



FROM IDG

AUGUST 1, 2018 | CIO.COM

FEATURE

How health care orgs are addressing AI challenges

Health care companies are embracing artificial intelligence for everything from drug research to diagnostics, but challenges centered around privacy, data and the AI 'black box' remain.

By Maria Korolov

Artificial intelligence is poised to revolutionize health care, in everything from drug research, to diagnostics, to health care delivery — all the way through to billing and accounting. According to Accenture, the AI health market will reach \$6.6 billion by 2021, and will save the industry \$150 billion a year by 2026.

But to get there, organizations need to overcome distrust of the technology, meet integration challenges, and deal with privacy and security issues.

Boston-based Beacon Health Options, which provides behavioral health treatment to more than 40 million people in all 50 states, faced all those challenges when it first turned to AI to help improvement treatment.

"Beacon is the largest company solely focused on behavioral health in the country," says Christina Mainelli, the organization's executive vice president and chief growth officer. "Serious mental illness, substance abuse — we're dealing with some very, very relevant issues in health care today."

Mainelli says that Beacon Health is using AI to identify the patients who need the most help, and get them the interventions they need. Following is a look at how Beacon Health and other organizations are navigating AI challenges unique to the health care industry.

The black box challenge

For health care organizations, the biggest worry with AI is that it can be hard to see why the system makes the decisions that it does. A mistake can literally mean the difference between life and death. This is known as the

"black box" problem of AI.

Intel recently surveyed 200 decision makers from health care organizations, and the black box problem — the lack of trust in AI — was the obstacle that stood out, says Jennifer Esposito, general manager of health and life sciences at Intel. In particular, 30 percent said they were most worried about a fatal error

"I don't think there's a solution right now for how you resolve the black box problem," she says.

Despite that, 37 percent of the respondents said they were already using AI, and more than half said they would adopt AI within the next five years. Most surprising to Esposito, the most popular use of AI was in clinical uses, undertaken by 77 percent of those embracing AI in health care.

Beacon Health faced just this issue of trust when it began rolling out its AI-powered risk assessment system. The organization took a three-pronged approach to address it.

First, a multidisciplinary team of clinicians, technologists, operators, community health care providers and other stakeholders was brought together to come up with the basic guidelines and algorithms for the AI to follow, starting with the factors that were already known.

"We know which variables contribute to those who are at highest risk of admission or psychiatric event," says Beacon Health's Mainelli.

But while a human being can look at seven to ten factors when making a decision, the AI can look at thousands of factors at once — and even more when unstructured data is

pulled in, she says.

Then the team runs experiments. First, they pick a historical time period and run the algorithms on that data to see whether it accurately predicted how often those patients were admitted or had psychiatric events.

"We found that when we use AI, our predictions are 220 percent more accurate," she says.

Once the AI was being used to identify at-risk patients, the team was then able to look at potential interventions, and follow up with the recommended interventions.

"We validate that the patients truly benefited from our intervention, and then we can deploy it as part of our business model," says Mainelli.

Like Beacon Health, the University of Pittsburgh Medical Center addresses the "black box" challenge of AI systems by getting clinicians and administrators involved in the process of creating the algorithms.

"One can't expect that a group of data scientists will develop models in isolation and have the folks in the trenches start using them right away," says Oscar Marroquin, the hospital's chief clinical analytics officer, adding that getting clinicians involved gives them a sense of ownership of the models.

The hospital currently has five AI-driven models in production, and 60 other active projects in the works. The hospital is also going to increase the complexity of the models already in use. Applications include predicting patient readmissions and condition-specific algorithms to predict patient outcomes for asthma, congestive heart failure, and low-back pain.

"The results have been very good," he says. But the hospital was aware from the start that changing behavior isn't easy. "Getting folks involved in the whole process helps them understand the strengths and weaknesses of the data and the complexities associated with consuming it, which helps with clinician buy-in."

Another approach to solve the trust problem is to show why the system made the decision that it did. Take, for example, the use of AI in medical imaging.

At Massachusetts General Hospital and Brigham and Women's Hospital Center for Clinical Data Science, AI has been used to interpret radiology images since 2016.

"We put heat maps over the tops of images to help readers understand where the machine is looking," says Mark Michalski, the center's executive director.

That makes it easy for users to decide how much to trust the system.

For other applications of AI, solving the explicability challenge isn't as easy, he says.

For example, the center recently began using machine learning to answer operational questions, such as how many beds are going to be needed, and questions about population health, such as the likelihood that a patient will be readmitted.

These decisions are based on electronic health record (EHR) data, so there's a limited set of variables to be considered, he says, and the system can show the specific factors that had the most impact on a particular prediction, which increases user comfort with the predictions in most cases.

In fact, one potential problem is too much trust in the system, because it's easy to get too comfortable with it.

"We want people to continuously understand why the systems are working the way they are working," he says. "I'd say the one issue that's the scariest one from my perspective — one of the ones that's going to be the greatest challenge — is ongoing validation and safety testing."

Messy data, legacy systems

Artificial intelligence feeds on data. In health care, that data can be very messy, which can lead to compounded problems when put to use with AI.

"Classical ML and AI systems take as inputs fixed-length feature vectors in order to produce predictions, instead of the structured — relational databases — and unstructured data such as images and notes that reside in medical systems," says Jeremy Weiss, assistant

professor of health informatics at Carnegie Mellon University.

AI technology vendors who work with health care data need to be prepared for these formats.

"We are often faced with unstructured data in behavioral health care, such as case notes and survey data," says Beacon Health's Mainelli, who was impressed with how the vendor the organization picked, Cambridge, Mass.-based Cyft, was able to handle all the different formats.

Beacon Health already works with outside partners, getting the data out to Cyft and back again in usable form was not a problem, as it may be for many health care organizations.

"When the list [of at-risk patients] is returned back to us," Mainelli says, "we can integrate that into our patient management system — it is fed into our care coordination system so it becomes part of the clinician's workflow."

Beacon Health is next planning to bring in third-party data, such as pharmacy records, into its AI analysis initiative.

That could be a bigger challenge, says Carnegie Mellon's Weiss. "Lack of standardization of protocol and policy across institutions can result in yet more data integration issues."

But initiatives under way around the world to standardize medical information can help.

In the U.S., for example, the federal government has issued standards for the electronic health record (EHR).

The Healthy Nevada Project, sponsored by Renown Health and the Desert Research Institute, with AI technology from SAS, is one health care initiative that is making use of EHRs to map the health status of the state's residents.

"We are using analytic modeling and statistical machine learning techniques to examine a variety of issues — team dynamics in the operating room, predicting patient no-shows, and forecasting emergency department load, just to name a few," says Jim Metcalf, the project's chief data scientist. "All of these analyses based on EHR data."

Similar work is being done at hospitals and health care organizations around the country, he says.

"Based on EHR data alone, data scientists are creating statistical models to identify patients most likely to go into septic shock or suffer from undiagnosed high blood pressure or be readmitted after discharge," he says. "With this insight, we can personalize medicine to the individual like never before."

Privacy and security

When it comes to patient privacy and secu-

riety, health care is one of the most regulated industries. AI, with its demand for large volumes of data, is a particular concern.

For Beacon Health, this meant choosing a vendor that was fully HIPAA compliant, so that the data was processed in a secure way.

"That level of certification is table stakes to move into this business," Mainelli says.

In the emerging health care AI market, the issue of health data privacy is being handled in several different ways. Some vendors, such as Cyft, encrypt the data to ensure only those with appropriate permission can view it.

Other vendors are looking at anonymizing or tokenizing the data, so that the AI system never sees any personally identifiable information. A third approach is to perform all the processing locally, so that the data never leaves the health care organization.

For example, VisualDx has recently released an AI-powered diagnostics tool, DermExpert, to identify skin conditions. The app is installed on an iOS mobile device, and the doctor takes a picture of the patient's rash and gets an instant evaluation, without transferring any data to a third-party or the cloud.

The app is part of another trend for AI use in health care: enabling non-experts, such as primary care providers, emergency room and urgent care doctors, to perform preliminary diagnostics.

AI's biggest challenge

The biggest challenge of using AI, says Beacon Health's Mainelli, isn't technological. The big challenge is how to adapt the business to make full use of it.

"How you will use it to reduce the cost of care, and improve patient outcomes?" she says. "It's a new way of doing business."

For other organizations looking to make the same journey, she recommends starting out with a clear idea of the problem that needs to be solved, and how the insights provided by AI will be used.

"If you don't act on the data, you haven't brought a tremendous amount of value," she says.

Then start small, and iterate rapidly.

"Leveraging AI in health care is new," she says. "I think it creates an incredible opportunity, not just to impact those suffering from serious mental illness or substance abuse, but to impact total medical costs."



Maria Korolov

Maria Korolov has been covering emerging technology and emerging markets for the past 20 years.