

OpenTheo

AI and the Soul of Medicine

October 28, 2021



The Veritas Forum

Dr. Eric Topol (Scripps Research), Dr. Rosalind Picard (MIT Media Lab), and Dr. Jon Tilburt (Mayo Clinic) in conversation with moderator Caroline Chen (ProPublica) about the role of AI in the future of healthcare. • Presented by the Veritas Forum at Harvard University and the Mayo Clinic. • Please like, share, subscribe to, and review this podcast. Thank you.

Transcript

Welcome to the Veritas Forum. This is the Veritas Forum Podcast. A place where ideas and beliefs converge.

What I'm really going to be watching is which one has the resources in their worldview to be tolerant, respectful, and humble toward the people they disagree with. How do we know whether the lives that we're living are meaningful? If energy, light, gravity, and consciousness are in history, don't be surprised if you're going to get an element of this in God. Today we hear from physician scientist author and founder of Scripps Research Translational Institute, Eric Topol, and the founder and director of the Effective Computing Research Group, as well as Professor at MIT, Rosalind Picard, also Professor of Medicine and Biomedical Ethics at the Mayo Clinic, Jon Tilburt.

Together they explore current trends in AI and medicine and how these two are met with equal parts skepticism and optimism in a wide-ranging talk titled AI and the Soul of Medicine, moderated by award-winning healthcare reporter Caroline Chen, presented by the Veritas Forum at the Harvard Medical School. So for our panelists, I thought we would maybe start with defining some terms here because I think that a lot of people when they hear artificial intelligence, like a lot of different things come to mind, and I think the term just captures so many different things. So I wanted to just start by asking you, when you hear the term artificial intelligence in medicine, what comes to mind for you and to what extent do you interact with AI in your day-to-day work? And maybe I would start with just in the order you guys are on my screen, Ros, Jon, and Eric.

Well, I'm the one non-MD on the panel. I've got the useless kind of doctorate that doctorate of science, but I'm super grateful to all the healthcare workers here. The AI that I actually work on building has had a evolving definition.

I think it's been evolving faster than anything biological over the last few decades. What the founders of the field of AI called AI is very different than what we have now. And so I think the being a little bit flexible about it may be most useful right now, where we just refer to it as things that appear to be done very intelligently by machines, usually taking multiple complex inputs a bit more than we can understand all at once and producing some kind of a decision or action diagnosis or course of steps to be followed in response.

Also, we work on AIs that do lots of kinds of detection of events and forecasting of events in my lab. And for example, I'm very excited. I got permission to mention this, even though it's not being public until tomorrow.

For years, I had this nutty idea that I could look at my wrist and see if I might be able to tell the difference between if one night after you're working hard, you're just feeling run down or if you're actually coming down with something. And now with lots of physiological data and curable studies at Duke and Columbia and a lot of different partners on this funded by DARPA and HHS and Hard Work at Empatica, full disclosure spin out I have found. They've now got a fully running AI algorithm on a wearable that is highly sensitive and specific for telling if tomorrow you're likely to have a viral respiratory infection such as influenza or COVID-19 and it just got CE medical approval, which is the European equivalent of the FDA.

So this crazy dream could an AI do this to an algorithm to now European acknowledgement that it works. That is super cool. John? Yeah, thanks.

AI I think means a lot of different things to a lot of different people. So thank you for being thoughtful about naming definitions as a good starting point. When I think of AI, I think of it in basically in two different tropes.

The first trope is really just high powered predictive analytics in the case of my practice as applied to data sets that have incredible volumes of data. So for instance, like taking all the data inputs from an EKG and predicting an arrhythmia. Then I think there's more of a cultural trope for AI on which we layer all sorts of potential hope streams expectations and fantasies and worries.

So you can think of it as sort of WandaVision medicine and the this more futuristic perhaps still sci-fi orientation of the American imagination that we could somehow engineer ourselves either to a utopian or dystopian future. And probably somewhere in the middle there is a sort of a somewhat more approximate possibility captivated by what Christine Rosen, the cultural critic would call techno solutionism, where if we just have enough data and we just have enough smart people from Silicon Valley working on

stuff, we'll solve society's problems. And so I think multiple tropes of the term get used in often sloppily.

So I'm glad we're talking about that. Well, I think those are really good ways to define it. Carolyn, I consider the subtype of AI of deep neural networks are deep learning to be the main thing of the day.

And there it's taking big data sets as inputs and typically we're talking about medical images, but it could certainly be speech and in the future more text and getting outputs that interpret with machine support through the neural networks to really make the physician or clinicians life so much augmented because it's basically seeing things, finding things, interpreting things that might not be picked up by human eyes because as John mentioned, it's such immense data. And that's just one part of the AI spectrum. Absolutely.

I wanted to ask you guys, you know, any of you who have a have thoughts about this, where in the medical spectrum right now, you feel that AI is either the most mature or where it's growing the most rapidly. And I don't know if the answer to both parts of that question may be the same. You know, when you think about diagnosis, treatment of any sort or say drug development, do you have a sense of where AI has been applied the most, has matured the most is sort of the same across the board or has there been sort of one area of medicine where you feel, you know, AI has advanced the most or where alternatively, there's been the most excitement lately.

Maybe I can start. I think we're still early in all this. I think that's what because of all the interest, the promise is in excess of the validation.

But having said that, I think we have enough evidence with almost every specialty of medicine, whether it's radiology, ophthalmology, gastroenterology, pathology, dermatology, I mean, it's pretty striking how much machines can do to support clinicians. So, you know, I think that part is really moving pretty quickly. And we're starting to see now, you know, randomized trials, prospective trials.

I think the other point that John brought up is about predicting. And there, you know, it's not so much different from what we used to have before, you know, it's like a machine learning fancy analytics and whether predicting better. And I don't know.

But the other thing that we you touched on Caroline is drug discovery. And we have all these claims about these AI discovered drugs. But actually, they're not so much discovered.

You know, there's one drug now in for COVID that's been repurposed that was facilitated by AI, which is good. And actually has a EUA by the FDA using conjunction with another repurposed drug. It wasn't AI catalyzed.

And then we have two drugs or three that have been expedited through AI. So it's starting to make a big impact potentially. We don't have any drug that was truly, you know, developed by AI that's out there in treating patients.

But we do have, you know, jumping through what could be considered two or three years in a matter of two or three months, that kind of thing. So I think we're making really good progress. But, you know, it's nothing like what a lot of expectations or assumptions are.

Ross, did you want to jump in there? Yeah, I would like to add, I think the most stuff is happening where there's the most data and the opportunity to mine it and try to, you know, do things like safety, improvement, forecasting. There's also cases right now where the AI is running autonomously making decisions and helping people today. And these are not as well known.

But I, one of them came out of our lab and out of our company in Patika also. There's other things out there too that do automated detection of a medical event. For example, we automatically detect using multiple physiological signals if a person is having a generalized tonic clinic seizure, if their patterns match those of generalized tonic clinic seizures or grand mal seizures.

Now, if you don't know about those, if you're having one, you're unconscious. So you can't call for help at that point. And actually the number two cause of all years of potential life loss of all neurological diseases is a thing called pseudop that happens after usually after a generalized tonic clinic seizure.

So if, and it's the case that the death rate is a lot lower if somebody's there at that time, but right then when you're having one, you can't call for help. So the device actually detects if you're having one or thinks you're having one. If it's a false alarm, you can cancel it because you're conscious.

If you're unconscious, it calls somebody on your carelessness. And hopefully they know to get there quickly. And then when the human is there, the device doesn't save your life, but the device, you know, summons the human who hopefully can give some first aid.

And most seizures are not life threatening, but in the case that you do stop breathing after one, it's significant to have somebody there helping give first aid. So that's a case where the AI is doing the real time detection and monitoring constantly, right? Better than people do. In fact, it's much more accurate at detecting these events, more sensitive at detecting these events and recording them.

But it's also higher on false alarms. People know false alarms. So the ideal optimal system here is the human with the AI working together.

I wanted to ask one last big picture question here, which is, you know, Raza's you were

talking and describing, you know, the new system that you're working on, you know, are you run down or just tired or that it's so personalized to you and same thing with the system that you were just talking about. My question is, you know, with AI, do you think that most of the goals are driven towards more and more personalized medicine? And I know that Eric, this is something that you think about a lot versus, you know, improving population health at a wide scale, because we have abilities to look at, you know, large data sets broadly. And I'm curious, when you talk about this sort of like WandaVision, I often think about, you know, health insurance companies, you know, wanting to improve health across, you know, their whole pool, like for them, arguably, they want to see sort of broad volumes of patients improving.

That's what's going to help their bottom line as opposed to, you know, single patients here and there. So I'm curious about your thoughts about AI's goals. Do you see them more and more being drilled down towards having, whether it's diagnostics or tools or treatments that increasingly help individuals being sort of the dream versus having, you know, vast population health shifts, you know, where we can sort of raise the bar for everyone somehow.

What do you hear from the people you talk to? I guess my response would be that it feels like a bit of a bait and switch, which is that it's sold as a personalization strategy, but then it's actually propagated within sort of the medical industrial complex as something that's going to be cost effective and efficient and scalable. And then it's, and then the marketing of it sort of continues to say that it's personalized. You know, every generation has their new sort of fascinoma about how medicine is going to be revolutionized, including, you know, evidence-based medicine, personalized medicine, embryonic stem cells, you know, the Affordable Care Act.

And all of these things have a place. It just seems to me that they're more incremental and more, they're more, their added value is a sliver. And to me, diagnostic uncertainty is a small part of the job of doing medicine.

And I don't want to ask I don't want to outsource touch to to a robot. And I don't actually even want to outsource conversation. One of Raza's colleagues, Sherry Turkle, has written about reclaiming conversation.

And I think that's the real exciting thing about what they're doing at MIT is really, you know, infusing the humanity into technology. And I think that's what we should be focused on. But I guess I just worry that sometimes we we sell it too hard as personal as when, in fact, the thing that's going to make it live or die has other drivers underneath it.

Eric Horoz, do you guys have thoughts about that? I'm suddenly struck by the irony of the term personalized because usually it just means that it's going to adapt its behavior to you. Right. And we do that, for example, in our mental health forecasting machine learning algorithms right now because it just goes better results.

If we know that, for example, when you're having more social interaction, it might improve your mood and somebody else that much social interaction might decrease their mood. And you know, these things are very individual, the kind of interaction when it happens. What the algorithm can do is learn the patterns that are associated with certain predictions for certain people.

But actually, there's no personal, you know, there's no person, right, in the algorithm, in the thinking about you, in the caring about you when it does that. So it's really more just individual adaptation than person. I, to get really personalized, there still seems to be a need for a person.

And it's interesting, like even when we've built chat systems, right, and we can craft the computer to sound empathetic, right, and people say, wow, sounds like it really cares about me. And after a month of interacting with it five minutes a day, the working alliance inventory that's usually measured between a patient and a therapist or somebody, a coach, somebody who's helping you. But the working alliance between you and the chat bot actually goes up, right, and shows bonding and caring and feeling cared for.

And yet people know that it's just an algorithm generating it. And recently, there was a really cool study, Rob Morris and team did published in JMR on, they gave the empathetic responses to people, told them in a randomized control trial that some of them were from the bot and told them that others were from the peer, we've been trying to use the AI to bring peer responses to real human responses. And when people were told the same response that the human generator was actually from the bot, they didn't like it as much.

They didn't think it was as good, even though the content was identical. So just knowing that it came from a real person gives it a real boost in the quality that you attribute to the response. I just added a couple of things I had Caroline on the front.

So I never have liked this term personalized because, you know, that's like a monogram or something. It doesn't make any sense to me. So I like the term individualized.

And I think it has a double entendre. That is, firstly, you know, we talk about the clinicians getting benefit from AI, but actually another layer dimension, which is equally as important as the users like what Roz has mentioned, whereby you can now, not in the US, but in other countries, you can get an AI kit at the pharmacy to determine whether you have a urinary tract infection accurately. You can diagnose whether your child has an ear infection without doctorless.

And you can, you know, get skin cancers and skin lesions diagnosed through algorithms. And, you know, we have one FDA-approved deep learning algorithm, which is detecting atrial fibrillation from a smartwatch. So, you know, there's another site for helping

individuals with AI.

Now, the other thing that's notable is that every one of us is unique. And so most of medicine is we treat everybody the same. They have the same diagnosis or, you know, we give them the same drug.

And the point here is that we have the newfound potential that no human can do of assimilating all the data for a person at every layer, you know, biologic layers, physiologic anatomic environmental, the whole shooting match, and giving them feedback as to how to prevent an illness they may be at high risk for or better managing a condition that they already have. That to me is the real deal individualized medicine. We haven't gotten there yet, but we can see a path that we will get there eventually.

So that is exciting. And, you know, that I think is what's in store in the years ahead. The way medicine moves, it'll take a long time, but it's it's it's coming.

Yeah, so you guys have brought me straight into this topic that I wanted to talk about, which is what is replaceable or not replaceable in the clinical encounter, and how what can be enhanced and, you know, what you feel can't be enhanced. So just exactly what you were saying, just now, Eric, you know, there are ais now that have been approved by the FDA or that can maybe be even more accurate and do what a clinician can't do. I'm curious, you know, John, as as our clinician on board here, what if anything, would you say you would not you would feel as missing? Like, if you had a device that could tell a patient more accurately, then, you know, a human could that they have an infection or they have, you know, a certain condition.

Does that count as a clinical encounter to you? Or when would you be comfortable with saying they don't need a human interaction versus that's unacceptable to me to not have a human interface? Sure. Can you kind of talk through your thinking there? I probably would answer this differently a year ago, to be honest, because we've experienced such a catapulted transformation in telemedicine in healthcare, and there's a bigger chunk of telemedicine that I think works that I thought it would. So I'm just be honest about that.

And I think there, and for the part that in which the cognitive burden of estimating risks of adverse outcomes, I don't think there's anybody arguing that we want the best tools and play analytic tools in place to do that, to do it reliably, to do it transparently, and to do it with the kind of high quality, you know, output that we would expect and maybe higher quality than humans can do, because if you can't do it better than a human, it's harder to blame the machine at the end and hold the machine accountable, right? You have to hold their programmers accountable. So I think for the part of medicine that's diagnostics, great. I think there's no question.

For the kind of patients I see, so last week, I saw African-American mom, age 40, college

graduate, three kids, three jobs, insomnia, chronic pain, taking care of a dad on dialysis. For that kind of patient, it's not clear to me exactly how much AI changes what I think I should do. And even if there was sort of a big hero six kind of robot that could help me touch that patient even when I can't, then all of a sudden I'm missing out on learning from her and growing as a human being, even if it's slightly more convenient for her and maybe she produces a scooch of oxytocin by having the machine touch her and she has just as good a psychological state.

I think the practice of medicine is still slightly impoverished in that circumstance. But of course, I might be wrong because I probably was a bit of a curmudgeon about telemedicine and it works pretty decent a lot of the time. Yeah, Ross, as you have been inventing and working on various products, how do you think about building humanity or the human process into the things that you work on? Is this something that you guys have conversations about as you invent devices? Yeah, my thinking also has changed a bunch on this over the years.

When I was first attracted to AI, everybody just wanted to build the greatest machine that humans could build. And the human mind was the final frontier to build something as great as our mind. And the more I learned about the human mind, the more I learned about emotion, the more I learned about how people figure out what matters to them and how respecting that really matters.

And when we build something that simulates what people do and gives them exactly the same input stimuli, but they don't respond the same way as when it's a real person. You have to respect that, right? That there's something different and unique that people have that all the stuff we build, even if the outputs textually are identical or visually are identical, there's something missing there. We don't know how to build it all.

It's very humbling as the more we learn about what are humans. We're quite unique. And at the Media Lab, where I work most of the time, that's a picture behind me right now, I'm in my office at home right now.

But there we are. Focus is on building the kind of future that makes people's lives better. And yes, that means advancing a lot of AI, but advancing it not to make the biggest, baddest, strongest, most powerful AI on the planet, rather to think over and over about what actually makes people's lives better, and then to build AI's that serve that purpose.

So we've really kind of flipped it on its head over the last couple of decades. And it's really more thinking about extending human intelligence and extending human health and extending human capability and extending human empathy and compassion. And Eric, you have this whole book that's called How AI Can Make Health Care Human Again.

Were there anything that, as you worked on that, you found surprising or unexpected as you were working on that topic that you hadn't expected as you started on that? Well, I

mean, really, Caroline, it was the impetus to do the book, which was the kind of captain obvious was we can make things more accurate because medicine is solely inaccurate, shallow medicine to rescue the field, which is notorious for not bonding the human bond. That is, when I was finishing up med school in the late 70s, which that defines me as an old dog now, but medicine was very different. It was a precious intimate relationship between you and your doctor.

And it was a very different look. And that happened over those decades, which I knew was that it had eroded that relationship, that human interaction, because it became a big business. And things like electronic health records and physicians and all clinicians being data clerks.

And then there was these RV units that you would basically be like billable hours for lawyers and all this stuff was happening. And it was destroying the relationship. And so what was the epiphany for me was, hey, you know what, we could get this back.

And so that really was the reason I did the book and basically it gave a chance to review what we can do with AI in the short term. But what we need to do with it in the longer term, which is restore humanity and medicine. And also what played a role was just as an experience, where I opened the book, which is going through a knee replacement, which was almost five years ago and still never properly recovered.

But I was in touch with the just devoid of empathy with the people that were caring for me. And I can't even use the word caring. So that's really, you know, her first-hand experience adds to that as well.

Yeah, I'm sorry, you went through that experience. My last question was just about bias and sort of, you know, we've heard a lot about bias being built into algorithms, whether intentionally or not. I guess my sort of two-part question here, because I don't want to presume either way.

So have you guys seen any examples of bias, you know, in algorithms in medicine? Or conversely, are there ways that you think that algorithms can actually help prevent bias or lower bias and help you actually bring the principles of equity and help to enforce them in medicine? I would say we probably don't even agree necessarily always on what equity looks like. And so enforcing them sounds a little dubious. But I think there's broad support for understanding what are the ethical means to move forward with new technologies and thinking about the similar principles and things like that, where large groups of thought leaders across the globe have kind of come together and tried to outline what those principles ought to look like.

Getting there, of course, is a harder thing than just articulating them. And I think there are real challenges there. I also think that we shouldn't lose sight of the fact that it's hard to articulate the appropriate sort of ethics of the means when we can't articulate

the ethics of the ends.

So to me, the harder questions lie with the ethics of the ends, like what is this for, what's the meaning of medicine? How does our pursuit of incremental prediction quality fit with the overall goal of what medicine is supposed to be? And if we can't answer that question, then all of those questions about the means to me are super important. They're not necessary, but they're not sufficient to answer the question of sort of how do you approach AI ethically? I'm really hopeful because the algorithms are, well, sometimes they're hard for people to understand they're transparent, right? You can run in a biased dataset or an unbiased dataset with different kinds of bias to find your flavor, and you can see exactly what happens. And that is really helpful, right? Whereas with a person or some other systems, you know, all bets are off, right? Whether they're acting in a biased way or not.

So it's testable. In some cases, it's very transparent. And there's a whole lot of work in AI right now going on toward explainability and transparency.

And I think with so many sharp people jumping on it, I'm very hopeful that the AI will be less biased in the kinds of ways we don't want to see. Although on that caveat, I want to just say it's all biased, right? There's always some kind of bias built in. The harder question I think are the ones John may be addressing, which is what are the outcomes we want? What are the optimality criteria that we are optimizing these for? We can optimize for all kinds of things, but what should we be optimizing for? Yeah.

I just want to remind very quickly, remind all of our audience that we will be taking audience questions not too long from now, probably in about 10 minutes. So if you have questions, you can start putting them into Slido. But I do have just a couple more questions for you.

So I don't know if I'm the average patient because I am a health care reporter. So I think I tend to ask a lot of questions to my doctors. And I like to ask, you know, how they've come to a diagnosis.

I will ask them the mechanism of action of the drug that they're prescribing to me. I tend to ask a lot of, you know, nosy questions. But I do think that I like to know that my doctor understands what they are saying to me generally.

So my question is, when we start to lean more on AI, how important do you think it is that a doctor understands what's sort of going on under the hood? And I also think about this as like a journalist who really is kind of obsessed with fact-checking, right? So like how much, both from the standpoint of like being able to explain to a patient, you know, this is how we came to the decision either that, you know, you have this diagnosis or we're going to give you this treatment, you know, that goes beyond because the algorithm said so. But also from a trust standpoint, you know, like if I was a doctor,

would I say like, I trust this just because, you know, some chief medical officer in my system decided that we're going to buy this software. Like to, to what degree does a doctor have to really understand in depth what is going on? How important do you think that is? Because I don't think in med schools, are they going to be like, here's how you fact check the algorithm that you're going to depend on.

So I'd love to just hear your guys' thoughts about, you know, the depth of understanding that clinicians should have or how they can sort of come to terms with something that they may not be able to understand deeply. So, I think that human in the loop story is really important that you don't want to make any key decisions and just trust a neural network, an algorithm to, you know, basically lead to an automated diagnosis and then treatment. So that's why oversight is critical.

And then that gets to your point, which is there should be familiarity with the nuances of AI. And that should be part of every medical school curriculum and training and postgraduate education. It hasn't started yet here, but it should because we're already, you know, in the early phase of seeing these getting integrated in care of patients.

So you don't, kind of rule number one is you don't trust, don't trust the AI implicitly. Like you don't fit, you want the facts checked? Well, yeah, same sort of thing. But on the other hand, you kind of touched on this another way, Carolyn, which is the idea about the explainability.

And you, I think, you know, when you got to a point earlier, when you were saying, do we need to have all these algorithms fully explained, when you ask what's the mechanism of action of a drug? When we don't even know what the mechanism action of many drugs are? Okay, we have no idea, we don't have a clue and we use them every day. So the point here is that, you know, maybe it doesn't have to be fully explainable. We're seeing more and more AI deconstruction of neural networks to use that as a reverse engineering to understand what's going on.

But if you don't if you can't explain it, and you don't understand what is going on inside this box, that's not a good thing. So we have to get education going in that regard. I would think of this as sort of a chain of trust, where you've got this sort of the distal distal user, the patient, their family, etc.

And you've got a practitioner with whom they interact in a health system and a really smart electronic medical record and a lot of expert recommender software doing stuff. But behind all that, there has to be sort of, I would say, layers of sophistication and confidence that at each layer, those the people to whom say, say, I as a practicing physician, am just smart enough to understand what the sort of the AI output programmer people do. But those people are still those people are connected to even sort of more sophisticated science about sort of neural networks and everything.

And if any of those links in the chain of trust are broken, you risk a lot. And you know, doctors don't even know the test characteristics of a typical test, and nor could they explain it. So for them to be sort of computer science with kids, and have the emotional intelligence that we hope we can really claim when AI is doing more work for us, I think is unrealistic.

But I agree that we've got to try to move in that direction, because the technologies are already out in front of us. Great. Rose, did you have anything to add to that before we turn to audience question? Just very briefly that I encourage people to don't just stop at accuracy rates, go out, ask what kinds of data was it trained and tested on, and what kinds of errors does it make.

And I think that's helpful for getting a little bit of sort of input output and level understanding without having to know all the details, what's going on inside. Great. Can I make one more comment? Oh, yes, yes, please.

I love talk about trust. Colleague and I wrote a paper about trust in AI probably a year ago, a year or two ago. And really, I think that it represents the way in which we can either flatten or distort the use of words.

So trust goes way back in medicine, and it means a particular thing, and it means something related to human-to-human relationships. And even applying the word trust in its most robust sense to human-machine relationships in some ways is either a category error or a flattening of the full robust meaning of the term. And that kind of sort of potential for distortion is the kind of thing that when we get swept along by the sort of the excitement about what a technology can do without looking under the hood, as you say, we tend to our vocabulary implicitly drifts before we realize what we've done.

And then we conflate the full orbbed meaning of trust with like reproducibility. Well, there's a lot more to trust than just reproducibility. There's this idea of entrustment.

And nobody would say that we should entrust ourselves to a machine. And yet, when we start using the word, there we go, right? So, just a point about sort of the lexicon of popular life, I guess. Yeah, thank you for that.

I never considered that before, so I really appreciate you bringing up that. I'll be more careful to use that word in future. What is going to turn now to questions from the audience.

And as we are taking these, please feel free to continue to add your questions or to upvote questions that you would like to be asked. And I'm going to just take a minute here to look through these questions. So, there are a couple questions here that are related to the workforce.

And so, I'm just going to combine a couple here. So, are there... Do you guys have any

thoughts about whether or not AI will be replacing the workforce, the physician workforce, you know, E.G. will AI replace doctors, and would there be any particular areas of medicines that might be relegated to AI to a greater degree or a near timeline than others, for example, radiology diagnosis? Well, that's an interesting question, because I spent, oh, I don't know, a year and a half or more with the UK. The NHS commissioned me to oversee a team of about 50 people in the UK about their workforce and the future of AI and its influence and other, you know, digital technologies.

And we came to the conclusion that no, this would not reduce the workforce. There's no physician clinician that's going to be replaced. And you can read the report.

It's a pretty, you know, it's a succinct report, but it's got a lot of, I think, supportive data and evidence in it. But what we can see is a reduction in growth of the clinicians, because right now we have unchecked growth in many countries. This is the only area of employment, healthcare workforce that's growing out of proportion to everything else.

And so, while we like to stimulate jobs, we're overcooked here, because that's what leads to the profound expense in healthcare. It's a labor force story. So we want to use AI to basically do a lot of back or back office functions, the coding and billing and a lot of stuff like that, not front line medicine.

And then what you the term you use about leaning, leaning on machines, yes, because we can reduce the reliance, you know, the full reliance of clinicians. So I guess the answer to that question is, will it be a long term workforce influence? Yes. So for example, if you get rid of keyboards, what we call keyboard liberation, which is one of the most important short term needs that AI can get us there.

I'm confident of that. But that will just redeploy clinicians to do what they want to do, which is take care of patients, which is why they went into medicine in the first place. So you see this kind of shunting to better activities, but not the the net reduction of people's the need for the talent and the people.

There's another question here that I'd love to hear all of your personal perspectives on, which is, I'm curious about the idea of healing, be it physical, spiritual, or emotional and bore. What do you think is the role of AI in all of these spheres? I can tackle that first, I guess. My bias perspective, which I probably should come out with beforehand is a white male Protestant Christian position.

And I think medicine matters for a variety of reasons. But to me, it matters most deeply because I think it's marrying greater realities. Greater realities related to kind of relationality that's deep in the nature of things.

And that to the extent that we accompany and console our neighbor in a finite human journey, we ourselves become more united with the deep nature of things. And I don't

ever want to outsource that because as soon as I do, I lose something in my own becoming. And I hope that the path of AI helping medicine is not the same path as the electric laundry washer or the toaster, which was the marketed hope of having free time to do all sorts of amazing things.

But then just cramming our lives full with more technological pre-occupation and allowing our our neighbors and our friends and our family to die alone in the ICU in the case of COVID. That's a tough question. I'm optimistic about a lot of things with AI, but I'm not grapple and grasping to come up with examples where it's actually been healing.

It's very easy to come up with examples where it drives us nuts. We're very frustrating, where it just does catastrophically stupid things. And we worry about those things.

So I still think the safest thing is to couple it with a smart human and look for ways that can help the smart human get the load of laundry done, but not falsely promise something that we're not going to have happen. I mean, it's all in the indirects. If you get people to be more autonomous with their health and you help clinicians not having to review in order amounts of data, or you tee it up for them, and you have time with people and you're doing the things that you really want, which is caring for patients, all that are indirects towards the human interactions, the bonds, the presence, trust, I'll say trust, and I think the empathy.

So it isn't a direct, it's a lot of indirects. Yeah, thank you. There's another question here that I think John might be interested in as our resident ethicist here.

Do physicians have an ethical obligation to use an AI-based tool if it has been shown to improve outcomes? That's a big if, but I think to the extent that providing reliable, diagnostic information is part of the job. Yeah, that's part of the job, but it's just part of the job. Yeah, I think we're going to see when things become a standard of care, and it isn't being used in an individual, it will be a medical legal precedent.

So just to become clear, I'll give you an example. About half of people with diabetes never get screened for retinopathy, a leading, if not the leading cause of blindness. Now we have deep neural network systems in grocery stores where there's no ophthalmologist, but just basically someone that can have a person get their eyes imaged, the retina, and then get cloud algorithmic interpretation accurately whether they have diabetic retinopathy to prevent blindness.

Now as that becomes, you know, routine and that becomes something you could do on a smartphone, you know, something that becomes basically of little expense, and someone goes blind because their physician never did it. You'll see a lawsuit, that kind of thing. So when things get validated, that's an example of a prospective well-done validation.

We don't have that many of those, by the way. There's like about like 12. Then we'll start to see as it gets, you know, basically embedded in medical care, there'll be challenges about why didn't I get that care? You know, why did I get a wrong diagnosis of my scan that led to this unnecessary operation kind of stuff? So we're still early, but when things get really well validated and change practice, that is standard of care, that's when you start just to see the questions come up.

John, were you going to add something? I guess part of the, to me, the challenge of sort of racing ahead with AI is it's sort of, in some ways obfuscates us from fixing the problems of the present and puts our hope and attention and energy elsewhere. So I think there's this fundamental irony, which is we're asking human beings to become more robotic while we're trying to humanize machines. And we're not really fixing that problem.

We're actually outsourcing in a broken system things that we haven't fixed. And I think Eric's a little bit more optimistic than I have that will get to the point where we can reclaim that sort of humanizing part of baguette medicine with an incremental add of AI on top of industrialized medicine. I would say forever, it's like COVID, right? If we had spent one tenth of dollars on national security that we had spent on public health infrastructure, we'd be in a different place today.

If we spent one tenth of our energy and attention in dollars on humanizing healthcare instead of improving AI analytics for healthcare, I think we'd be in a different place. We don't have a chief empathy officer at Mayo Clinic, but we have a president and CEO of the Mayo Clinic data platform. And that tells you what's important.

And unfortunately, I think it becomes a distraction and it keeps us from fixing our other problems. I would like to comment on that because I think this is really quite important. The medical profession has not ever stood up for patients in itself.

It has no representation. Basically, it's represented by trade guilds like the American Medical Association and all these other societies. And so, you know, I wrote a piece in the New Yorker.

Now, it's not quite a couple years ago. And basically, what it's about is that we, this became a big business because we let it happen. We never stood up.

We never had solidarity or representation. These trade guilds all they care about is maintaining reimbursement for their constituents, which is not the issue. So, if you have AI in full gear and it's now, you know, the overlords, which are the administrators and health systems like the ones John just referred to.

And he has this passive stance that, you know, you can't do anything about it. I have a very different view about that. But if you let it happen, the overlords will say, well, see

more patients now that you have machines to help you read more slides, read more scans and on and on.

So, there needs to be a revolt against that to use the gift of time for care. That's what has to happen. Otherwise, the default mode is it will get worse if that's even possible.

We have a global clinician burnout. We have the highest levels of depression and suicide in the history of the medical profession. And if we don't stand up for patients and for restoring the humanity in medicine, then we get what we deserve and our patients get what they deserve.

They don't deserve. So, that's why it's time to not be passive anymore. And that's something we have to, you know, the idea that we just let the overlords continue to rule the roost that time has passed.

That was great. I think you guys should write a joint manifesto. The preacher has a great use.

Well, thanks for articulating that because, yeah, I've been hearing this from all friends who are physicians also that it's just a huge and growing problem even before the pandemic. And I'm, you know, I think the stresses of everything recently has only made it worse. And it's great when a leader like you speaks out and says what you just said.

I think we need to figure out how to convert that into the change. And then the ideal that people dream of needs to be articulated. And then we need to be building that future and optimizing for that.

Anything's optimal given the right criteria. You know, this is this problem of what are the criteria and if the criteria is real one-on-one human time together, physician and patient, and that's what's most healing, then, you know, we need to optimize for that. Absolutely.

I nominate Dr. Tilbert for Chief Empathy Officer. We should start a campaign. I'm just keeping an eye on time here.

We'll probably have time just for a few more questions. So I'm just going to quickly look through here and don't want to want to make sure that we have the most recent ones as well. You know, there's some questions here which I'm just combining some questions, which relate to particular populations.

So, for example, what was your opinion on using multi-competent robots in elder care and someone else who asked about a patient with stage four lung cancer, how would AI address care. And I'm just curious, as you think about particular populations that you have either worked with or, you know, Ross and Yuri Keiske created for, are there particularly vulnerable populations or special populations where you think AI has a particularly good opportunity to work with or conversely where you would say this is the

time for AI to back away from? I don't think there's a simple answer for even a whole population or type of people I got started doing. A lot of the work we're doing now got started working with people on the autism spectrum.

And many of them came to us seeking technology that they could wear or use to augment their abilities to read the emotions of others, interestingly, to interpret your facial expressions while you were simultaneously talking. They were overloaded by the 10,000 different patterns that were assaulting their visual system. And they wanted help with that.

In a population like that, you know, to help them connect to their physician, if you will, you know, this could be useful. However, to just, none of them were actually seeking to replace people, though, with this technology, although they, many of them love to interact with robots. So we, again, we keep finding there are these sweet spots where it's helpful.

Many of us, I know, who have elderly loved ones that we haven't been able to access because they're in a care facility during this pandemic. You know, I don't know if anybody else has, but I'm, you know, I kind of wished I could have had some big fluffy thing, give my mom a hug, right? You know, I give her this virtual hug over Zoom. Mom, I'm giving you a big hug and squeeze now.

Close your eyes. Can you feel it? You know, it's, it's not the same thing as being there. We have to be creative.

If a robot just walked up and did that, you know, even with a person with memory stuff, it's just not the same, right? We know that. And yet, just like telehealth has replaced a lot of face to face now, it's, I don't think it says good, but it has some advantages, right? It has some features. So I think we just need to be very creative about looking for where the technology can help, but not seeing it as this evolution of humans are going to turn into machines, and that's going to make a better future.

I do not buy that one. Yeah, one thing that the remote monitoring movement and the robots for the elderly has reminded us is how alone together we are and how challenged our society is in terms of community and fragmentation. And so long as those are those remote monitoring devices and the robot and the elderly living room remind us of the painful imperfection of our presence, that they're a band-aid, they're not, they're not permanent healing, then they, maybe they're a good reminder to us.

And I've been working somewhere with Native American communities trying to facilitate care and improving our care in making Mayo Clinic more hospitable and less overtly cold and white to those in our area, in Arizona. And I think if we said we want to do a chatbot-enabled telemedicine, virtual AI-infused whatever with your community, that would be prudentially unwise. And it's not because Native people are categorically opposed to

technology, and it's not because it wouldn't eventually benefit them, but there's just too much history there to lead with that.

So I think it's much more of a prudential question in terms of where we are in the present than some sort of categorical absolute. Yeah, thanks for that example. That's, that's helpful.

There's a question here about liability. So it's one of the biggest hurdles to AI, be it self-driving cars, facial recognition, etc., is being implemented as a liability question. How would that need to change? And I guess this is, I guess liability either for, you know, a medical group or hospital, or even individual doctors here.

I guess another way of saying this, one of you guys mentioned oversight earlier, I'd be curious what you guys think about the role of the FDA as well, or other sort of oversight bodies? Well, we have some problems with the FDA because they have essentially a conduit for the companies to put forth proprietary algorithms that aren't published to the medical community. And then they give them approvals. Usually they're so-called, you know, 510Ks, which are not the ultimate validation.

They're just basically retrospective data sets, which most AI work has been done on. And so they get out there and jolly, you know, for example, radiology algorithms that are out now. And the medical community doesn't even know the data.

They haven't seen it. And, you know, the FDA has done their job supposedly, which get it approved. And it's real, it's, it's really unfortunate.

There's, there's this lack of open science. And so we have a real chasm between what is the regulatory science and what the medical community needs to feel comfortable to have that oversight. So one thing, you know, as you earlier got to is understanding, you know, neural networks.

Another is the one that you're implementing in your health system to care for patients. And you don't even have an ability to get to it. I'm in agreement.

We need a lot more transparency around the data as well as the algorithms, because the, you know, if it was trading down something very different, then what you're actually going to be running it on. And it usually is, right? Then the generalization errors when it's tested in the future could be quite bizarre. And surprising to people who thought they understood what was going on.

And that's dangerous. That's worrisome. I think there need to be these safeguards built in.

It depends how risky the decisions are made on it, you know, how much of that we need. What about the question of liability? Like I, you know, I'm out of my depth here because

I'm more used to reporting on drugs. But I think, you know, they're, they're sort of really clear, sort of labels that come on drugs where, you know, this is what you can use this on.

But I don't know how it works with sort of an algorithm, like you said, like, is the hospital responsible if they then apply it to a set of patients that was not what it was tested on if the company wasn't clear enough. What happens then? I'm not sure you have the perfect experts in the room on this. I think we all agree that there may not be the right transparent regulatory sort of mechanisms to help us establish that chain of trust that society would need to really rely heavily on AI the way we anticipate we might want to.

I also think that my understanding is that AI might put lawyers out of business, so I'm not exactly sure who I would consult. If it were the reliability of AI itself that we were like, there's something really kind of circular there, so I'm not really sure. But there are important questions.

It just seems like it's one additional question about getting the means as sort of transparent as possible. Yeah, there's another. You can't ever let your guard down.

So let's say you have, which we don't have, a grand validation. You know, it's gotten through randomized trials. You know, I'll give you an example.

There's been several randomized trials of colonoscopy, you know, with machine vision, so it picks up the polyp much better than the gastroenterologist could who would have thought the gastroenterologist would be leading the charge in AI, but they they have for randomized trials. Most of them have been done in China, but the point here is that you use the algorithm, you implement it, you put it across all the health system, and then what happens? There's potential for adversarial attacks. There's a chance that there could be just a glitch in software.

So you can't ever let your guard down. No one has really talked about this post implementation phase, but it is yet another dimension of the uncertainty that we have to keep in mind. You know, as each of you are thinking about the next, say the next 10 years, as you are working on your either as a clinician, you know, on the products you're working on, or on your research, there's a question here, which is, I'd love to hear how your faith or worldview informs or enhances your work.

And I would love, you know, I know that's a huge, huge question here, but I would really love to hear how this is, you know, how your faith or your worldview is going to be guiding you as you move forward. Let's say for the next decade in this field, who wants to take that first? Well, so that's a big question, Harold, and that's surprising coming from you. But, well, it's coming from the audience here.

All right, right, you threw you, right? Yeah, I mean, for me, I'm in clinic each week, you

know, I spend a day in clinic each week, and that's what is my chance to be in touch with what the helmet needs of patients. And, you know, I think that helps me in the years ahead to find out how so many of them want to be more in charge and autonomous. And they, they, I think we can get them there.

And I think we'll get of the virtual health coach that gets all their data. It isn't for everyone, but it's for a majority of people that would want that function to help basically coach them to better health or preventing illnesses that they otherwise might get. So, you know, I'm seeing that in the years ahead, we will get there.

It doesn't get there. We're just from deep learning, we need hybrid models to get there. So, it's a challenge even for AI scientists to find the path to do that.

It works like Ross's work, you know, in one dimension, if you want to understand, you know, a seizure or, or maybe a diabetes, a sensor, but to do it all, the holistic view of person with all the relevant data and all the different layers of data, we are ways to go. But I'm excited about that because if people, when people go to see a clinician, it's a, it's a ice pit view of their, it's not the real world, it's for minutes or even an hour or whatever it is, it's not their real, you don't get a high frequency or continuous sample of relevant metrics. So, I'm hoping that medicine will change in that way.

It all relies on massive data per individual and interpreting that. It's a big challenge, but I'm confident someday we will get there. And that will be a segue to a much more accurate individualized medical approach and far less of the mistakes, which hasn't been emphasized yet today.

But there's a very high, you know, five million serious diagnostic errors a year in the United States, at least, which is, you know, an enormous number of serious errors. So, we've got to stop that stuff. We can do better than that.

I seem to be consistently working on problems that relate to people with stigmatized conditions. Recently, a doctor came to me and said, could, could we help build an AI that helps people to be aware of how empathetic they do or don't look, the medical students and the physicians, when they're interacting in particular with patients who have substance abuse situations? He said that in his experience, these patients don't tend to be treated with as much respect. They're stigmatized.

The doctor at the end of the day is like, ah, you know, you, why would I want to, you know, you're, you're hurting yourself right now. And so, while that's certainly not true of everybody, there are disparities. And what I see in epilepsy is a lot of stigma.

What I see in autism has gotten better. But there's so many conditions where it's not good. And I, I grew up as an atheist.

I over time embraced belief in God and then eventually became a Christian. And one of

the things that changed in me was a view of other people as all of us equal, all of us equal, all of us made in the image of God, which is kind of hard for me to grasp at the beginning. But it was a great equalizer.

It was a very powerful equalizer. This is smartest, most accomplished, most educated, you know, coolest, whatever was on, on par with the one who is stigmatized or picked on at school or whatever. And I didn't realize how profound this was until I was speaking in Beijing to a group of people.

And I was talking about our autism work and our epilepsy work and all the cool things that's led us to learn about the brain and what a privilege it's been. And one of them said, how can you say that? How can you work on that? And I'm like, what do you mean? Like, I, and the translator was helping me, like, am I misunderstanding something here? And this person had a different worldview. And their worldview was from their religion that these people with these diseases were being punished in this life for something they had done in a previous life.

Now, this was a person at the top university in Beijing at Xinhua, and they called themselves the MIT of China. And they were, they were quite sincere that their worldview was that some people with certain diseases are not as, you know, they are rightfully being punished for something that makes them inferior. And therefore, we shouldn't even waste our time helping them.

So this was just shocking to me. And I realized this is a worldview choice. It's a worldview perspective.

And, you know, the Christian worldview is very clear, or Judeo-Christian, you know, God, all of us made in the image of God, and all of us equally loved by God. And that has inspired my desire to try to rectify some of those inequities that are out there. Far from achieving it, but it's a goal that I really enjoy trying to work on.

Thank you. John, do you want to close us out? I love the Simon and Garfunkel song "Homeward Bound." I think it's a great encapsulation of the loneliness and longing of the modern secular human soul and this desire that I just need someone to comfort me. And for me, I think of the C.S. Lewis quote, this is something like hell is being given over to yourself and the kind of lonely despair of sort of that really sinking emptiness of wondering whether there's nothing else and there's no one else.

And for me, the Simon and Garfunkel song hints at it and at least the confessions of the tradition that I know at the best say that my only comfort in life and in death is that I'm not my own and that I belong to another and that relationality is the deep nature of things. And for me, that gives me a lot of hope and hope that the kinds of conversations we have like today can continue. If you like this and you want to hear more, like, share, review, and subscribe to this podcast.

And from all of us here at the Veritas Forum, thank you.

(gentle music)