CONTENT PROVIDED BY

ADTMAG
APPLICATION DEVELOPMENT TRENDS

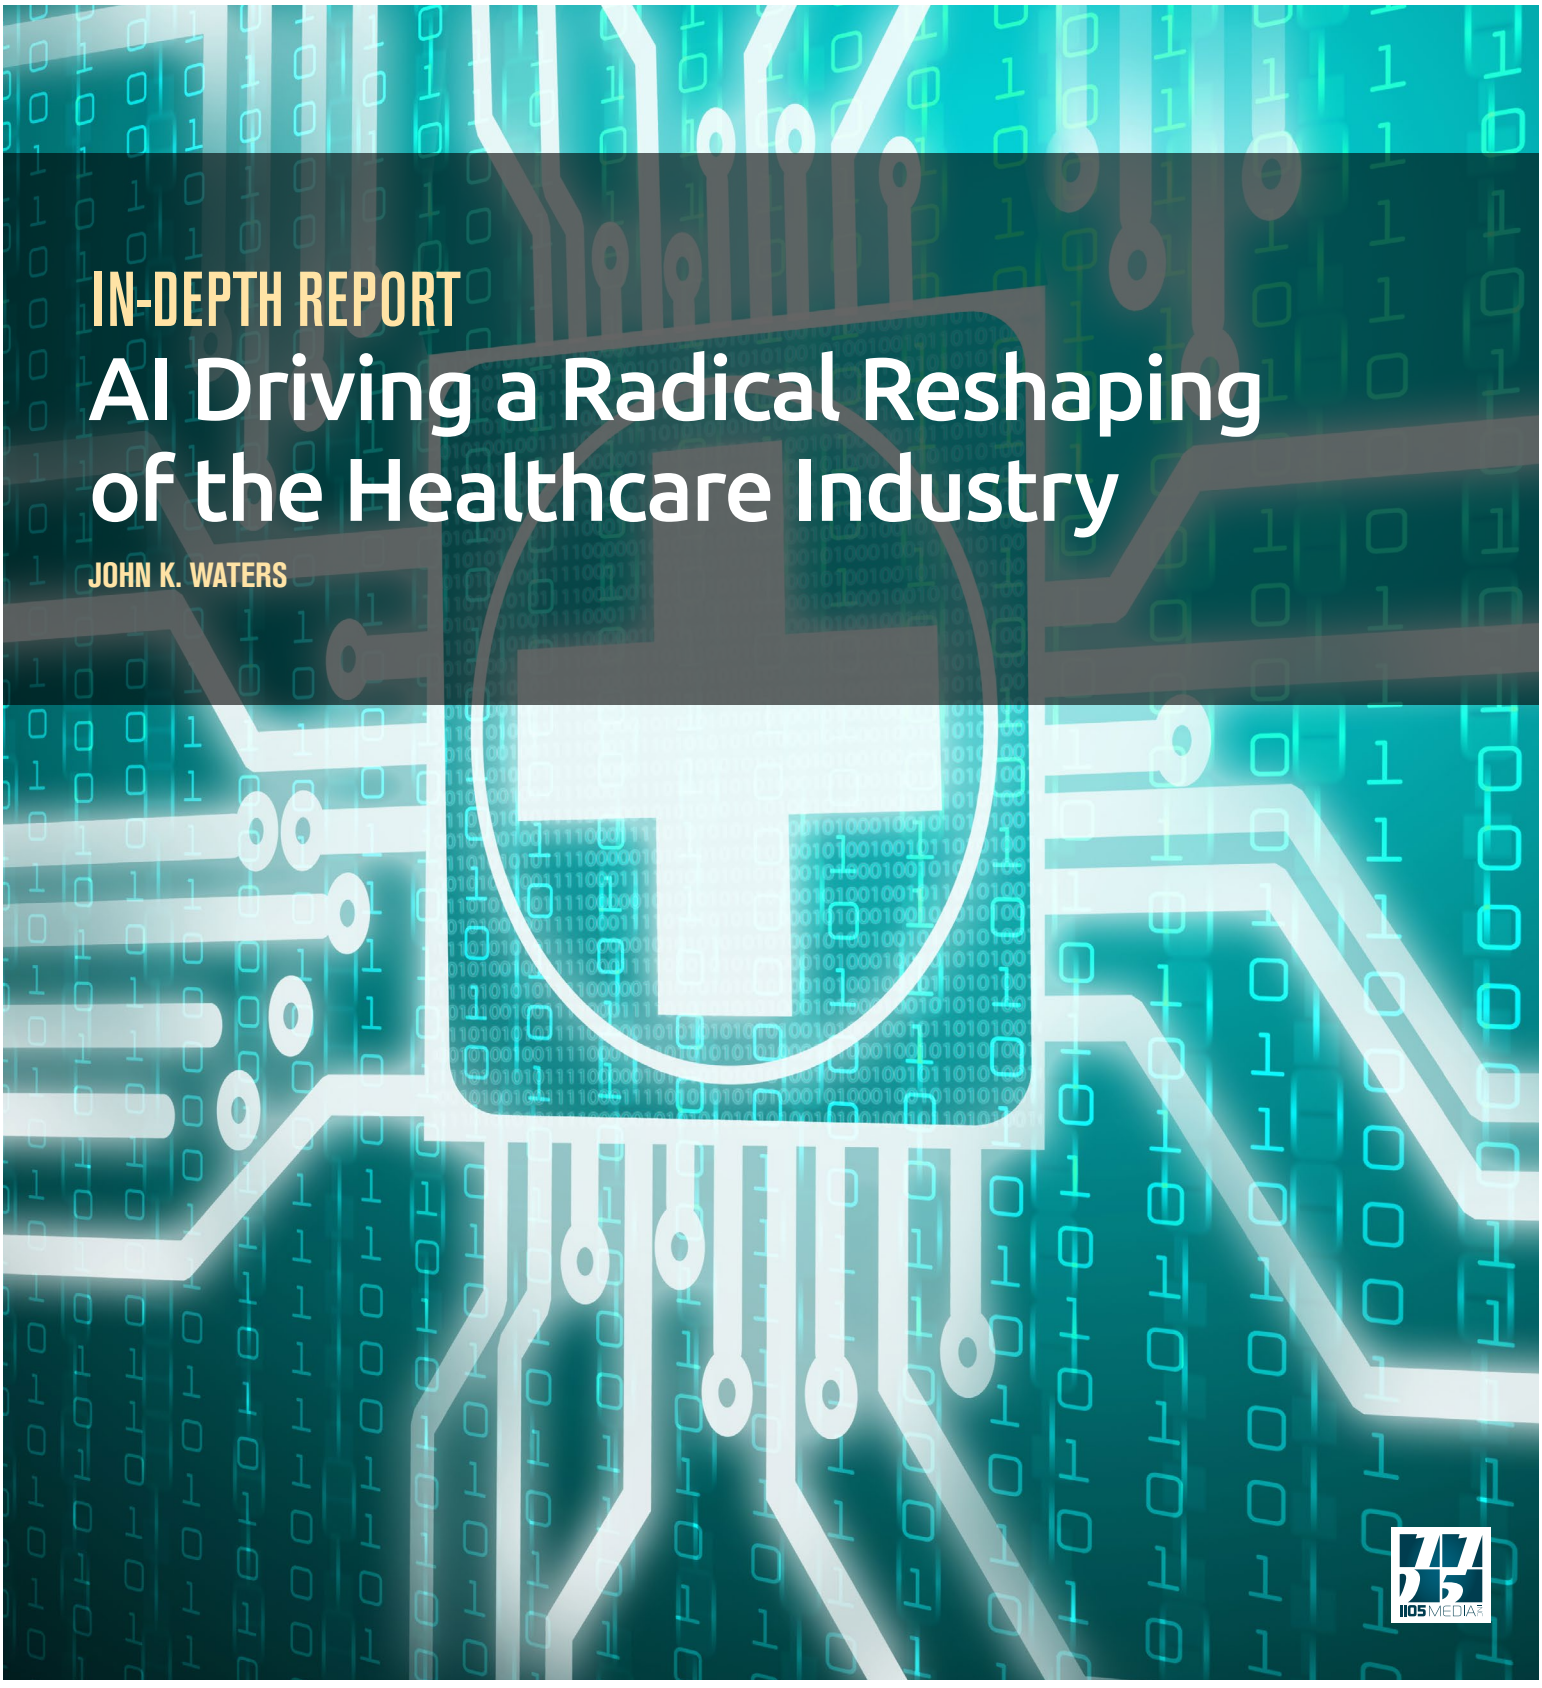

IN-DEPTH REPORT

AI Driving a Radical Reshaping of the Healthcare Industry

JOHN K. WATERS

1105 MEDIA

Executive Summary

The combination of Artificial Intelligence (AI) and Big Data is driving a radical reshaping of the healthcare industry. AI and cognitive initiatives in a growing number of organizations are uncovering the true value of vast and ever-expanding stores of digital health data. New systems and software are appearing that enable complex algorithms, deep analytics, and predictive capabilities to tap these data stores and deliver better health outcomes, while dramatically increasing business and administrative efficiencies.

Industry watchers are expecting the AI-for-healthcare market to exceed $6 billion over the next few years, and investments are pouring in from traditional high-tech companies, pharmaceutical and life sciences firms, and an ever-expanding field of startups. From image recognition systems that read X-rays better than radiologists to smartphone apps that speed up the verification of insurance coverage, AI-based healthcare solutions are emerging at a rapid pace.

New systems and software are appearing that enable complex algorithms, deep analytics, and predictive capabilities to tap data stores and deliver better health outcomes.

But AI-for-healthcare is presenting CIOs and IT leaders with a unique set of challenges. These are disruptive solutions that will, at some point, provide core, industry-standard capabilities that must be adopted, but the current generation is just beginning to prove its value. And in healthcare, there's a lot more at stake than in other industry verticals.

Business decision makers don't need to be experts in the cluster of technologies that comprise AI, but they do need to understand what they are, what they can do, and what they can't do to accurately assess their value to their organizations and use their capabilities successfully.

Definitions

Artificial Intelligence: The term "artificial intelligence" was coined originally to describe machine intelligence, as opposed to the natural intelligence of human beings. But true AI has yet to be achieved, and the term is commonly applied today to a broad category of applications that mimic human cognitive functions—things like learning, pattern recognition, and problem solving.

"It's worth keeping in mind as you're making decisions about how you're going to invest in these technologies that there is no true intelligence in them," says Jay Swartz, Chief Science Officer at Blackbox AI, better known as the Watson Whisperer. "We're heading there, and we're getting close, but we're not there yet. These applications operate at a surface level of the

data: they can identify a car or a chicken, recognize that a chicken is a living thing and a car is a piece of equipment. But they don't have fond memories of riding around in the back seat of a convertible or gathering eggs from the hen house, which come with the way human beings understand cars and chickens."

Machine Learning: The solutions the marketplace is currently calling "AI" are, in fact, based on Machine Learning (ML), a data analytics technique that uses algorithms to parse data, learn from it, identify patterns, and make decisions with minimal human intervention. ML is about getting computers to learn and improve with experience without being explicitly programmed. For practical purposes, it's the current state of the art in AI, and the terms are virtually synonymous.

Deep Learning: A subset of ML, Deep Learning (DL) employs a collection of algorithms that utilize layered networks, called neural nets, to carry out the processes of machine learning. DL is the key to many of the advanced capabilities in healthcare-related AI solutions.

Machine Learning (ML) is about getting computers to learn and improve with experience without being explicitly programmed.

Both approaches are made possible by computers that are fast enough to support the "training" of large neural nets using extremely large data sets; large-scale, interconnected computational layers; and compute- and data-intensive heuristics and numerical optimization techniques. Modern graphics processing units (GPUs) and other processor advances, for example, make parallel processing faster, cheaper, and more powerful.

Natural Language Processing: Natural Language Processing (NLP) is an area of AI that has already had an impact on healthcare. NLP is the automatic manipulation of natural language, such as speech and text, by software—in other words, it's a way for humans to communicate with intelligent systems without typing. NLP enables voice-activated interactions, chatbots, and voice-to-text solutions. This is a technology advanced and adopted out of necessity, it could be argued, to help facilitate the transition from paper-based records for healthcare providers.

Cognitive Computing: Another term that finds its way into conversations about AI is cognitive computing, which refers to systems that simulate human thought processes. These are self-learning systems that use ML algorithms for data mining, pattern recognition, and natural language processing to solve problems without human intervention. The most famous cognitive computing system as of this writing is IBM's Watson platform, which relies on DL algorithms to process information by

One of the primary benefits of AI in healthcare is its ability to turn data overload into information that is clinically relevant and/or actionable.

comparing it to a teaching set of data. The more data it is exposed to, the more it learns, and the better it gets.

Capsule Networks: And then there's the recent talk about Capsule Networks, a new DL innovation developed by Geoffrey Hinton, known as the Father of Deep Learning. A "capsule" is a set of neural layers nested inside a single layer, resulting in what is being called a "CapsNet." Hinton just published a paper on this concept last year, but the AI community is very excited about it and because of the warp-speed pace of innovation in this space, healthcare CIOs and IT leaders should add it to their lexicons.

Drinking from a Firehose

The amount of patient data available to medical professionals today is truly staggering, but that data is only useful if it can be sifted, sorted, refined, and applied—tasks that are effectively beyond human capability. One of the primary benefits of AI in healthcare is its ability to turn this data overload into information that is clinically relevant and/or actionable. In fact, all these technologies require access to vast amounts of data to "learn" and get smarter about the specific domains in which they are applied.

Some healthcare-oriented AI solutions are initially fed data from medical text books, hospital care guidelines, and the details of millions of clinical cases. But ultimately, they must tap into electronic medical records (EMRs). The list of data types stored in EMRs today includes exam and procedure reports, lab values, pathology reports, imaging data, and data that is automatically downloaded from implanted devices—not to mention admission, discharge, and transfer information, and billing information.

And that firehose of information is only getting stronger as the Internet of Things (IoT) becomes an increasingly important piece of the new health-care puzzle. New generations of apps and wearable devices, and solutions that allow patients to upload their own data and images to their EMRs—things like photos of a healing wound taken with a smartphone camera, blood sugar levels from home monitoring devices, weight logs, and activity trackers—are already amplifying the healthcare data stream.

"As you put together your AI strategy, you must include consideration of the impact of IoT," says Ahmed Banafa, lecturer in San Jose State University's engineering department. "We're seeing new products and services every day creating new business models, and especially relevant in healthcare, new ways of engaging customers."

A Team Sport

AI will not be diagnosing patients or replacing doctors any time soon, even though some solutions are outperforming humans in certain narrow applications (image recognition in radiology, for example). AI should be thought of as support technology—highly sophisticated support technology, to be sure, but tech that's meant to provide medical professionals with heretofore unavailable insights from Big Data for their evaluation, and sometimes to catch mistakes before they happen.

The key concept here is something Swartz calls the "human-in-the-loop model," in which a medical professional is always considered as an element of the AI solution.

"We're not ready to rely on AI 100 percent," he says. "You must have a trained professional in the mix, checking the AI's work, so to speak. Eventually, these systems are going to be so good that that role won't seem as essential as it does today. But for now, you need an actual person woven into the implementation and application of this technology, always."

The key concept is something Jay Swartz calls the "human-in-the-loop model," in which a medical professional is always considered as an element of the AI solution.

The implementation of an AI solution in a healthcare organization—or any organization—should be considered a "team sport," Swartz says. This is the place where the IT professionals must work closely with the business decision makers. As demand grows for AI implementations, so will demand for "translators," businesspeople with strong technical backgrounds who can bridge the gaps between biz and tech.

"This is disruptive technology that many people—doctors, nurses, administrators, patients—are not yet ready to trust," says Banafa. "Which is understandable. In healthcare, we're talking about life-and-death situations, patient privacy and security, and big cultural shifts. It's critical to keep these concerns in mind as healthcare organizations role out their AI projects."

Use Cases

One of the biggest challenges business decision makers face when evaluating AI- solutions in general is the pace at which they are emerging. "We're at a very interesting point in the evolution of Machine Learning," says Swartz. "The discoveries and innovations are coming faster than our ability to assimilate them. And deep learning is advancing so quickly new techniques are appearing almost weekly. It's not uncommon now for a project to start with a set of components that effectively become obsolete before it's completed."

AI applications sift through mountains of medical records and images to detect rare and dangerous diseases sooner than any human could.

Which is to say, AI-in-healthcare is a work in progress. In fact, many of the leading providers have yet to release commercial versions of their solutions. And yet, it would be hard to find a healthcare organization not being affected in some way by the advent of AI, either clinically or administratively, as a growing number of providers begin offering solutions in a range of product categories:

Medical Records Management

The management of electronic medical records (EMR) isn't the kind of AI application that grabs headlines, but it is proving itself in several organizations through solutions that collect, store, reformat, and trace patient data to provide faster, more consistent, and more cost-effective care. If the essential problem facing modern physicians is their only-human ability to sift through an inhuman amount of patient information, these solutions could prove to be game changing.

Predictive Analytics

AI applications that synthesize clinical, labs, demographic, and behavioral data to predict how likely a patient will be to be readmitted to a hospital and/or to prevent patient emergencies are showing promise. Not only do these apps have the potential to improve the quality of care, they could provide patients with a clearer picture of their health.

Early Diagnosis

This is one that gets the press, and perhaps rightly so: AI applications that sift through mountains of medical records and images to detect rare and dangerous diseases sooner than any human could. Researchers at Stanford University, for example, have trained a DL algorithm with 130,000 images of moles, rashes, and lesions to provide a way of diagnosing skin cancer at its earliest stages, as well as screening.

Treatment Design

This application of AI relies on what is effectively a partnership between healthcare professionals and software. Combining patient records with reams of clinical trial data and medical journal entries, the AI app suggests treatments, to which the physicians apply their experience and insight. A study by the University of North Carolina School of Medicine is worth noting here: The school's researchers compared the treatment recommendations from IBM's Watson platform on 1,000 cancer cases with the recommendations of oncologists. Watson gave the same recommendations for 99 percent of the cases.

Hand-Held Consultation

There are now more cell phones on the planet than people, so it was just a matter of time before apps appeared to dispense medical advice through these devices. When AI and speech recognition is applied to this concept, the result is a chatbot-enabled solution that provides "consultation" services based on the symptoms a patient reports, his/her personal medical history, and current medical knowledge. This model also has the potential to cover a large part of clinical outpatient services when applied via an in-house system.

Hot Hardware

None of this would be possible without intensive computing power. The use of GPUs in particular has facilitated AI/ML advances. A GPU has a massively parallel architecture made up of thousands of smaller, more efficient cores that are designed to handle multiple tasks simultaneously. Tensor Processing Units (TPUs), Google's new chip designed for the company's TensorFlow framework, was specifically developed for neural nets. Another type of GPU, the General-Purpose GPU (GPGPU), is showing up in DL research and development projects.

These chips significantly reduce the time required for training an ML solution. What would take a week using a CPU might take a day using a GPU, or a few hours using a TPU.

Among the hurdles CIOs and IT pros are likely to face as they evaluate AI/ML solutions for their healthcare organizations are unrealistic performance expectations.

The 100 Percent Trap

Among the hurdles CIOs and IT pros are likely to face as they evaluate AI/ML solutions for their healthcare organizations are unrealistic performance expectations.

"What's the metric you really need to look for as you evaluate machine learning solutions?" Swartz asks. "A lot of people are stalled by the belief that you want to drive to 100 percent, when what you really want to drive to is above human performance. Getting to 100 percent is impossible, of course, for man or machine, so why demand that level of performance from an AI solution? If your research shows you that your physicians perform at, say, 80 percent, then an AI system that performs at 90 percent should be very valuable to your organization. This is the metric you really need to look for in the use cases."

Conclusion

Artificial Intelligence is the confluence of multiple advances, all of which must be taken into consideration by business decision makers charged with determining the value of an AI solution for healthcare. Some of these technologies are still in the experimental phase; some are already changing healthcare in significant ways, tapping into the ever-expanding stores of digital health data to complement the experience of physicians and other healthcare professionals.

It's true that there is a lot of hype around AI, but its fast and furious evolution is likely to push it past that hype very soon. Recent investments from big players and startups support that prediction, and may be seen as a kind of validation beyond the headlines. There's a lot going on in this space, and CIOs and IT leaders considering investing in it would be smart to approach these disruptive technologies using a strategy that includes close communication among the organization's technologists, subject matter experts, and business decision makers.

John K. Waters has been covering the IT beat from Silicon Valley for more than two decades. He is an editor for Application Development Trends, and contributes to a range of online and print publications. He's also the author of more than a dozen books.