

AI is 'next revolution' in health care. What are the industry's ethical responsibilities?



John Banja is a professor in Emory University's Department of Rehabilitation Medicine and a medical ethicist at Emory's Center for Ethics.
EMORY HEALTHCARE

By [Melanie Lasoff Levs](#) – Associate Editor, Atlanta Business Chronicle

(Healthcare Section)

April 21, 2023

The term “artificial intelligence,” particularly when referring to health care, may evoke science fiction images of robots performing surgery or computer-generated patient exams. While these scenarios could happen in the future, today’s intersection of AI and health care is more nuanced. As technologies and protocols are being developed, so too are ethicists determining the role of AI in the industry.

[John Banja](#), a professor in Emory University’s Department of Rehabilitation Medicine and a medical ethicist at Emory’s Center for Ethics, weighs in on ethics of AI in health care.

How would you describe AI when discussing where it fits in in a healthcare space? What we largely talk about in health care with AI are these machine-learning models. They typically perform one of two functions. One is image recognition. You use them to detect breast cancer, lung nodules, a Covid pneumonia. You also have natural language processing. The day is not too far away from when a model will be able to extract key words in doctors’ notes and create a bill and send it to the insurance company, keep track of payment, all those computational functions that AI is very good at.

The other is prediction or recommendation models. For example, there's a model out now that looks at the back of the eye of a diabetic patient. People with diabetes are at risk of retinal damage. This model is a camera that takes a picture of the back of the eye and reads the retinal image and makes decision whether the retina is OK or if the patient needs to be referred to ophthalmologist.

Five years ago, Silicon Valley was saying these models are going to overtake health care. That has turned out to be hyperbole. It's much more difficult to create a reliable model than we thought five years ago. The technology is not up to par yet, but it will be five, 10, 20 years from now.

What areas are and will AI be used? I suspect the first big use of AI in our health care system is not going to be in clinical care or treatment. I think it's going to be in administrative functions. We will use AI to schedule patients. They're working on AI models that will predict no-shows and will use that prediction to double-book a particular timeslot. We'll schedule two patients in there and will make sure if they both show up, we will have backup staff available, and will use an AI model to do that.

My own prediction is the first department that AI is going to take over in hospitals is accounts payable. That's a computational function that is just perfect for AI.

Are those images of robots doing surgery just science fiction? When the day comes that AI replaces a doctor, that will be an interesting day. It may not be all that far away, 10, 20 years. We'll see how people feel about that. It's going to be a very rocky and tumultuous ride when that technology gets up to the standard of care and equals a board-certified physician. When that happens, you can bet hospitals will import that technology because you can run it 24-7, you don't pay it vacation benefits or time-off.

How will the health care industry adapt as AI evolves? You're not going to bring this new technology into a big health care center and expect that everyone will adapt to it in 48 hours. There will be loads and loads of glitches. People will have a huge learning curve with these technologies. There will be breakdowns. These models will be wonky and glitchy. There will be mistakes. There will be legal liability issues. Suppose you go for a mammogram and the model says you're clear of cancer and misses a lesion. Maybe now you're going to file a lawsuit. Who's going to be liable for that? The computer software programmer? The hardware manufacturer? The hospital? The doctor? Legal scholars are scratching their heads about that one.

What role does ethics play in using all this technology in health care? One ethical issue with AI is fairness. Is it fair to everyone? Let's say we're going to develop a model that will diagnose skin cancers. We will want the model to be able to tell the difference between a benign skin lesion and melanoma. You will get hundreds of thousands of images and feed them into model and tell the model "this is a malignant skin cancer, this one isn't." What these machine learning models do is break that image down into pixels and create a statistically generated prototype of what a particular malignant skin cancer looks like. When you feel you've educated the model on enough images, you provide it with an image and ask it, "What is this?" That's how you test the model and see how accurate it is and compare it with the accuracy [of a number of] board-certified dermatologists.

It turns out the model is really pretty good in terms of diagnosing [skin cancer] on white persons' skin. It is not at all good at diagnosing on Black persons' skin. One of the reasons is we haven't educated the model enough on dark-skinned people. We don't have a good representation of a Black person in the data set. That is a problem that we are having to deal with. That type of lack of representation is really important.

What are the human responsibilities then? Even if you don't have an ethical molecule in your body, still, if you're developing a model, the accuracy of that model will be very important to you because you're going to want to lease or sell it to someone in the marketplace. And if a competitor has a more accurate model than my model is, all the business is going to go to them. We will be trying to create the most accurate model. The question then is how you achieve that kind of accuracy, especially if you want to be sensitive to problems like I'm saying about fairness. The model should not discriminate against anyone. One of the solutions is you hire computer science firms, AI developer firms, culturally diverse teams whose membership represent minority groups and groups that are sensitive to a history of discrimination and oppression. Make sure the developers are constantly auditing the models out there, and constantly evaluating and assessing their own model. It's like the FDA assessing a drug and it comes out in the marketplace. We still don't know whether that drug, which may have been tested on 25,000 people [in clinical trials] and gotten good results, now it will be used on 15 million people. That will be a far better test of how good that drug is.

How does the medical establishment work with med students on these new technologies?

Medical students know a heck of a lot more about machine learning models than the faculty does. I gave a talk at a conference and identified this as a major challenge for our medical students coming through. They're the next generation of health care professionals. They're the ones spending their careers rocking and rolling with this new technology and machine learning stuff. The professor has to develop this curriculum now to make room for how we are going to use these technologies. The problem is these technologies are still in the research phase. We haven't implemented them yet. In about five years, they will be pretty common. Now is the time we need to start thinking about all these technologies and how we are going to train our future cadre of medical professionals.

What other specialties down the road will benefit from this technology? The first splash will be in the specialties that rely a lot on pictures. That would be radiology. Each of the applications is quite narrow in scope. If you develop a model to look at mammograms, that's all it's going to do. You've got to get a model for each one of these diseases. These models are showing a lot of promise with image recognition technology. Another one that may appear even sooner than image recognition is natural language processing models. Doctors and nurses today spend up to 30% of their time logging data into medical records, just doing documentation. The thinking is that these models are going to be able to relieve health professionals of a heck of a lot of that stuff.

You've done a lot of work in the "medical errors and their disclosure" space. How can AI play a role in that issue? If the models become as sophisticated as we think they will be, they should help out enormously. They won't make anywhere near as many mistakes as health professionals do today. Moreover, in something like breast cancer, these models might get so good they'd be able to detect breast cancers before a radiologist will actually see them on a mammogram. If we have models that detect a breast cancer when it's just a couple of molecules old, that means earlier detection which means early treatment and better outcomes. We hope that errors will plummet, therefore harms that result from errors will also diminish, therefore medical malpractice insurance premiums should plummet because there will be fewer cases. That raises issues, though, of who will be liable when the model will make a mistake and that will be a question, especially in the early days of the technology. This is the next revolution in health care. Let's hope we can control it and manage it.