

Speaker 1:

Join us for our AI + Health Online Conference on December 8 and 9, 2021. At this wide ranging online conference, the Stanford Institute for Human-Centered Artificial Intelligence, Center for Continuing Medical Education, and Artificial Intelligence in Medicine and Imaging, are convening experts in leaders from academia, industry, government, and clinical practice to explore critical and emerging issues relating to AI's impact across the spectrum of healthcare. Content will be relevant to practitioners, researchers, executives, policy makers, and professionals with and without technical expertise. This is a paid event that includes student and Stanford affiliated discounts. By participating in this activity, physicians can earn a maximum of 10 AMA PRA Category 1 credits. To learn more about the registration fees, please visit aihealth2021.stanford.edu.

Ruth Adewuya, MD (host):

Hello. You are listening to Stanford Medcast, Stanford CME's podcast where we bring you insights from the world's leading physicians and scientists. If you're new here, consider subscribing to listen to more free episodes coming your way. I am your host, Dr. Ruth Adewuya. This episode is part of our hot topics mini series. Today we are talking about artificial intelligence in medicine. I am joined by Dr. Curtis Langlotz. Dr. Langlotz is professor of radiology and biomedical informatics and the director of the Center for Artificial Intelligence in Medicine and Imaging at Stanford University. Dr. Langlotz's laboratory investigates the use of deep neural networks and other machine learning technologies to help physicians detect disease and eliminate diagnostic errors. He has led many national and international efforts to improve medical imaging, including the Medical Imaging and Data Resource Center, which is a U.S. national COVID-19 imaging research repository. Thanks for chatting with me today.

Curtis Langlotz, MD, PhD (guest speaker):

Thank you, Ruth. It's a pleasure to be here with you today.

Ruth Adewuya, MD (host):

Given your clinical background, how did your interest in artificial intelligence and medicine begin?

Curtis Langlotz, MD, PhD (guest speaker):

I've always been interested in computers. When I was very young, didn't read very much but read a lot of science fiction. I read a book called Gödel, Escher, Bach, which was a Pulitzer Prize winning book that came out when I was an undergraduate here at Stanford and dealt with concepts like self reference and symmetry and intelligence and really got me interested in AI. Then I took a couple of courses in computational linguistics, thinking about we process text and think about the words that we use from a professor who's now retired named Terry Wingrad. He did some really early work on interpreting language as instructions to a robot to put the red pyramid on top of the blue cube and things like that. That got me very interested. Ultimately I enrolled in the Master's program in AI here at Stanford and was one of the few biomedical people at the time. That really got me interested in medical AI.

Ruth Adewuya, MD (host):

AI in medicine can sometimes feel so nebulous for clinicians who are not trained in this field. What should physicians know about the current state of research in AI in medicine?

Curtis Langlotz, MD, PhD (guest speaker):

I would say it's an incredibly powerful tool, but it's just a tool and it depends on how we use the tool. Just to give you an idea, when I was doing my PhD as a graduate student, it might take you your or whole PhD to build a system that could work on one patient. Now, with the right training data and these new machine learning algorithms, in a matter of a week or two you can produce a system that's even more powerful and more accurate than what we could do back then. These new technologies and their capabilities are a real step wise advance into what's happening.

Curtis Langlotz, MD, PhD (guest speaker):

But, on the other hand, some of the headlines I think are sensational. They're a little bit overblown. When you see, even about our work, that this algorithm can find pneumonia on a chest x-ray better than a radiologist, that's maybe true but, in fact, usually the algorithm plus the human is better than either one alone. As radiologists, we train to recognize 200 or 500 different things on a chest x-ray, not just pneumonia. These technologies are going to transform the way we practice medicine. It will be everywhere, but I think it's going to be a slower evolution. The first changes we'll see will probably be some of the things that'll help us do the things that we find less enjoyable to do as part of healthcare or alert us to things that we might not otherwise recognize.

Ruth Adewuya, MD (host):

I'm glad you brought up radiology. The number of FDA cleared AI or ML enabled medical devices is far higher for radiology than any other specialty. Why do you think radiology is such a focus for applications of AI?

Curtis Langlotz, MD, PhD (guest speaker):

These technologies are fantastic at computer vision. Our images are digital, so it's really just an array of numbers, if you think about it. An image is the various brightnesses or colors, and those could all be represented as numbers so it's an ideal kind of input to these neural network technologies. These technologies are very good at pattern recognition. They're very good at measurement and quantification of things. They're very good at finding a needle in a haystack. Those are all things that, particularly the needle in the haystack and quantification, are not particularly strengths of humans. There's a good complementary there in the computer vision space. Then radiology, specifically, I think it's because our data has been digital now for many, many years in standard format so it's relatively easy to generate large amounts of training data. But I think of the over 300 FDA cleared algorithms, more than 200 of those are in radiology. The next most common specialty is cardiology, which is also clearly a digital imaging kind of specialty. That's where these technologies have been used the most.

Curtis Langlotz, MD, PhD (guest speaker):

A lot of the innovations which occurred about five to eight years ago happened because of a database called ImageNet, which is about 14 million just natural scenes, photographs from the internet, which were used to develop competitions for computer vision. The revolution occurred back in about 2013 or so when those were clearly making massive improvements with these new neural network technologies. Everyone, I think, at that point saw that these were going to have a profound effect on medical imaging as well.

Ruth Adewuya, MD (host):

One leading AI researcher notably said that we should stop training radiologists now. You have instead predicted radiologists who use AI will replace radiologists who don't. Can you elaborate on this?

Curtis Langlotz, MD, PhD (guest speaker):

Yes. I think you could replace the word radiologists with any clinician, really, in that statement because I think these technologies will ultimately be everywhere. It will take some time. I talked earlier about this notion of not just one thing on a chest x-ray, but 200 things. I think humans will always be needed, and this is not just in radiology at all across specialties, needed to integrate findings to follow disease over time, although increasingly AI algorithms can do that, integrate the clinical information, do some strategic thinking about the patient. One way to think about it is using an aviation analogy. I think we all are happy that there's an autopilot to help the pilot when we fly in a passenger airplane. The other hand, we wouldn't want the computer to be completely flying the plane because there are strengths that each has.

Curtis Langlotz, MD, PhD (guest speaker):

I think we'll all be getting better versions of autopilot so that these automation tools will be doing more and more of the things that we're not very good at. In flying, it has to do with fine corrections to stay on course and some of the boring parts that humans aren't very good at. Yet you still need humans there to respond, to look out the window and respond to urgent events. I think that's going to be true throughout medicine as well so that we all can work at the top of our license. Again, taking from my world radiology, let's say repeated measurements of lesions and comparing them over time volumetrically. That's something that's not particularly interesting inherently to do all of those measurements for a human, but a computer can do that very rapidly and can plot it out for us. Then we can analyze it and make conclusions from it. I think that mix is where we're headed.

Ruth Adewuya, MD (host):

Can AI models assist in training new radiologists, particularly in low resource settings or when physician time is otherwise limited?

Curtis Langlotz, MD, PhD (guest speaker):

The notion of how AI will affect education is really interesting. I don't think we really know the answers there yet. The jury is still out. First of all, there's a lot of good research going on about using AI to help understand where our learner is in their learning journey and then to present them with the right material at the right time. I think that's really still fairly difficult to do but, again, there's some early work in that area. We talk a lot about how AI is going to affect education of our clinical trainees: residents and fellows and so forth. Just because we have a calculator doesn't mean we shouldn't still learn how to add and subtract. There's this sense of maybe the trainees shouldn't have all the power tools on day one of their training so that they learn some of the basics inherently before they then use these more powerful tools.

Curtis Langlotz, MD, PhD (guest speaker):

Then there's the notion of educating physicians about AI. There, I would use an analogy to what happened when MRI first came out in radiology, which is a new realm of physics that we weren't really being taught about. We had to learn a lot in our training about radiation physics, but not so much about MRI physics and protons and spins and phase encoding and all of that. But we ultimately did have to learn that because some of the artifacts that are produced when you have MR images are very important to distinguish an artifact from a real lesion in the liver, let's say. You need to understand a little about the physics to know about how the devices can fail you. I think that can be mapped directly

onto our view of AI, which is that we're going to need to learn more about how it works as physicians so that we can recognize when it might be leading us astray.

Curtis Langlotz, MD, PhD (guest speaker):

Then your last point about underserved areas. I can tell you that, for my specialty radiology, about two thirds of the world is underserved by radiologists. Even if AI can be used to help train radiologists, I think the main role, at least in the near term, will be to augment the skills of other professionals, non-radiologists, who will need help in interpreting some of these images and so that we can provide better quality diagnostic interpretations in areas that might not otherwise have it.

Curtis Langlotz, MD, PhD (guest speaker):

Then the other area where these technologies are having an impact is that there are now much less expensive devices. You may have heard of this Butterfly ultrasound device you connect to your smartphone. There's also now a portable MRI machine that you can wheel around the hospital and is good for certain applications. Those are devices where the image quality would not have been high enough, really, to be useful. But with some of these AI image enhance techniques, you can manage these with these smaller devices and still provide diagnostic information. AI and clinical education are going to be interacting like that, I think, in various ways as we move forward.

Ruth Adewuya, MD (host):

That's truly exciting. To AI researchers, applications of AI in medicine may seem limitless. Is basic science research in AI in medicine important? Or is it an advantage of this field that those new technologies can be created while much of the foundation is still not well understood?

Curtis Langlotz, MD, PhD (guest speaker):

Absolutely. This kind of basic science in AI and in medicine is important and necessary. I would go back to the example of ImageNet. What we had from those competitions that were trying to recognize a spider or a rabbit or some object in a natural scene, we immediately saw those could work in medicine. But in medicine, our images are much higher resolution than those photographs. They're often three dimensional, four dimensional. They may be multi-modality, different types of images, CT, MR, x-ray, multi-channel, we need to compare to the prior. There are many ways in which medical imaging is very different than that. We need to do a lot of basic research to make those problems tractable with similar kinds of techniques.

Curtis Langlotz, MD, PhD (guest speaker):

Many of the images in ImageNet were labeled using what's called a mechanical Turk. They just used people who were paid small amounts on the internet to go and look at those images and apply labels. For medical images, it's much more costly to label because we need highly trained specialists to provide labels for training data. There's a lot of good research going on using what's called self supervised learning. For example, you can create an artificial task, like you can remove part of an image and then ask the algorithm to hypothesize what that missing piece might look like, like a missing piece of a jigsaw puzzle. Just in doing that, that algorithm learns a lot about what is present in medical images, generally. That's a way to dramatically reduce the need for human level training. Then the other area is explanation. These algorithms will be working in concert with humans and having the human understand a little bit about why the system came to its conclusion is going to be very important. There's a lot of research going on in that area as well.

Ruth Adewuya, MD (host):

Yeah, because I think one common critique is around the explanation piece of AI models.

Curtis Langlotz, MD, PhD (guest speaker):

Yeah.

Ruth Adewuya, MD (host):

I imagine that that poses a challenge for AI in medicine.

Curtis Langlotz, MD, PhD (guest speaker):

It does. There are good examples of automation bias, essentially where humans are more likely to believe the machine whether or not it may be correct. But I would reframe the explanation question a little bit and think about it more as trustworthiness. Let's think about Tylenol. Do we know how Tylenol works? Well, we may not know the mechanism, but we know from clinical trials that it works. There are many medications like that. Some of these algorithms, it will be the same. We fundamentally do know what's going on inside the model. It's fundamentally mathematics and an optimization problem and that's well understood. But for a given case, we don't always have a clear explanation that a human would understand as to why the algorithm came to its conclusion. Yes, explanation is important as much as we can obtain it. But even if we can't, if we have good clinical trials to say, "This is going to help in a real clinical setting," that creates this sense of trustworthiness that I think is important.

Curtis Langlotz, MD, PhD (guest speaker):

The other thing I would say about these algorithms is that the way that they're not like Tylenol. Just imagine if, every time you got a shipment of Tylenol, you needed to check whether it was still working. In a sense these algorithms, just because they've gone through FDA approval and they've been proven on a particular data set, doesn't necessarily they'll work in my clinical setting. If it's been trained at Stanford, it might not work at Emory University and it might not work at Harvard. Or it's trained there, it might not work here. There's more due diligence that's required on the front end before we implement these.

Curtis Langlotz, MD, PhD (guest speaker):

Then the other corollary to that has to do with data drift. In my world, there are new imaging devices that come online. There are new patient populations that we may see due to a clinic that opens or closes, new diseases like COVID. All kinds of reasons why, over time, these algorithms may not work as well as they do on day one. Due diligence on the front end and also monitoring or some kind of analytics over time to make sure they continue to work on day 100 as they are on day one.

Ruth Adewuya, MD (host):

How do you think about potential biases in AI models? What is the significance of diverse data sets to reduce bias?

Curtis Langlotz, MD, PhD (guest speaker):

Yeah. Great question. First of all, the word bias has multiple meanings. One is a definition that a statistician or a computer scientist or a data science would have. I have this algorithm and how can I reduce the amount of bias that it's showing in its results? But, in a way, that's too late if you're thinking

of about bias at that point where you've already chosen your data set and the question. Bias really can enter these machine learning algorithms at any stage, including at the very beginning. What question are you asking? What data set are you using? Is it representative? Is it going to give a good answer for everyone? Why is it giving different answers for different populations?

Curtis Langlotz, MD, PhD (guest speaker):

I'll give you just one example. Our Center for AI in Medicine and Imaging, we have over 120 faculty now that are doing this kind of work. But our first algorithm that came out of the center, this was some time ago now, the work of David Larson and others was bone age. You're taking an x-ray of the child's hand and you're looking at the cartilage versus ossification and determining the physiologic age of the child and comparing that to the chronologic age to look for developmental delay. That algorithm was trained against a reference standard called Greulich and Pyle. That data was based on about 300 white children who grew up in Cleveland in the 1950s. Is that getting the right answer for everyone? Well, we know that bone development is different across different ethnic groups, so probably not. It's interesting because the way that you think about that question ... Should we say, "Oh, there are these differences. Let's develop different models for different ethnic groups. We'll get different data sets, different reference standards." But there's really some context there as to why those differences exist.

Curtis Langlotz, MD, PhD (guest speaker):

If we're building different models, are we just accepting that there may be some societal factors that are leading some people's bones to develop less well than others? Maybe we should fix those societal factors rather than just accepting them and building that model. But on the other hand, if it's a genetic difference that we know statistically is present and it just makes sense to have different models for different groups, then we ought to go ahead and do that. It takes some very careful thought about the context in which these algorithms are developed, and very early in the process, and think carefully about what questions are we asking and what data sets do we need so that we get appropriate answers for everyone.

Ruth Adewuya, MD (host):

How can healthcare systems prepare to take advantage of new AI tools as they are developed?

Curtis Langlotz, MD, PhD (guest speaker):

This is going to be a new revolution, a new area of investment for health systems, without question; investing in the people and the infrastructure that makes it easier to deploy these models. If you think about the electronic medical record, it's table stakes today. If you want to have your data in digital form, and it has so many advantages, you need to have an electronic record. In order to make it easy to deploy these models, to make it plug and play, whether it's a vendor model or a model that may come out of a laboratory, you need to have certain infrastructure where you can place the model. Often it's in the Cloud. You place the model. The clinical data the model needs can be sent through interfaces to that system. The model can be executed and then the results transmitted back into the clinical workflow, wherever it may be.

Curtis Langlotz, MD, PhD (guest speaker):

That's actually a really interesting point of study today and, again, we'll be talking about this in the course in December, is that implementing these is a lot more than just thinking about the model itself. It's all about the environment around. It's really a change, often in the way that we work. Who should be

receiving the result? What should they do with it? How many false positives are there? Is that going to be a distraction? Is it going to take actually more time versus less? Really thinking about these implementation challenges as a performance improvement type of a project where you're looking at implementing the model, tracking over time the effects, and thinking about it in a system wide way. Then the other factor I think that we'll all just need to consider is this need for due diligence and think about the generalizability and assure that it's going to be working for our population and over time.

Ruth Adewuya, MD (host):

When you look back at when you first started your Master's program and this was new, I imagine there was skepticism around artificial intelligence in medicine. Fast forward to today, I imagine there's still a bit of it. But do you generally find that clinicians are more receptive to the idea and the role of artificial intelligence in medicine? Or do you feel like you're constantly having to tell the story how its important?

Curtis Langlotz, MD, PhD (guest speaker):

Yeah. We've gone through many called AI winters, where there was a lot of expectation and then the expectations weren't met. It really went through a difficult time. I think this time is different in the sense that, as I was saying earlier, the technological advance in terms of the accuracy and the speed with which we can develop these systems is really immense. There's a real step up in what we can accomplish with these technologies. That part is real.

Curtis Langlotz, MD, PhD (guest speaker):

There still are, even given that reality, probably some expectations that are maybe a little higher than they need to be. There probably will be some consolidation in the market, but this is going to continue to have a profound effect on the way that we work, I think, increasing over time. But ultimately will be a benefit to the physicians who use these technologies. Even though, particularly in my specialty, there was a lot of anxiety early on about, "Well, am I going to be replaced?" I think that those anxieties now are, as we learn more about the technologies and how they work, those have really abated. We're just starting to think about what's the best way for us to use these to benefit us and the practices that we have and the patients that we serve. I'm really optimistic overall about these technologies and just really excited and feel privileged to be able to work in this area today.

Ruth Adewuya, MD (host):

That truly sounds exciting. What do you see as the next frontier for research in AI applications in medicine?

Curtis Langlotz, MD, PhD (guest speaker):

One of the biggest is multimodal data. Today there are a lot of applications that focus on data from the electronic health record, or it may be data that is imaging data, or it may be data that's genomic data. But when we start to have the ability to combine those different data sets and look at a patient from multiple dimensions and think about their diagnoses and be able to make predictions, that's incredibly powerful. The other we already have talked about, and that has to do with translating these models and the implementation science that's required to make decisions about which models are ready to implement and what are the effects going to be and what are the technologies and the infrastructure that we need to make them accessible to our patient care activities. I think those are the two biggest areas that I see that would be the next frontiers

This transcript was exported on Nov 08, 2021 - view latest version [here](#).

Ruth Adewuya, MD (host):

Really exciting stuff. What is one key takeaway you would have for clinicians on AI in medicine?

Curtis Langlotz, MD, PhD (guest speaker):

This is likely to have an impact on your practice and you ought to take a little bit of time to learn about it. It's truly exciting. As you learn about it, you can think about what are the ways in which these technologies can affect my practice and might benefit my practice. Then think about, particularly at a place like Stanford, we'd love to hear about that. Maybe we can work on some of those problems together because that kind of interdisciplinary interplay is really what helps us make progress, where you have clinicians who are the experts at care who can identify the problems, then you have the data science, computer scientists, AI experts. Those teams of people working together on real problems, not only develop new methods that advance the science, but also develop new technologies that can really help our patients. I would encourage everyone to learn a little bit about it and think about how it might help their practice.

Ruth Adewuya, MD (host):

Excellent. Thank you so much for that. Thank you for chatting with me today.

Curtis Langlotz, MD, PhD (guest speaker):

Thank you. It's been a pleasure.

Ruth Adewuya, MD (host):

Thanks for tuning in. This podcast was brought to you by Stanford CME. To claim CME for listening to this episode, click on the Claim CME button below, or visit medcast.stanford.edu. Check back for new episodes by subscribing to Stanford Medcast wherever you listen to podcasts.