

AI in Healthcare

Setting a Common Perspective

artificial intelligence

[JANUARY 18, 2018]



FOUNDED BY BRIGHAM AND WOMEN'S HOSPITAL
AND MASSACHUSETTS GENERAL HOSPITAL

WORLD MEDICAL
INNOVATION
FORUM™



The 2018 World Medical Innovation Forum focuses on the advancements and opportunities of AI in healthcare. This document aims to provide participants with a common framing of AI in healthcare.

Defining Artificial Intelligence

If intelligence is understood to mean traits that enable humans to independently perceive, contextualize, interpret, and learn from information in a way that can be applied towards adaptive behaviors resulting in actions or communication, then Artificial Intelligence (AI) is simply one or more of these traits manifest in a machine in order to achieve a given goal.

That goal may be to:

- **replicate** what a human could otherwise do but in a way that is preferable for reasons such as improved productivity, quality, or safety
- **surpass** humans in tasks that humans cannot feasibly do such as to make optimal decisions with information that is too complex for humans to be able to absorb, interpret, or act on under real world conditions, or
- **augment** human intelligence to produce a capability that is superior to either the machine or the human alone.

Machine Learning (ML) is often used interchangeably with AI but more accurately it is a form of AI that relates to the capability of a machine to improve its performance on a task by iteratively repeating the task, assessing its own performance relative to a desired outcome, and making adjustments as needed. It is a subset of AI, albeit an important one. Other possible attributes of AI include natural language processing, spatial navigation, machine vision, logical reasoning, and pattern recognition. Of course many of these are used in combination with each other and ML is a common component of many AI systems.

Scope of AI

There is a paradox in AI that once a machine achieves a form of intelligence, that capability is then no longer considered by many to be a form of AI. This idea was pithily captured in Tesler's Theorem.

Optical Character Recognition (OCR) for example demonstrates one of the hallmarks of intelligence, namely perception. But as OCR has improved and become embedded in our machines, it is now thought of as software feature and not as AI.

Avoiding this paradox, the scope of the definition of AI at the Forum is broad and includes not just cutting edge innovations but also existing AI capabilities as they relate to healthcare.

“Intelligence is whatever machines haven't done yet.”

LARRY TESLER, 1970



AI in healthcare

AI in healthcare is being driven by three trends. First there has been an exponential growth in useful health related data including from EHRs, personal monitoring devices, genomic information, social information, diagnostic information, and more. The diversity and complexity of these data necessitates the use of AI.

Second, there is continuing growth in computational power driven by advances in processor design and computing architectures, and in AI methodologies, most notably relating to deep neural networks. These enable the use of AI.

Finally, there is a broad recognition that AI can address many of today's health care challenges across cost, access, and quality of the system. These include the high cost of finding new medicines, workforce imbalances, variation in care and outcomes, and more.

This confluence of ability and need is driving investment and innovation resulting in a positive feed-forward loop. Consequently, over the last few years a large number of healthcare AI applications have emerged: from finding new uses for old drugs to discovering new drugs; from guiding clinicians to personalized care pathways to engaging patients on their healthcare with automated bots; from identifying patients at risk by combining all kinds of data about them including unstructured data and social information to unearthing fraud in provider claims.

Most AI applications in healthcare including the examples above may be grouped into one of five contexts: R&D, care delivery, patient engagement, population health, and administration, and in each the goal is either to replicate, to surpass, or augment human intelligence.

The boundaries between these categories are fluid and applications may move from one to another or occupy more than one slot. The following are illustrative examples.

	Replicate	Surpass	Augment
R&D	In-silico drug screening and optimization	Discovering new drugs based on pattern analysis in published literature	Predictive tools to improve clinical trial recruitment
Care delivery	Machine vision to interpret radiology images	Optimize patient triage in the ER	Complex clinical decision support
Patient engagement	Clinical bot for at-home diagnosis	Patient monitoring and real-time interventions	Elder care robot to support care plan and rehabilitation
Population health	Personalized and contextualized behavioral nudges	Early warning of epidemics	Patient risk stratification using multiple datasets
Administration	Automation of prior authorizations	Detection of fraud and misuse	Optimization of patient scheduling and staff resourcing

Societal questions

Like any other enabling technology, AI creates new questions for society. These questions are particularly acute in healthcare given the vulnerability of patients and the desire to protect privacy. How society responds to these questions and how that is manifest in the actions of legislators and regulators will make a significant difference to how quickly and how well we are able to benefit from AI innovations in healthcare.

Four questions for us to consider as we approach the 2018 Forum.

In what way will AI change what humans do in healthcare?

The pace of development of the underlying enabling AI and the applications across the landscape above varies dramatically. An interesting observation made by Hans Moravec and others is that what most lay people intuitively think is easy or hard for machine intelligence is wrong.

When does it matter if intelligence is artificial?

The level of realism that can be achieved by AI agents in remote interactions is rapidly reaching the point that is hard to distinguish from a human. This has far-reaching second order implications in healthcare such as relating to the patient-clinician relationship and for payment models.

The well-known Turing test assesses a machine's ability to interact using natural language conversations in a way that is indistinguishable from a human. In 2014 a software program over a five-minute text conversation convinced 10 of 30 judges that it was a teenage boy and there have been other successes as well.

Bots are not only getting better at parsing questions and finding the right answers, they are also getting better at demonstrating human like qualities such as typos, delays, and idiosyncrasies.

Do we need to know why?

The most successful ML approaches in the form of deep networks can behave like black boxes, in the sense that while we have a statistical sense of how inputs relate to outputs, it is often not possible to precisely determine why a particular input generates a particular output.

While this is important in other industries as well, it is particularly important in healthcare. Consider, for instance, if a patient dies after following AI guidance, not knowing why a particular pathway was recommended will confound issues in the ethical, legal, and regulatory spheres.


What does AI mean for privacy and legal responsibility?

Information collected by and inferences made by an AI embedded in a provider system will be covered by the usual patient data protections. However, similar systems may well begin to exist outside the traditional healthcare system. If a non-clinical AI can infer the health status of an individual with a high degree of accuracy based on non-protected information, that has implications for privacy.

“The main lesson of thirty-five years of AI research is that the hard problems are easy and the easy problems are hard. The mental abilities of a four-year-old that we take for granted – recognizing a face, lifting a pencil, walking across a room, answering a question – in fact solve some of the hardest engineering problems ever conceived... As the new generation of intelligent devices appears, it will be the stock analysts and petrochemical engineers and parole board members who are in danger of being replaced by machines. The gardeners, receptionists, and cooks are secure in their jobs for decades to come.”

STEVEN PINKER





World Medical Innovation Forum **Shaping AI in Healthcare**

Harnessing this broad range of AI opportunities in healthcare and addressing the complex associated societal questions requires a multi-disciplinary approach. The Forum is an opportunity to influence the shape of the future of AI in healthcare by tapping into and contributing to the perspectives of clinicians, investigators, healthcare and life-science executives, HIT and AI experts, regulators, and investors.