Point/Counterpoint:

Artificial Intelligence in Healthcare

Antonia F. Chen, MD, MBA
Moderator

Adam C. Zoga, MD
Point

Alexander R. Vaccaro, MD, PhD, MBA
Counterpoint

Projected to quickly become a more than \$6 billion market, artificial intelligence (AI) in healthcare is top of mind for many healthcare institutions. How can it help improve outcomes and support institutional budgets and goals? To delve into this and more in the AI field, *Healthcare Transformation's* Associate Editor, **Dr. Antonia Chen**, leads a discussion between **Dr. Adam Zoga**, a physician and diagnostic radiology specialist at JEFFERSON UNIVERSITY HOSPITAL and **Dr. Alexander Vaccaro**, leader of the ROTHMAN INSTITUTE and professor of neurosurgery.

Dr. Antonia F. Chen: How will artificial intelligence (AI) transform healthcare?

Dr. Adam C. Zoga: I am going to reference **Andrew Ng**, who is on the forefront of AI in healthcare and the former chief scientist at BAIDU and an adjunct professor at STANFORD UNIVERSITY. **Andrew** says that in 15 years from now, AI is going to be as necessary as electricity for members of the healthcare team. I agree that we are going to be reliant on AI. It is going to be so built into our day-to-day workflow and treatment algorithms to the degree that we are not going to know what we did before AI was there to help us.

In the short run, I really think AI is getting close to being ready to take on some of the tasks that we as members of the healthcare team do not particularly enjoy doing and find irritating, such as charting. However, I think AI is a long way from playing a huge role in decision making during hospital rounds, but I think AI is ready to start helping us with things such as workflow management in the next few years. Diagnostic imaging is certainly a place in healthcare where AI could help manage some of the daily tasks that might seem mundane or routine, but it can also help radiologists avoid errors by alerting us to something we may have overlooked.

© Antonia F. Chen et al. 2017; Published by Mary Ann Liebert, Inc. This article is available under the Creative Commons License CC-BY-NC (<http://creativecommons.org/licenses/by-nc/4.0>). This license permits non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited. Permission only needs to be obtained for commercial use and can be done via RightsLink.

Dr. Alexander R. Vaccaro: In my mind, AI is divided into two separate systems. One is computer-aided detection and the other is computer-aided diagnosis. When AI was first introduced, surgeons wanted radiologists to use this type of technology to look at images and figure out if an abnormality existed, thus implementing detection. This would be confirmed by a human, such as a radiologist, and they would agree or disagree with the assessment, which would provide a high sensitivity so no abnormality was missed.

Now, the quantum next step is computer-aided diagnosis, where AI would actually tell you the diagnosis and then you would have to decide what to do with the information. This is comprised of five steps. One is pre-processing, which is a binary step that decides if a finding is present/not present. Second, once the computer learns what you are looking for, then it segments it into areas of anatomic geography. The third step is candidate detection, where it is decided that a lesion exists based on certain features that you want to extract. In the fourth step, you classify the lesion, and in the fifth step, the system provides the diagnosis.

Spine surgeons are in the process of developing an AI system to classify spinal fractures. The first thing the software does is identify the fracture by looking at the vertebral boundaries and the cortical end plates to see if there is an area of irregularity. Once it detects a fracture, the software can then classify the type of fracture, such as type A or type B injury. The surgeon can then agree or disagree with this classification. It can help in detection and diagnosis. The goal is to use this technology to teach others in the world, which can lead to improved treatment. With AI, certain regions of the world that do not have the resources that we have can bypass the need for outsourcing and can capitalize on using this technology to improve patient care.

Dr. Chen: How can AI be used to improve healthcare in areas outside of our current workflow?

Dr. Zoga: We have to remember AI is going to play a much greater role in the healthcare system than simply radiology. It is going to make contributions throughout the healthcare management team. I really think there is

Artificial intelligence expert Andrew Ng also co-founded Coursera, a web resource that offers free online courses from top universities. Coursera is currently the largest MOOC (Massive Open Online Courses) platform available, partnering with more than 80 universities and offering about 400 courses. According to his website, www.andrewng.org, **Ng's** goal is to connect everyone to **a great education for free.**

“Spine surgeons are in the process of developing an AI system to classify spinal fractures.”

huge potential for AI to aid workflow management, particularly in my specialty of diagnostic radiology.

For example, we learned that last year more than 80% of extremity X-rays ordered through Medicare were not interpreted by radiologists, which is problematic if there are unexpected findings beyond the specialists' organ system expertise. AI might alert a urologist looking for a ureteral stone to the possibility of a lung nodule. AI can also play a huge role at triaging an imaging worklist before a radiologist ever sees it. As an example, shoulder X-rays taken in the radiology department or in a shoulder specialty clinic can first be evaluated through AI to look for urgent findings such as discontinuity of the osseous cortex, articular dislocation, or pneumothorax. These images should not wait until the end of the day to be read, and maybe these studies can be sorted to the top of the worklist with the help of AI.

Another goal of AI should be picking up incidental lesions, which include important findings beyond what you were looking for on an imaging study. AI needs to be able to pick up a lung nodule, and then help classify the lung nodule using the Fleischner criteria and recommend the appropriate follow-up. The new AI models are based on deep learning, and deep learning can do more than just make an observation. It can declare a confidence in the observation that it has made, and then it can recommend the next step. Then, the findings observed by AI can be reviewed by qualified members of the healthcare team, who will ultimately decide on the best course of action.

I also think that AI will be useful outside of imaging, and will be helpful in medical record management. AI can use its deep learning to guide the ordering of appropriate laboratory tests on a daily basis systematically, and it may help with the development of some treatment algorithms.

Dr. Vaccaro: Adam brings up an important issue for many surgeons and that is detecting lesions that are outside of our bailiwick. As orthopedic surgeons, we are very good at reading musculoskeletal radiographs, so we tend not to rely on radiologists always. However, what happens if there is a lung or kidney lesion? I remember in my first or second year of practice in the early 1990s, a patient had

“AI can also play a huge role at triaging an imaging worklist before a radiologist ever sees it.”

an unrelated spinal lesion, and it was very subtle—I missed it. I remember the patient having discomfort, and I could have relieved that discomfort a lot earlier if I detected it and initiated treatment.

In the field of medical oncology, you absolutely do not want to make a mistake in diagnosis, especially from imaging studies such as computed tomography (CT) scans looking for pulmonary lesions or mammography looking for breast lesions. It would be beneficial to have systems in place where AI, through computer-aided detection, would detect a lesion and then is able to validate it before burdening a radiologist in order to minimize the occurrence of false-positive findings. This would decrease wasteful follow-up secondary imaging studies or procedural interventions. You could save so many lives with this because AI may assist with early detection of malignancies where we can still modify the natural history of the disease.

Another current not so rare problem is when a physician orders an imaging study and a report is generated, the ordering physician never sees the report. This type of technology could end that by detecting a suspicious lesion, then notifying the radiologist to recommend further imaging with the consent of the patient, with a report generated to the referring physician. With this process, you know that the abnormal lesion was followed up on appropriately. Hopefully, it will be automated in such a way that once the initial imaging is performed, the lesion is immediately picked up and then further imaging is done while the patient is still in the imaging suite.

Dr. Chen: How do we know how AI arrives at assessments and plans?

Dr. Zoga: No one really knows how AI arrives at assessments and plans, and we need to understand the process much better before we use it in situations where lives are at risk. What is scary is that everything we do as healthcare professionals is based on experience. It is based on publications, and we have learned throughout history. Now, we are relying on a new system, and we are putting faith in the idea that its assessments are accurate and its confidence is valid.

Thankfully, there are new tools that can trace the decision tree back through the AI neural network, referencing publications and experience with alternative

AI by the Numbers

AI within the healthcare market is projected to reach

\$6.6 billion by 2021—
a compound annual growth
rate of **40%.**

Source: Frost & Sullivan

treatment outcomes. However, we need to understand why and how AI systems report their confidence before widespread implementation.

Dr. Vaccaro: I am going to use a very simple analogy, which we are experiencing now in surgery. When I began training, I trained as an open surgeon. As time went by, imaging got better, and we were introduced to three-dimensional technology that could be used in the operating room. This included intraoperative CT and similar technology, which allowed us to visualize structures hidden from direct view. Now we can even make a hologram of the desired anatomy as we operate. Over time, better nerve and vessel detection systems were also developed, allowing us to use guide instruments more safely and effectively. Our eyes are depending on more than direct sight to operate safely on a patient. By starting as an open surgeon, it has helped me become a better minimally invasive surgeon.

Now I am teaching the fellows, residents, and medical students how to operate minimally invasively. However, they may not have the same benefit I did in learning from open surgery. They are learning a way different from how I learned. AI will do the same thing; it will help us leapfrog and develop a new paradigm in image detection and diagnosis. The negative aspect is that the steps that we may need to get there may be frustrating and timely. Now, with robotic and virtual-assisted rounds, a machine can be provided with a list of symptoms in which the robot then provides a diagnosis. We as physicians did not have to go through the deductive or inductive process to understand exactly how that diagnosis was made.

The key is to not make us too dependent on AI. It is imperative that we continue to learn medicine the traditional way, so that we understand the process by which AI arrives at diagnosis and assessment. That way, we can use AI to help us make more accurate diagnoses with less potential for false positives and negatives. If AI does make a mistake, we as humans should have a means of detecting this and correcting any potential error before a patient is harmed.

Dr. Chen: What parameters will AI optimize for? Are we looking for the best outcome in terms of scores or for decreasing mortality, and how do we control that?

Dr. Zoga: It takes a long time to get outcomes data, but if AI can help us generate and use some of the quality metrics that we deal with on a daily basis, then it is a win from the outset. If AI can help triage studies that might have urgent findings so that they get interpreted more expediently, then that is also a win.

We have a study that JEFFERSON radiology resident **Ali Syed** is presenting at the RADIOLOGIC SOCIETY OF NORTH AMERICA meeting in November, where we let AI look at chest X-rays. **Dr. Syed** built in a deep-learning model that would

“Now, with robotic and virtual-assisted rounds, a machine can be provided with a list of symptoms in which the robot then provides a diagnosis.”

identify shoulder dislocations. Hopefully, no physician will miss an anterior shoulder dislocation on a shoulder X-ray, but the story is a little different for chest X-rays. When a trauma patient comes in and gets 5–10 different imaging studies where there are innumerable images, such as CT scans of the entire body, it is easy to miss a peripheral diagnosis such as an anterior shoulder dislocation on a chest X-ray. His model worked, and AI picked up a number of anterior shoulder dislocations on chest X-rays without a radiologist or help from clinical history. This is where AI can be useful—doing some of the dirty work that we do not really have the manpower and the time to do ourselves.

Dr. Vaccaro: The simplest things to measure are process measures, such as blood loss, length of stay, and so on that occur related to a hospital stay, and how to prevent a patient from being discharged prematurely, which may increase the risk for re-admission. AI can analyze all data points, synthesize and analyze the data and develop prediction models to alert the care team to potential problems. From a diagnostic perspective, this may allow an earlier detection of disease and the commencement of more time-efficient treatment. Another service AI can provide is acting as a watchdog to detect red flags in a patient's medical record that were missed or may be missed due to the subtlety in presentation. Remember that as a physician, you are responsible in general for what is in a medical record, and with the vast amount of information present in an electronic medical record, this may be a daunting responsibility, depending on the number of consultants and information downloaded. AI can extract pertinent information and then notify the caregivers of its importance and possibly initiate earlier intervention. We need a helping hand, and that is where AI comes in. It will be the helper not only to expedite but also to make care more efficient, make care less fragmented, make care more cost-effective in the long run.

Dr. Chen: With regards to privacy and AI, how do we make sure that what AI finds is protected and who owns those data?

Dr. Vaccaro: As far as I know, there are four major players in the AI world: IBM, AMAZON, GOOGLE, and MICROSOFT. I would like to use the example of IBM with the Watson computer. I expect that IBM (Watson computer) will house and own the data. It will then, through the AI process of deep learning, learn from these data and create more effective diagnostic and treatment algorithms.

In terms of personal privacy, algorithms can be written to be HIPAA compliant and therefore protect patient privacy. However, AI can learn from a patient's medical information and combine that with others to treat disease processes best.

“From a diagnostic perspective, this may allow an earlier detection of disease and the commencement of more time-efficient treatment. Another service AI can provide is acting as a watchdog to detect red flags in a patient's medical record that were missed or may be missed due to the subtlety in presentation.”

Thus, IBM will own knowledge created by its Watson computer, but laws can be written to protect individual rights. The keeper of the information will benefit and grow from their services.

Dr. Zoga: I am not really concerned about HIPAA compliance and AI. I think we are already implementing the measures that need to be in place for patient privacy. Information that is acquired in the healthcare setting is distributed globally. When a patient gets on a magnetic resonance imaging scanner or a “proton spinner,” the acquired data get stored in a computer and go to data storehouses, while they are simultaneously being analyzed to create a diagnostic image. Protected health information (PHI) is already in the hands of IT systems throughout healthcare. These data are well protected and should remain protected when AI is being utilized.

Dr. Chen: Then, who controls the role of AI in medicine?

Dr. Zoga: I believe it is going to be a commercial product that we as a healthcare system can choose to purchase and use. We purchase innumerable commercial products that use and store PHI from many vendors throughout healthcare. I don’t see AI as vastly different from the electronic medical record or a CT scanner.

Dr. Vaccaro: Let’s use IBM Watson as an example. Chat bots are little miniaturized AI robots that you can ask questions about the symptoms that you are having, and then they will have a conversation with you explaining the differential diagnosis. These bots will help you diagnose your illness, and then it will tell you who to see to seek treatment.

Companies will commercialize these bots, and hospitals will then license this ability from AI companies. These AI computers will be in the hallways of our hospitals when we round; it will walk alongside us, and will be another member of the team to help with our diagnoses and treatment. Since this is a commercial entity, we will have to pay for that robotic house officer probably more than a house officer gets paid today.

Will the government get involved? Maybe, but the government usually is last to the table. Companies will develop the technology, and then maybe license it to the government so that it is more affordable.

Dr. Chen: To quote Dr. Klasko: “Any doctor who can be replaced by a computer should be replaced by a computer.” Who is going to be replaced in medicine then?

“These AI computers will be in the hallways of our hospitals when we round; it will walk alongside us, and will be another member of the team.”

Dr. Vaccaro: Hopefully, no one will be replaced, and I believe that we will actually add more people. It's not necessarily a bad thing to have a large portion of your GDP go to healthcare because it employs a large number of people. We will employ more people in the business of medicine, more people in IT, more people in computer science and healthcare. We may actually see more people go into the business of medicine, so I do not think it will necessarily replace anybody. It may actually grow the business and make the world healthier.

Dr. Zoga: It seems that everyone references radiology when they talk about AI. Speaking for my specialty, radiologists really are not afraid of losing their jobs to robots, at least not radiologists who perform and interpret tertiary imaging studies. For years, radiologists have known that we need to shift to a value-based system and away from a volume-based system. We need to be involved in the patient care team.

We made some mistakes in the last decade where we commoditized ourselves, and we are moving away from that. Of the current JEFFERSON radiology residents, five of them are involved in AI projects, and they are excited about the possibilities of what AI can do for radiology going forward. We are already using AI, and I do not think it is taking good jobs away or really any jobs away from the healthcare system. I think it is shifting jobs a little bit, and there may be some different types of jobs available, especially in radiology. It will help morph traditional film libraries into data centers with an IT bent, which will allow us to offer additional value to the healthcare system.

I think that AI might replace the pen and torn-up coffee-stained piece of paper from the tangible physical patient record. IT innovations have already replaced film that used to come out of X-ray units. In our ultrasound division, we use an imaging enhancement tool called Imorgon. When our technologist places his or her calipers on a lesion that he or she sees on an ultrasound image, Imorgan automatically identifies density and volume of the lesion and then auto inputs this information into patient reports. Thus, without anyone losing their jobs, we have created a more comprehensive and accurate patient report more quickly. Based on this, I do not know that anyone needs to lose their job over AI, but I think there is an opportunity for some of us to shift our jobs and understand a little bit more about how these computers work. Radiology provides the perfect substrate for AI to grow, and I am optimistic about our future, working with AI. ■■■

“Radiology provides the perfect substrate for AI to grow, and I am optimistic about our future, working with AI.”