver the core functions of primary care, not the functions themselves.”⁴⁰ Consider each of these pillars—what they mean and why they matter for health care AI (Figure 4).

First Contact: The first contact patients make with the health system is with primary care—whatever it is, PCPs see it first. AI applied upstream has the best chance of making the most impact on people’s lives and also creating positive ripple effects downstream.

Comprehensive: From cradle to grave, from outpatient to inpatient, from delivering babies to minor surgeries—PCPs do almost everything. This means that PCPs are subject matter experts in nearly every domain of medicine and can add value to practically any AI use case in health care. If you are a model developer and you can only have one physician on your team, it should be a PCP.

Coordinated: PCPs are the quarterbacks of the health care team, with unrivaled interconnectedness to every specialty and every part of the health system. This means that PCPs are the key stakeholders for and know all the other stakeholders in essentially every health care AI use case.

Figure 4. The 4 pillars of primary care and why they matter for artificial intelligence in health care.
Abbreviations: AI, artificial intelligence; PCP, primary care provider.

Pillar	What it means	Why it matters for healthcare AI
First Contact	The first contact patients make with the health system is with primary care – whatever it is, PCPs see it first	AI applied upstream has the best chance of making the most impact on people’s lives and also creating positive ripple effects downstream
Comprehensive	From cradle to grave, from outpatient to inpatient, from delivering babies to minor surgeries – PCPs do almost everything	PCPs are subject matter experts in nearly every domain of medicine, and can add value to practically any AI use case in healthcare
Coordinated	PCPs are the quarterbacks of the healthcare team, with unrivaled interconnectedness to every part of the health system	PCPs are the key stakeholders for and know all the other stakeholders in essentially every healthcare AI use case
Continuous	Unlike most other specialties, primary care sees patients over decades – PCPs are in it for the long haul	AI applied alongside continuous rather than episodic care has the best chance of earning patients’ trust, sustaining patient engagement, and improving long-term outcomes

Continuous: Unlike most other specialties, primary care sees patients over decades; PCPs are in it for the long haul. AI applied alongside continuous rather than episodic care has the best chance of earning patients' trust, sustaining patient engagement, and improving long-term outcomes.

PCPs Have Deep Relationships with Patients and Communities

Health care is fundamentally a social enterprise, powered by committed, caring, and collaborative connections between the humans involved. AI can be perceived by some people as a threat to that human connection. Among health care professionals, PCPs are best positioned to assuage patient concerns about AI, build confidence in new technologies over time, and play the role of a trusted champion for both patients and AI.

How Primary Care Will Lead Health Care AI

This final section will consider how primary care might lead and participate in health care AI.

Ensuring That AI Use Cases Are Relevant and Human-Centered

PCPs can lead by working with technology developers to bring their most promising AI solutions from “code to bedside”⁴¹ in support of the Quintuple Aim. It starts with articulating real-world problems and why they are important to patients, clinicians, health systems, and society. Aligning projects around real pain points ensures that technology and workflows are developed to address a genuine need. PCPs can keep industry focused on human-centered solutions, while demanding that AI be designed around augmenting human capabilities and supporting human-driven models of care delivery, rather than replacing human providers or subverting human relationships that lie at the heart of healing.

Using Quality Improvement Methods to Implement AI in Health Care Settings

Traditional research methods often do not support successful AI implementations in health care.⁴² Different from drugs or medical devices, AI technology is highly dependent on the data in a given system and how it is integrated into that system. As discussed previously, AI models trained on data from 1 health system might score well in that system but perform surprisingly poorly at another.

Furthermore, clinical workflows are highly variable between different specialty/service areas within the same health system and even more variable between different health systems and geographic regions. For these reasons, an implementation approach that seeks to understand rather than control real-world variations—based on quality improvement techniques—can help facilitate AI model and workflow design that withstands normal variation. Figure 5 shows a repeatable, mixed method, 6-step approach based on quality improvement methods—skills that are familiar and accessible to PCPs—that can be used to implement AI solutions in health care settings.⁴² Health systems can drive AI innovation by empowering PCPs to engage in translation and implementation work—supported by multidisciplinary teams of improvement experts, clinical informaticists, and data scientists—through new or existing quality improvement programs and institutional value-based funding sources.

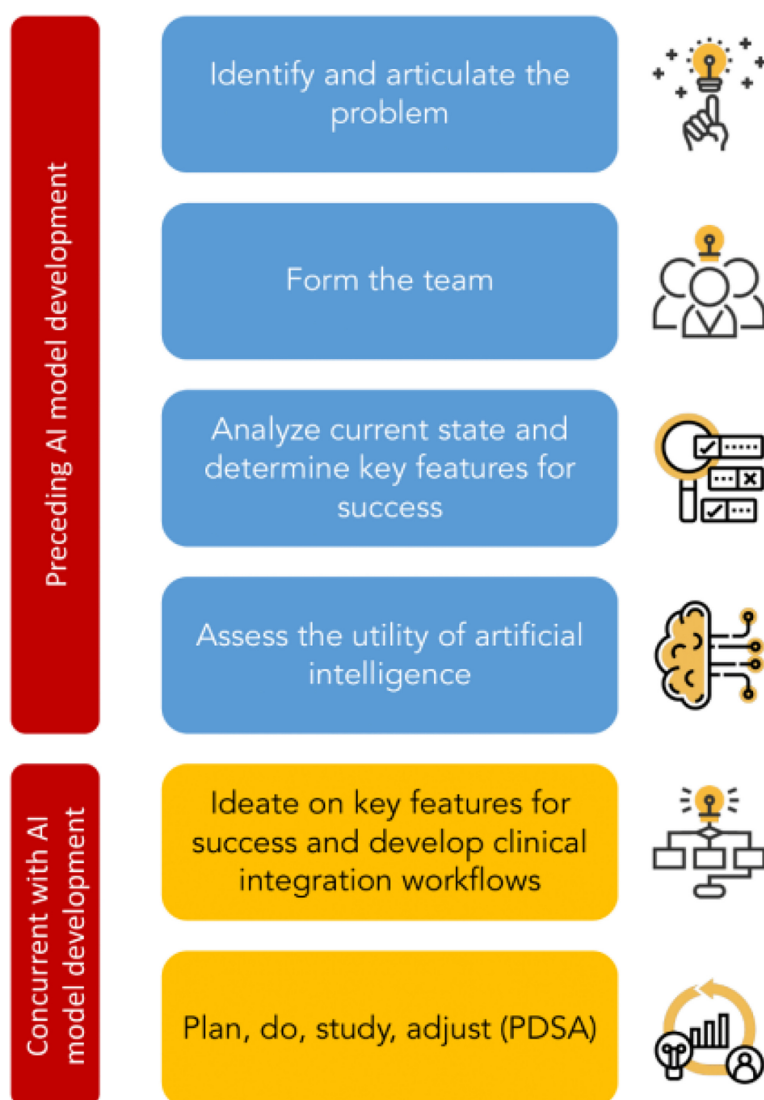
Advocating for Equitable and Accountable AI

PCPs can lead by applying a health equity lens to every AI implementation, asking: “Who might be left behind, or harmed, with this AI solution?” Similar to the idea of patient advisory councils, health systems and technology companies should establish health equity advisory councils to audit and correct race, gender, and socioeconomic biases in their algorithms. There needs to be significantly greater public and private investment in both research and advocacy for equitable and accountable AI in health care. Because a fundamental component of equity in health care AI depends on the integrity, representativeness, and interoperability of data, this vision could be advanced by establishing a nationwide health data architecture and governance framework for AI.⁴³

Education and Training

Readers who wish to learn more about AI equity and accountability can consider these books: *Automating Inequality* by Virginia Eubanks⁴⁴ and *Weapons of Math Destruction* by Cathy O’Neil.⁴⁵ Those who wish to engage in advocacy can consider joining the Algorithmic Justice League,⁴⁶ the Center for Humane Technology,⁴⁷ or Data for Black Lives.⁴⁸ Those who desire more formal training can consider doing the American Medical Informatics Association’s 10 × 10 virtual courses or health informatics certification,⁴⁹ completing an Accreditation Council for Graduate

Figure 5. A 6-step approach (“The Stanford Method”) for implementing artificial intelligence solutions in health care settings. Abbreviation: AI, artificial intelligence.



Medical Education–accredited clinical informatics fellowship, or getting certified by the nascent American Board of Artificial Intelligence in Medicine.⁵⁰

Conclusion

Primary care—as the dominant force at the base of the health care pyramid, with its unrivaled interconnectedness to every part of the health system and its deep relationship with patients and communities—is the most uniquely suited specialty to lead and participate in the health care AI revolution. PCPs can lead by partnering with technologists to ensure that use cases are relevant and human-centered, applying quality improvement

methods to health care AI implementations, and advocating for inclusive and ethical AI that combats, rather than worsens, health inequities.

The author thanks Ms. Grace Hong for transcription support.

To see this article online, please go to: <http://jabfm.org/content/35/1/175.full>.

References

1. Esteva A, Robicquet A, Ramsundar B, et al. A guide to deep learning in healthcare. *Nat Med* 2019; 25:24–9.
2. DeSilver D [Internet]. Chart of the week: the ever-accelerating rate of technology adoption. Pew

- Research Center website; 2014 [cited 2021 March 19]. Available from: <https://www.pewresearch.org/fact-tank/2014/03/14/chart-of-the-week-the-ever-accelerating-rate-of-technology-adoption/>.
3. Perrault R, Shoham Y, Brynjolfsson E, et al. [Internet]. The AI Index 2019 Annual Report. Stanford University Human-Centered AI Institute website; 2019 [cited 2021 March 19]. Available from: https://hai.stanford.edu/sites/default/files/ai_index_2019_report.pdf.
4. Kasparov G. Chess, a Drosophila of reasoning. *Science* 2018;362:1087.
5. Silver D, Hubert T, Schrittwieser J, et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 2018;362:1140–4.
6. Jiang HJ, Russo CA, Barrett ML [Internet]. Nationwide frequency and costs of potentially preventable hospitalizations, 2006. Healthcare Cost and Utilization Project Statistical Brief #72. U.S. Agency for Healthcare Research and Quality, Rockville, MD; 2009 [cited 2021 March 19]. Available from: <https://www.hcup-us.ahrq.gov/reports/statbriefs/sb72.pdf>.
7. Morgenstern JD, Buajitti E, O'Neill M, et al. Predicting population health with machine learning: a scoping review. *BMJ Open* 2020;10:e037860.
8. Yarnall K, Pollak KI, Ostbye T, Krause KM, Michener JL. Primary care: is there enough time for prevention? *Am J Public Health* 2003;93:635–41.
9. Winn AN, Somai M, Fergestrom N, Crotty BH. Association of use of online symptom checkers with patients' plans for seeking care. *JAMA Netw Open* 2019;2:e1918561.
10. Rajkomar A, Yim J, Grumbach K, Parekh A. Weighting primary care patient panel size: a novel electronic health record-derived measure using machine learning. *JMIR Med Inform* 2016;4:e29.
11. McCarthy J [Internet]. One in five U.S. adults use health apps, wearable trackers. Gallup website; 2019 [cited 2021 March 19]. Available from: <https://news.gallup.com/poll/269096/one-five-adults-health-apps-wearable-trackers.aspx>.
12. Stein N, Brooks K. A fully automated conversational artificial intelligence for weight loss: longitudinal observational study among overweight and obese adults. *JMIR Diabetes* 2017;2:e28.
13. Sinsky C, Colligan L, Li L, et al. Allocation of physician time in ambulatory practice: a time and motion study in 4 specialties. *Ann Intern Med* 2016;165:753–60.
14. Lin S. The present and future of team documentation: the role of patients, families, and artificial intelligence. *Mayo Clin Proc* 2020;95:852–5.
15. Smith SV [Internet]. How a machine learned to spot depression. National Public Radio website; 2015 [cited 2021 March 19]. Available from: <https://www.npr.org/sections/money/2015/05/20/407978049/how-a-machine-learned-to-spot-depression>.
16. Savoy M. IDx-DR for diabetic retinopathy screening. *Am Fam Physician* 2020;101:307–8.
17. Liu Y, Jain A, Eng C, et al. A deep learning system for differential diagnosis of skin diseases. *Nat Med* 2020;26:900–8.
18. Lin S, Shanafelt TD, Asch SM. Reimagining clinical documentation with artificial intelligence. *Mayo Clin Proc* 2018;93:563–5.
19. Rabinowitz S, Iroanya E [Internet]. Can machine learning reduce patient, provider and insurer administrative burdens? Healthcare Information and Management Systems Society website; 2019 [cited 2021 March 19]. Available from: <https://www.himss.org/resources/can-machine-learning-reduce-patient-provider-and-insurer-administrative-burdens>.
20. Holdsworth L, Park C, Asch SM, Lin S. The potential use of technology-enabled and artificial intelligence support for pre-visit planning in ambulatory care: findings from an environmental scan. *Ann Fam Med* 2021;19:419–26.
21. Lin S, Sattler A, Smith M. Retooling primary care in the COVID-19 era. *Mayo Clin Proc* 2020;95:1831–4.
22. Buolamwini J, Gebru T. Gender shades: intersectional accuracy disparities in commercial gender classification. *PMLR* 2018;81:77–91.
23. Fowler GA [Internet]. Black Lives Matter could change facial recognition forever—if Big Tech doesn't stand in the way. Washington Post; 2020 [cited 2021 March 19]. Available from: <https://www.washingtonpost.com/technology/2020/06/12/facial-recognition-ban/>.
24. Coley RY, Johnson E, Simon GE, Cruz M, Shortreed SM. Racial/ethnic disparities in the performance of prediction models for death by suicide after mental health visits. *JAMA Psychiatry* 2021;78:726–34.
25. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 2019;366:447–53.
26. Hammond G, Johnston K, Huang K, Joynt Maddox KE. Social determinants of health improve predictive accuracy of clinical risk models for cardiovascular hospitalization, annual cost, and death. *Circ Cardiovasc Qual Outcomes* 2020;13:e006752.
27. Unfairness by algorithm: distilling the harms of automated decision-making [Internet]. Future of Privacy Forum, Washington, DC; 2017 [cited 2021 March 19]. Available from: <https://fpf.org/wp-content/uploads/2017/12/FPF-Automated-Decision-Making-Harms-and-Mitigation-Charts.pdf>.
28. Gender Shades website [Internet]. MIT Media Lab; 2018 [cited 2021 March 19]. Available from: <http://gendershades.org/>.
29. Wong A, Otlés E, Donnelly JP, et al. External validation of a widely implemented proprietary sepsis

- prediction model in hospitalized patients. *JAMA Intern Med* 2021;181:1065–70.
30. Kashyap M, Seneviratne M, Banda JM, et al. Development and validation of phenotype classifiers across multiple sites in the observational health data sciences and informatics network. *J Am Med Inform Assoc* 2020;27:877–83.
31. Davis SE, Lasko TA, Chen G, Matheny ME. Calibration drift among regression and machine learning models for hospital mortality. *AMIA Annu Symp Proc* 2017;2017:625–34.
32. The Social Dilemma website [Internet]; 2020 [cited 2021 March 19]. Available from: <https://www.thesocialdilemma.com/>.
33. Ledger of Harms [Internet]. Center for Humane Technology website; 2020 [cited 2021 March 19]. Available from: <https://ledger.humanetech.com/>.
34. U.S. Capitol Riot [Internet]. New York Times; 2021 [cited 2021 March 19]. Available from: <https://www.nytimes.com/spotlight/us-capitol-riots-investigations>.
35. Vosoughi S, Roy D, Aral S. The spread of true and false news online. *Science* 2018;359:1146–51.
36. Dwoskin E [Internet]. Misinformation on Facebook got six times more clicks than factual news during the 2020 election, study says. *Wall Street Journal*; 2021 [cited 2021 Sept 6]. Available from: <https://www.washingtonpost.com/technology/2021/09/03/facebook-misinformation-nyu-study/>.
37. Taneja H [Internet]. The era of “move fast and break things” is over. *Harvard Business Review* website; 2019 [cited 2021 March 19]. Available from: <https://hbr.org/2019/01/the-era-of-move-fast-and-break-things-is-over>.
38. Lin S, Mahoney MR, Sinsky CA. Ten ways artificial intelligence will transform primary care. *J Gen Intern Med* 2019;34:1626–30.
39. Starfield B. Primary care: concept, evaluation, and policy. New York: Oxford University Press; 1992.
40. Stange KC. Barbara Starfield: passage of the pathfinder of primary care. *Ann Fam Med* 2011;9:292–6.
41. Stanford Healthcare AI Applied Research Team [Internet]. Stanford Medicine website; 2021 [cited 2021 March 19]. Available from: <https://med.stanford.edu/healthcare-ai>.
42. Smith M, Sattler A, Hong G, Lin S. From code to bedside: implementing artificial intelligence using quality improvement methods. *J Gen Intern Med* 2021;36:1061–6.
43. Matheny M, Thadaneys Israni S, Ahmed M, Whicher D. Editors. 2019. Artificial intelligence in health care: the hope, the hype, the promise, the peril. Washington (DC): National Academy of Medicine.
44. Eubanks V. Automating inequality: how high-tech tools profile, police, and punish the poor. New York: St. Martin’s Press; 2018.
45. O’Neil C. Weapons of math destruction: how big data increases inequality and threatens democracy. New York: Crown Books; 2016.
46. Algorithmic Justice League website [Internet]; 2021 [cited 2021 March 19]. Available from: <https://www.ajl.org/>.
47. Center for Humane Technology website [Internet]; 2021 [cited 2021 March 19]. Available from: <https://www.humanetech.com/>.
48. Data for Black Lives website [Internet]; 2021 [cited 2021 March 19]. Available from: <https://d4bl.org/>.
49. American Medical Informatics Association website [Internet]; 2021 [cited 2021 Sept 6]. Available from: <https://amia.org/>.
50. American Board of Artificial Intelligence in Medicine website [Internet]; 2021 [cited 2021 Sept 6 19]. Available from: <https://abaim.org/>.

