

Artificial intelligence: what's possible, why now?

David Wild

July 2018—When it comes to artificial intelligence, it can be difficult to distinguish hyperbole from reality. So how much can AI truly replace human tasks in society and, more specifically, in medicine?



Dr. Singh

Ajit Singh, PhD, who spoke at the Executive War College in May on what's feasible today and the use of AI in medicine, said answering that question requires understanding what drives growth in AI—the discipline of how to make computers do things at which, for now, people are better—and the limits of its implementation.

Despite advances in AI, there will likely always be differences between “us and them,” Dr. Singh said of humans and computers. While physicians need to adapt to AI, the technology will “absolutely not” replace them.

“Emotion, understanding, consciousness, and creativity—you will never be able to hire a robot with these emotional quotient elements,” he said.

One of the most important drivers of AI and technological innovation is diversity, and nature can teach us a lot about its role in the growth of AI, said Dr. Singh, a partner at Artiman Ventures, a company in Palo Alto, Calif., focusing on early-stage technology and life science investments, and a professor in the Stanford University School of Medicine. Prior to joining Artiman, Dr. Singh was president and CEO of BioImagene, a digital pathology company acquired by Roche. Before that, he was with Siemens for nearly 20 years.

With the largest diversity of species on the planet, the Galapagos Islands provide an instructive example of how innovation can take place, Dr. Singh said. One reason the islands are so diverse is they offer a “remarkable opportunity for mobility.”

“Three ocean currents meet around the Galapagos Islands’ coastlines, bringing in and mixing life forms and large molecules from different parts of the planet.” With so many different life forms gathering in one location, “genetic experiments” happen constantly and organically, Dr. Singh said. In some cases, the new life forms fail; in other cases they succeed.

As in nature, human innovation is also based on diversity, but of a different type. “What helps us innovate is cognitive diversity,” he said, “or transdisciplinarity.” The latter differs from multidisciplinarity, which means “bubbles” of multiple disciplines that often work separately. “Transdisciplinarity is when disciplines intersect, when they’re able to discuss and discard failures,” Dr. Singh explained. “That’s when innovation happens.”

The origin of AI dates to 1950, when British scientist Alan Turing proposed the idea of a thinking machine (*Mind*. 1950;59:433-460).

Top 10 AI applications in health care

- Robot-assisted surgery
- Virtual nursing assistants
- Administrative workforce assistance
- Fraud detection
- Dosage error reduction
- Connected machines
- Clinical trial participant identifier
- Preliminary diagnosis
- Automated image diagnosis
- Cybersecurity

His “Imitation Game,” later known as the Turing test, involved posing a series of logical questions to a human and a machine, both of which would be hidden from view. If the human interviewer could not distinguish between machine and human, the machine would be said to have passed the test and could be considered a thinking machine.

Other early AI innovators included Dietrich Prinz of the University of Manchester, who developed the first chess-playing program, and Allen Newell and Herbert Simon, who in 1955 created the Logic Theorist, a computer program that was able to prove 38 of the first 52 theorems in *Principia Mathematica*, a cornerstone text of mathematical thinking.

These and other critical developments in computing led John McCarthy to coin the term “artificial intelligence” in 1956 at the Dartmouth Summer Research Project on Artificial Intelligence.

With a history stretching back more than 60 years, why is there so much interest and growth in AI today?

“The basic ideas of AI and how to implement AI have not changed much,” Dr. Singh said. “The connection networks built in the 1980s are still the same, and the math has not changed. So what has changed?” Data, he said, which is now in abundance. “This has been the most critical thing we needed for AI to evolve to this point.”

Multiple academic disciplines now collaborate in the development of AI, he said, citing philosophy, mathematics, economics, neuroscience, psychology, cognitive science, computational engineering, control systems, and linguistics. “In order to make AI implementable, we had to have that transdisciplinarity coupled with large amounts of relevant data, and it took us several decades to get there,” Dr. Singh said.

The two dominant implementations of AI are symbolic and connectionist, Dr. Singh said, each with its strengths and limitations.

In the symbolic implementation, algorithms are used to mimic human intelligence. A connectionist implementation uses neural networks to build knowledge “from the bottom up.” The capacity of a machine to teach itself is limited by the complexity of its neural networks and by the amount of data it can learn from, Dr. Singh said.

Currently, one central processing unit can contain up to roughly 10^6 transistors, an amount that can facilitate about 1 billion simultaneous operations. A single machine can contain between 1,000 and 10,000 CPUs as well as up to 10^9 bits of random-access memory, leading to a cycle time of 10^{-8} seconds, or 10 nanoseconds, to complete a task. In contrast, the human brain has 10^{12} neurons, 10^{14} synapses, and a cycle time of 10^{-3} seconds, or one millisecond, to fire.

While we will soon have computer systems with as many “neurons” as the human brain, we are not likely to have computers with as many synapses, or interconnectivity, as the human brain, “not in the near future,” Dr. Singh said. Even if a human-level degree of neural interconnectivity could be built, it is unlikely that computers would achieve human-level AI.

"We don't learn everything from scratch after birth. There's a lot of genetic hardwiring that has taken place over several million years of evolution that has taught us to do certain things really well," he said.

For example, humans can understand context, while computers at present cannot. That limitation presents a problem with an application like speech recognition, Dr. Singh said, noting that AI systems can understand individual words from small vocabularies with 99 percent accuracy but cannot truly understand speech.

"If I were in a restaurant and hurriedly said, 'I'd like my coffee with dream and sugar,' most waiters would bring me cream and sugar, even though I said dream and sugar," he said. "Context is extremely important, and unless we're able to input a lot of context into AI systems, no matter how much vocabulary we input, it won't be able to understand speech."

Computers also have a limited capacity to perceive, unlike humans, who can identify patterns, understand scenes, and recognize objects relatively well even with poor lighting and with objects occluding their line of sight.

"If you had to recognize individual faces as they were walking into a conference room one by one, you could likely do that with very high accuracy, and even if you had a cluttered environment where you had partial occlusions and different lighting, the success rate would still be high, whereas for an AI-based face recognition system, the success rate would drop to 30 to 60 percent," Dr. Singh said.

Computers can identify images in a constrained environment, but they do a poor job at understanding a scene in an unconstrained environment, he said. "If you had to locate someone by looking at the back of their head in a picture taken from behind a crowd of people, you would still be able to do this much of the time," Dr. Singh said. "The accuracy of a computer in a task like this would be below 10 percent."

This limited capacity for pattern recognition in the absence of full context limits AI's role in medicine. Radiologists and pathologists did not just learn to analyze images—like other humans, "they learned to recognize patterns through 50 million years of evolution, four years of residency, years of fellowship, and maybe some life experience," Dr. Singh said.

Mimicking the impact that evolution has had on the perceptive abilities of humans would require that the computer analyze 7 billion patients' worth of data to diagnose breast cancer alone with as much accuracy as a domain expert. "Providing that amount of data is not feasible," Dr. Singh said. A connectionist-only approach with no a priori knowledge from humans will never work, especially for a problem that has an inherent high dimensionality.

Allowing the computer to learn from scratch by observing has its pitfalls as well. "If a computer were placed at a traffic light in New York City and learned the rules of traffic lights by observing human behavior, it would learn that red means stop, green means keep going, and yellow means speed up," Dr. Singh said. "It's no different with diagnosing breast cancer. The system will only learn the average of all humans. But who wants to be an average physician or be diagnosed or treated by an average doctor?"

Selectively inputting data from only the best physicians is an option, but that body of data would be too small, Dr. Singh said.

"It's a catch-22."

Despite these limitations, there will be successful applications of AI in medicine, including in time IBM Watson, for which the initial expectations were unrealistic, he says. In Silicon Valley alone, there are now about 40 startups that address pathology exclusively, in areas ranging from molecular profile analysis to NGS-based liquid biopsies to image analysis for breast, brain, and prostate cancer. If a startup dies, he says, someone else will pick up the asset and make something of it.

In medicine, there is plenty of AI activity to be learned from. The AI successes are modeling and capturing *existing* knowledge and making it available. Examples are arrhythmia recognition from electrocardiograms, coronary heart disease risk-group detection, monitoring prescription of restricted-use antibiotics, and early melanoma diagnosis.

The latter is the use of AI to classify skin lesions (Esteva A, et al. *Nature*. 2017;542:115–118).

The applications that are modeling and learning net *new* knowledge to improve human performance are promising, he said, among them Cellworks for oncology therapy selection, Genxsys decision support for GPs, and Zebra Medical for radiologists. “If you want to build net new knowledge, that can be done,” Dr. Singh said, “but you have to constrain the environment. You can’t create a generalist out of it, but you can create a specialist or super-specialist.”

Breast cancer is an unconstrained problem, he said. “Too much complexity, too many dimensions, the number of biomarkers, the type of images, family history, and so on.” Pneumonia and melanoma are more constrained, for example. “If the data you have available is large and the dimensionality is low,” you have an AI solution, he said.

In the near term, Dr. Singh predicted the AI industry will focus on 10 broad categories of health care applications. (See “Top 10 AI applications in health care.”)

Three virtual nursing assistant applications—Sense.ly, Tavie, and Ada—are working in the United States, England, and Canada, Dr. Singh said. Chatbots are and will remain another popular application of AI, with apps like Babylon Health, which mixes AI and live physician interaction through video and text. The United Kingdom’s National Health Service adopted the app as an alternative to the NHS’ 111-telephone helpline, which patients call for health care advice and to be directed to local or after-hours medical services. “Now, when you call in for a basic triage at the NHS, you’ll interact with a chatbot first,” Dr. Singh said.

Predicting the course of AI in the long term is more difficult, but one thing it will likely do is produce “a patient of the future that is different from the patient of today,” Dr. Singh said. “The patient of the future will be wired with devices that provide a lot of data and allow us to pick up trend lines.”

The evolution to a wired human has begun, in fact, with instruments that patients can attach to their mobile devices to examine their mouth, throat, eyes, heart, lungs, skin, and body temperature (TytoCare) to help clinicians diagnose a variety of conditions remotely, and mobile apps to aggregate data and monitor patients’ emotional health and alert caregivers when symptoms are problematic (Ginger.io).

Care delivery will also change as AI improves and proliferates.

“Device cameras will have very high resolution, systems will use natural language processing, and we will have robotics in our ecosystem.”

Among the recent innovations that have changed how care is provided are technologies that use retinal self-imaging to help clinicians detect illnesses ranging from glaucoma to multiple sclerosis (eyeSelfie, Rimokon) and an app for monitoring wound healing postoperatively (Parable).

Advanced robotics, more sophisticated learning systems, analytics, and communication speed will all improve health system operations, Dr. Singh said. Current applications include a system that helps prevent readmissions and save money by using predictive analytics to optimize where frontline staff are deployed (Care at Hand), autonomous robots that transport materials and clinical supplies throughout a facility (Aethon Tug Robot), and “smartglasses” that make it possible for clinicians to see high-definition images of the vasculature as they insert an intravenous device.

For precision medicine, the tools available or in development include the world’s largest database of human genotypes (Human Longevity), a cell culture platform that mimics the architecture and physiology of the human liver (LiverChip), and nanorobots built from designer DNA that deliver drugs to specific cell types (Wyss Institute).

Pathologists and laboratories can take immediate steps to integrate AI into their operations and care and get ahead of the AI curve, Dr. Singh said.

“Start digitizing and indexing your data now,” he recommended. “Even if you don’t use digital pathology, even if you never want to read on screen, start digitizing because those data will be critical someday.”

He also suggested storing correlated clinical data. Electronic health records store longitudinal data. “As newer, practical implementations of AI become available over the next decade, we will be able to do a lot more with the longitudinal data than we currently do. Hence, we must start storing correlated data—images, clinical data, and all—longitudinally. So run an experiment and say, ‘When can you get the correlated clinical data for the current images?’”

Another is to start a pilot project to verify the accuracy of coding. “Lab administrators often complain about coding being a problem, leading to claim rejections,” he said. “AI can already help with that now. There are systems that work rather well.”

Another pilot to implement: using chatbots for genetic counseling. “There are not enough genetic counselors,” he said, “and as much as 80 percent of their work can be done by a chatbot.”□

David Wild is a writer in Toronto.