

Will AI Replace the Doctor—or Restore the Art of Medicine?

Richard Low, MD | Association of Yale Alumni in Medicine | October 3, 2026

EDITED WORKING TRANSCRIPT — Based on the uploaded YouTube transcript. Obvious automated-caption errors have been corrected where identifiable; punctuation, spacing, and paragraph breaks have been improved. The text otherwise retains spoken wording, including repetitions and informal expressions. Check against the recording before quoting verbatim.

Video: <https://www.praxisemr.com/richard-low-yale-ai.html>

Introduction

[0:09] Wilhelmina C Korevaar MD, President, Association of Yale Alumni In Medicine: Okay, now it's on. Time to begin. Hello everyone. Hello. It's time for us to start our general meeting. If everyone could please sit down. I'd like to welcome everyone to the Association of Yale Alumni and Medicine general meeting this afternoon or this still morning. This morning we're going to have a speaker for the first portion of the meeting. Dr. Richard Low, 76MD, has flown here from Argentina to talk about healthcare innovation and entrepreneurship. The title of his talk is, "Will AI replace the doctor or restore the art of medicine? Lessons from a Yale education." A little bit about Dr. Low. After finishing medical school and an internship, he went to California where he [1:52] established a network of locum tenens emergency room physicians. He realized that medical records were important in ha physician communication one to the other and became involved with electronic medical records which he found quite disappointing. And so he came up with a concept of a neural network intelligence that would learn from each individual provider. And in a way, although I'm sure that this is my own brain's oversimplification of what he's done, uh, he kind of has created an interface that I would say is similar to how an individual lay person interacts with ChatGPT to try to figure out, you know, what their conditions resemble in terms of a standard condition that ChatGPT recognized. izes [2:56] and then over time through interacting with ChatGPT say that patient uh begins to develop their own better idea about what is going on and again forgive me because I I did try to study everything that Dr. Low has written about how his uh neural network intelligence system works but a lot of it was frankly over my head. So he started this in 1989 with a concept processor that was based on clinical concepts and went on to develop a an AI based medical record. So 1989 when you listen to Dr. Low. I am sure you're going to be inspired. You're going to be asking yourself, for example, as I did, well, gee, what how come I didn't come up with something like that? But again, I want to [3:57] emphasize he started working on this in 1989, so a long time ago. And I think that is one of the lessons that I've learned by reading about Dr. Low. It takes a long time for ideas to percolate. It takes a lot of clinical experience in the interim for those ideas to become relevant and it isn't until after a lot of time has passed that you can actually bring those ideas to fruition. So here to tell us more about AI and medicine I present Dr. Richard Low. Thank you.

[5:9] **Richard Low MD:** I'm truly embarrassed after this presentation that I just received. I thank you very much uh uh Wilma for inviting me back to Yale. It's very emotional for me. So if I don't say things properly, please forgive me. I don't give these things too often.

[5:10] Uh I am extremely grateful uh to have you here. I I thank uh Jaideep Talwalkar who happens to be here. He's very involved in health IT. In fact, I found out there's a lot of people here who are incredibly involved in health IT and and they're doctors. Thank goodness, medical doctors. We'll talk about that. Uh I'm very grateful to Dean Brown for her support of this gathering and for really taking us alumni as a part of Yale. Uh and I will obviously forget to mention some of you and I apologize if I do but I thank everyone that is here in the room today. My my my great friend Robert who is in charge of the international program here at Yale. It's too many people to thank. Obviously, I want to thank and [6:08] Arthur, Dr. Arthur, I don't know where she is, but she she put the Oh, there you are. for putting this thing together. You don't know how much effort that takes, but they did a great job. This has been an incredible week. Uh, and I there's a special pleasure to thank all of you who are on the Zoom today for being there. uh my work developing Praxis electronic medical records has obviously shaped my thinking but I won't be discussing Praxis today okay I will be discussing the ideas behind it here I'm supposed to press this button can you see can okay thank you beautiful Right. So, uh, when I arrived at Yale Medical School 54 years ago, one of the first people who spoke to our class was Dean Lewis Thomas, Dr. Thomas. As I remember it, he told us something that I have never forgotten, and I bet you many of you remember what he said because he must have repeated it for more than one class. There is no such thing as a stupid question. There is a stupid person who doesn't ask questions. Ask, ask, ask. Did anybody else hear that here? Uh, okay. Well, okay. One person did. I guess that tells you how old we are. Uh, uh, that is Yale. As you know, we had no conventional grades. Attendance to class was largely our responsibility. We were surrounded by extraordinary people in medicine, science, and the humanities and given an unusual freedom to decide both what we should learn and how we should learn it.

[8:05] But what Yale was really teaching us was something more important than anatomy, physiology or pharmacology. It was teaching us to question assumptions. And so more than half a century later, I want to question an assumption that has been running throughout medicine.

AI Replacing Doctors Is the Wrong Question

[8:25] Will artificial intelligence replace us as physicians? Um I think let me take this out. I think we're asking the wrong question. Let me turn that question around. Can artificial intelligence free us again to be doctors? Because as we all know, for the last 50 years, something has already been pulling us away from the practice of medicine and is not artificial intelligence. It's documentation is the hideous coding, CPT, ICD-10, SNOMED, RxNorm. is insane. It's prior authorization. I'm sure many of you have had to fight with that insanity. Uh it's quality reporting. Oh, I love that one. We could talk I I could talk about quality reporting for more than an hour. Is the endless justification.

[9:34] I mean, it's insane asking one of doctors who spent so many years to justify to insurers and algorithms what we had to do for a living. uh is the endless bureaucracy that requires us to justify our clinical decisions to insurers and increasingly to algorithms entities that have never seen our patients much less cared for them. Ironically, the computer has been one of the main forces drawing us away from medicine. I say this as someone who has spent more than 30 years developing medical software and who has always believed that computers could do precisely the opposite. And I wasn't alone in believing that we have today an amazing doctor uh that should probably be giving the talk instead of me. Uh Dr. S. Yin Ho [10:25] uh she's also a class uh an alumni 96. Uh she also has a masters in business. I mean she she's worked for many large electronic medical records software company. In fact, she was, you know, also the CEO, uh, of all scripts for a while. Um, uh, Dr. Ho wrote a book on electronic medical records that I encourage all of you to read. All doctors should be reading her book because it's not about computers. It's far more fundamental than that. In fact, she's got copies of the of the book and I recommend that you get her book. uh she reminds us that from the very beginning there were two fundamentally different ideas about what computers could do for doctors. The dominant model treated computers primarily as a

[11:21] repository of data but a smaller group many of them led by physicians like us imagined something far more clinically meaningful. The EMR could be a veritable medical instrument that could extend our memory, support clinical reasoning, and adapt to each individual doctor's way of practicing medicine. Unfortunately, as Dr. Ho so well explains in her book, that physician led vision did not prevail. The priorities of third party payers and the government did. What took its place was a healthcare IT industry led largely by technocrats with limited clinical understanding and increasingly shaped by bureaucratic requirements, coding systems, insurers and regulators

[12:04] rather than by our needs as doctors and that of our patients. And as and as electronic records became more universal, the forces shaping them also changed. Coding mattered, reimbursement mattered, regulations mattered, quality reporting mattered, government certifications mattered, and our needs as doctors and patients mattered less and less. Increasingly, we physicians sitting across from our patients were no longer the ones deciding what an electronic medical record should do. Those decisions were being made larger by insurers, regulators, administrators, technologists, and of course, the government. All far removed from our clinical encounter. But this outcome was [13:01] a choice. It was never unavoidable. Indeed, Dr. Ho was part of a small group of physicians who continued to insist on a very different relationship with the EMR, a technology that should adapt to medicine and not force medicine to adapt to technology. But the industry chose a different path. It began by redefining what the medical record was and for whom it was meant to serve. Even the language began to change. We originally spoke of electronic medical record or EMR, a tool designed to serve physicians and their clinical staff. However, eventually the dominant term became electronic healthcare record or EHR. This subtle change in wording even sounded beneficial. The health IT industry presented the new EHR as something broader and more complete than [13:50] a simple EMR catering solely to doctors and their patients. It was neither. This word this change also marked the fundamental shift in purpose. The medical record was no longer expected to serve primarily the doctors caring for their patient. It was now expected to serve everyone. insurance companies, regulators, certification programs, quality reporting programs, public health agencies, patients, the government, and an expanding collection of other parties who came to be known as quote stakeholders. And I don't know why I hate that word, but I do. And yes, the dominant EHR became increasingly effective at satisfying bureaucratic requirements.

[14:32] Clinic and hospital administrators found them useful. Doctors, well, we all know what happened. What the EHRs did not necessarily become better at was helping physicians practice medicine. Doctors began to spend hours in front of their new EHRs about double what was the case on pen and paper. What began as the record of a personal confidential encounter between a clinician and a patient gradually morphed into an instrument serving an entire healthcare bureaucracy micromanaging our medical care. The new EHR was simply the instrument by which this could be accomplished. And of course, a system designed to be all things to all people rarely serves any one of them particularly well. Insurers already had [15:15] systems designed around their needs. Regulators, of course, had their own and administrator had their own billing and management systems. What medicine was losing was a medical record designed around the needs of doctors and their patients. The health IT industry solved this and continues to solve the wrong problem.

Doctors Forced to Adapt to Computers

[15:34] Okay. Uh instead of teaching computers to understand doctors, it forced doctors to speak computerized templates, menus, checkboxes, drop-down lists, required fields, codes, structured data, endless clicks. Medicine squished into a mold. The industry did not design technology around how we physicians think. Instead, it forced us our clinical thought into forms computers could process. For decades, we doctors have had to carry the burden of the machine's limitations, and we still do today. This

is awful. But now, artificial intelligence can reverse this relationship. For the first time, computers can work with language, context, and meaning, not merely codes and structured text. We no longer have

[16:36] to conform to the machine. The machine can begin to adapt to us, to how we reason, how we exercise judgment and how we practice medicine. But adapting to us requires more than processing our words. It requires understanding how we transform an encounter with a patient into a medical record. And that brings us to a fundamental question. What is exactly a medical record?

A Transcript Isn't a Medical History

[17:04] Suppose a patient comes to see us and talks to us for 20 minutes. Say an ambient AI system records and transcribes every word uttered. Let's say the transcription is perfect. And let's say that it also adds the appropriate labels and headers and other documentation. Question, do we now have a medical history? No, we don't. We have a transcript. And a transcript is not a medical history. If a transcript truly were a medical history, yeah, then we would not be needed. In fact, a court reporter could generate such a transcript with a bit of training. And the reason for this is fascinating and we may not have thought about it. This is because the patient does not write the medical[17:57] history. We do. We're not simply transcribing words into ambient or from ambient. We listen, we interpret, we question, and ultimately we decide what matters from all that was said or even from what was not said because the patient provides us the information. But the words are only part of it. And here we bring another hero that we happen to have here today. And unfortunately, you can't see the folks that I would like to show you because they're here in the room, but you can't see them. Most of us do know Dr. Lisa Sanders or at least know her work. Anyone who reads the New York Times has almost certainly encountered her celebrated diagnosis cases. You probably also know she's a fellow Yale alumni and a professor of[18:47] medicine here at Yale. As Dr. Ho recounts in rushing headlong, Dr. Sanders put it beautifully. How a patient reacts often says much how a patient reacts often says much more than what the patient is actually saying. That's beautifully said. We combine the patient's words with what we observe, what we know about our patient, what we remember, what we suspect, what does not make sense, what our patient did not say, but what her face, her posture, or even a moment's hesitation revealed. We combine all that with what we have learned from hundreds of other patients before this one. And from all these clues, we then construct the clinical history. This is certainly not a compressed transcript. It is an act of human[19:40] interpretation of medical thinking. And that act of interpretation does not end with the history. It culminates in the diagnosis. And here's another surprise that you may not have thought of. A diagnosis is not simply what the patient has. A diagnosis is what you think the patient has. very different. Yes, the patient may bring us a symptoms but we have to do the thinking and this brings me to this place that many of you remember and I know that younger graduates probably do not remember this place Fitkin Amphitheater. Remember Fitkin? Every Friday afternoon at 5:00 we would all gather at Fitkin Amphitheater. How many of you remember that? For the clinical pathology[20:34] conference, a case would be presented, the history, the physical examination, the laboratory findings, the imaging. Then the discussant, an expert clinician would reason through the case, constructing a differential diagnosis, weighing one possibility against another and searching for a coherent explanation of all the findings. Eventually though, the clinician had to commit to a single diagnosis. This is what quote this is what I think the patient has. That was the discussant's final diagnosis. Not really not reality itself, mind you, but a representation of reality constructed with that within that particular clinician's well-trained mind. And then came the moment of truth, the pathologist. The slides were appeared on[21:25] the screen. We saw the tissues. We saw the cells. We saw the lesions. And there it was, the answer, the truth, the final diagnosis. I love those conferences. They were enormously compelling and educational. But looking back, I think they also taught us something subtle about the nature of medical knowledge. They encouraged us to believe that the diagnosis existed somewhere outside of the physician's mind. that if only we looked deeply enough and

measured precisely enough, we would arrive at the objective truth. And of course there is an objective biological reality. The pathologist did not invent the cells. The tumor was there. The microorganism was there. The inflammation was there.

[22:12] But at the point of care, we don't encounter this that biological reality directly, do we? We encourage it. We encounter its clues. And from those clues, we construct the diagnosis in our minds. The patient has the disease, but we hold the diagnosis in our mind. A diagnosis is our mental representation of what we believe is happening inside another human being. Sometimes we're right and sometimes we're wrong. And most often we must diagnose and act before the pathologist ever arrives with a final slide, don't we? Artificial intelligence is a profound shift and not just in medicine. For more than two centuries, the extraordinary success of physics and engineering[22:57] encouraged us to think of knowledge primarily as the discovery and measuring and measurement of objective facts. The figure seen on the left side of the image. Break reality in small into smaller pieces. Measure them accurately. Discover the rules connecting them. Eventually you will understand the whole. That way of thinking gave us the scientific method, the industrial revolution and of and most of the miracles of modern medicine. And it also gave us this steam engine and it gave us the computer which many of us still think of as a machine the way many scientists think of viruses as chemical products rather than as living things. But neither viruses

[23:41] nor computers are fit neatly into those simple categories. They may be both or neither. A diagnosis is a fact. A radiographic finding is a fact. A symptom is a fact. But what those facts mean for our patient at this moment here and now is a matter of clinical judgment, of empathy, of feeling. In other words, the figures on the right side of the screen. And we are used to practicing medicine in the fascinating space between uncertainty, between what we presently understand to be happening and what is objectively happening. That's the true art of medicine, isn't it? The facts remain objective, but their clinical meaning takes place in our mind where they are organized into understanding of our[24:29] patient. Or put it simple, facts exist in the world, but our understanding of those facts exist solely in our minds. And until now, computers could only store and process facts. Now, the age of artificial intelligence is teaching machines to work with meaning, to work with our minds, the images we see on the right side of the page. The steam engine does not care about what you think. Artificial intelligence does. This is a profound change whose impact we have only begun to appreciate. AI can already work with language, context, and patterns of thought. Artificial intelligence allows us to ask a far more meaningful question. What do I need to understand about my patient right here and now to make the clinical decision?

[25:19] And here we encounter a remarkable paradox. Artificial intelligence. Sorry about that. Artificial intelligence. Uh oh, I need help. Let me see. Maybe I can do it here. F5. No, I think it is correct? Oh, okay. It's correct. I apologize. And here we encounter a remarkable paradox. Artificial intelligence, one of the most advanced achievements in engineering, may help restore the humanistic side of medicine. And you know, it's really funny. Many of you remember this too. I was fascinated that Yale always picked students who had whose undergraduate majors had been in the humanities, music that were musicians in our class and artists and and philosophers and and I thought, boy, that means we have a lot of idealism. Yeah. But it[26:16] means something else too. And now now it becomes very clear can't practice medicine without that. We're not just scientists, we're humanists. And I don't know any other scientific endeavor that is like that. We're not one or the other. We're both. And think about this. What I'm about to say now you many of you may know who this gentleman is. CP Snow famous British philosopher. In 1959 he famously described the two cultures the growing separation between the sciences and the humanities. He was fascinated by the fact that scientists couldn't understand humanists and humanists could not understand scientists. I sometimes joke that a scientist cannot understand Emanuel Kant because if he understood Emanuel Kant he couldn't be a[27:02] true scientist but medicine has never belonged exclusively to either culture. He's all and doctors always belong to both. Interesting, isn't it? Artificial intelligence may help you reunite these two intellectual worlds that have been drifting apart for

generations. Medicine needs the scientists to ask what are the facts and needs the humanist who ask what do these facts mean for my patient for this particular human being. And perhaps artificial intelligence can begin to bridge those two cultures because it is forcing science and technologies to confront questions that once were left largely to philosophy and the humanities. What is meaning? What is understanding? What is judgment? What is authorship? What is suffering? What[27:53] makes us human? And ultimately, what does it mean to understand another human being? These are not peripheral questions in medicine. They're the very heart of medicine. And that brings us

Why Ambient Scribes Aren't Enough

[28:07] back to ambient AI. Let me begin by saying that at first glance, the ambient AI scribes appear to offer exactly what physicians have been waiting for. They promise to save us time and to free us from one of the great absurdities of the conventional electronic medical record, turning away from our patient in order to feed the computer. But most of these systems are built around narrow conception of clinical documentation. capture the conversation and generate a note. In other words, they approach a clinical documentation as a mere transcription problem. But as we have seen, it is just a transcript of sounds in the room, not really a clinical history. Because the medical record[28:51] ultimately takes form in the physician's mind. It contains what we noticed, what concerned us, what we doubted, what we concluded, and what we decided to do. So the limitation lies not in artificial intelligence but in the narrow task we have assigned it to accomplish. If we design the AI merely as a scribe, it will listen only for words. And unlike a human scribe, it does not know how we think or at least not as well as our trusted assistant does. But if we design it as a medical assistant to the physician, you must also learn to recognize what that physician understands and how this particular this particular physician thinks, how you personally think. In that sense, AI can learn to read the unique mind of[29:42] each physician. In the last three years, I counted more than a hundred medical articles describing the problem of authorship. And one of the caveats in relying solely on simple describing and confusing it with clinical history. Several of these papers appeared in the last two months alone and in fact several have been generated in this very institution by amazing authors and they are physicians. Uh they're brilliant, brilliant authors. Uh, in fact, we have the honor to have Jaideep Talwalkar here. Uh, Donald Wright, I think he's on Zoom, I believe. Uh, Ted Mel Melnik had the pleasure of meeting yesterday. Uh, Lisa Zang, uh, Dr. Ritchie Verma, and there are many

[30:36] others. And in fact, I apologize if I don't mention all of the ones, all of the papers that came out of here because I did my my best to try to find them, but there's just too many. And then I learned from Dr. Jacob that in fact his paper that came out two months ago was was sent out two years ago. So I can only imagine there are many more Yale papers and if I so together these studies document this extraordinary promise but they also reveal the persistence of the central challenge. A transcript tells the machine what was said. It does not tell the machine what mattered. A patient may spend 10 minutes discussing something we considered incidental and[31:23] 30 seconds mentioning the symptom that challenges that changes our entire differential diagnosis. The transcript doesn't know that we do. A patient may tell us that the pain is a little better. We may write that the patient is deteriorating. That's not necessarily a contradiction. Perhaps the pain improved but the weakness progressed. Perhaps examination changed. Perhaps as Dr. Sanders has explained so well something in the patients expression or reaction concerns us or perhaps we have seen this pattern before and suddenly something clicks. The real challenge for medical AI is not simply can the machine listen to our patient. The crucial question is can the machine[32:07] understand what we understand. If it can begin to do that we arrive at a very different future for artificial intelligence in medicine. Instead of creating an artificial doctor who to replace us, we can create an extraordinary capable assistant that frees us to be better physicians. And that brings us back from philosophy to the daily reality of practicing medicine. Let's consider another experience that has become both routine and rout and routinely infuriating. We examine our patient and determine that the

an MRI is medically necessary. And then as all of all of us know painfully well, an entirely different process begins. The insurer wants documentation. Has the patient completed six weeks of conservative[32:57] treatment? Which medications were tried and failed? What was physical therapies completed? Was prior images performed? Does our patient satisfy this payer's particular criteria? It infuriates me just to say it. Please submit the records request is denied. Well, we all remember the horrible thing happened in New York only eight months ago. Some p some doctors cheered and I feel very bad that they cheered, but I understand very painful for us doctors when we're fighting for our patient. We have to deal with that garbage. Uh, and now the process escalates. Please submit the records and somebody has to do that. Please, more information is required. Somebody[33:42] has to do that. The request is denied. Now you have to explain our clinical reasoning. Now, please appeal. Send it all back. Who has to do all that? Well, unfortunately, we do. Or worse, somebody in our staff has to do that. All that has to be paid for. But why should we devote our precious time or the costly time of our billers, practice administrators, and other highly trained members of our staff to work that requires information processing rather than clinical judgment. That work is not free. Every hour devoted to satisfying administrative requirements consumes resources that could support patient care, compensate caregivers, or make care more affordable. it reduces the[34:27] income of those providing the care and increases the cost of the care for patients and here can bring a very simple idea let's say you own a roulette table in Vegas at worst you go to the bank and you want some money and you say I need so much money for this afternoon games and if at worst the bank will send a technician to calibrate the table because if that table is working properly you have no risk. We of course are in the opposite extreme. And that's why insurance companies exist, right? If we knew those risks, they wouldn't exist, would they? Who would exist? Hey, we would. People who give care. Let's think about that. Could artificial intelligence do that? I believe it can.

Let Machines Handle Administrative Burdens

[35:15] Um so so we allocated the resource backwards. Physicians care resources clinical judgment. The machine's abundant resources information processing. Physicians should decide. The machine should assemble, retrieve, compare and[35:34] explain. So let's repeat this with AI. Um one more time. Let's say we believe a patient needs an MRI. Then let the artificial intelligence engine find the facts. Gosh. Okay. Sorry about that. As you can tell, I don't do this very often. [snorts] Let let the let's just decide a patient needs an MRI. Then artificial intelligence let the artificial intelligence engine find the relevant facts, determine the insurance requirements, assemble the evidence, uh the authorization, everything that I just said before, appeal it. And if they if another computer because there are computers that are are are representing the insurers, right? If they deny it, ah, let machine argue with[36:41] machines. Let them fight with each other. We don't have to be involved. AI can do that. Right. But it has to be, you know, programmed, right? And and by the way, who has to do that? Yes, we Dr. Ho. But you guys have to, okay? And we're going to talk about that. Um, so why are we okay? So the same principle applies throughout medicine. Let AI prepare the coding. Let AI create a concise medical summary that tells me at a glance what is happening with my patient. Let AI draft meaningful letters for our consultants, referring physicians and patients informed by our knowledge and express in our individual clinical voice. Let AI read incoming documents, extract important findings, and transform them into searchable data,[37:36] trends, and flowcharts. Let AI explain our instructions and afterwards in a clear language and the patients own language. Uh, AI can do all that. Yes. After decades, I don't know if you realize, but all the computers are looking backwards. After decades of physicians learning to accommodate to the machine, perhaps it is finally high time for the machine to learn from us. If we do it correctly, the result will not be less human medicine. It will be far more human medicine. The physician who no longer spends time finishing the

notes can think. The physician who no longer uh has to spend 20 minutes searching through a chart can listen. The physician who no longer has to compose a[38:46] third justification for insurance can talk to the patient. the physician who's not staring at a template can look across the desk. That is what I think that's why I think that the I believe that the debate about whether artificial intelligence would destroy the art of medicine has the future almost exactly backwards.

AI May Restore the Art of Medicine

[39:11] Artificial intelligence may be our best opportunity to restore it. And that brings me back to Yale and to Fitkin once again. more than half a century ago gave a group of young medical students something extraordinarily valuable. Freedom. Freedom uh freedom to learn, freedom to question, freedom to think. Lewis Thomas told us not to be afraid of asking stupid questions. So perhaps the question we should be asking is not will artificial intelligence replace us. Perhaps the question is what has prevented doctors from being doctors and can artificial intelligence finally set us free? I believe it can. The industrial age taught us how to measure the patient. The first information age taught us how to collect enormous[39:58] amounts of data about the patient. Perhaps this new age of artificial intelligence will teach our machines to understand what that what that information means to those of us caring for the patient. For decades, we were forced to adapt to machines. For the first time, machines can adapt to us. If we get this new relationship right, artificial intelligence will not mark the end of the art of medicine. It may mark its return. But that leaves us with two this major responsibility. We physicians cannot stand back once again and allow the next generation of medical technology to be designed without us. We allowed that to happen with the electronic medical record. and we have lived with the consequences. This time we must be[40:43] deeply involved. And by the way, unlike EMRs that you need to know coding and all that stuff, in artificial intelligence, all you need to know is how to read and write and be smart. And I have a feeling over there there's a lot of smart people here. So that's not an excuse. All right. Uh by the way, doctors know more than programmers or developers of software. And I include myself in that. I learn from our clients every day. I'm a student of our clients because there's no way even for somebody who's practiced medicine to understand what happens out there unless you talk to the person who has to suffer through it. Um we must speak clearly about what we need technology to do for us and for our[41:29] patients and what we must never allow it to do in our place. The opportunity before us is fundamentally different. Earlier computers required us to communicate in their language. Artificial intelligence can work with human language, context, and meaning, which is our language. AI can learn from us if we're willing to teach it. So, let us teach it not merely how to document medicine, code medicine, or justify medicine, but how to support those of us who practice medicine. If we physicians help shape this technology, AI need not pull us farther away from our patients. It may finally help us bring us back. Thank you. Any questions? Remember what Dr. Thomas said, no stupid[42:21] question.

Go ahead.

Questions and Answers

[42:21 Jaideep S. Talwalkar, MD, Professor of Internal Medicine and Pediatrics, Associate Dean for Educational Technology and Innovation] Thank you, Dr. Low, for your uh comments and your emphasis on humanism. uh a lot of thought provoking things that you said, but one of them is challenging something that I've been thinking is just a fundamental part of what we need to teach people right now. And I'm wondering if this how you feel about this. So based on the EHR experience and how we didn't express our voices enough to shape the way those tools look, I'd really been arguing that physicians need to be on the front line of understanding new technologies, how they work, so that we have a seat at the

table. uh but you're arguing for systems that are so intuitive, so embedded in what we do[43:01] that we just need to focus on being doctors, what we already need to know. So I wonder how you feel about how much a trainee today or someone starting their career today needs to actually know about computer science and how AI works versus just focusing on the traditional aspects of medicine.

[Dr. Low] Thank you so much for that question, Jaydeep You know, this new world of AI does not really require no encoding or or computers. In fact, that's almost irrelevant. You don't need to know how a pen works or how paper works or how the keyboard works. But it does require creative and intelligent people. And by the way, I think that the humanism that we were all taught had to do with creativity to think out of the box. We have to know what we don't know, which is what we were trained as doctors. We're always knowing what we don't know because if we knew that there would be no need for us. I don't really see patients going to the computer say I'm I I have I think I have cancer. Let me have the computer fix me. No, they're going to come to you. All right? And so what you're doing at school is incredible because you're teaching medical students to start getting involved in computer develop in software development and this is crucial. All right. Uh uh this is absolutely crucial. Uh I don't think they need to know much about computers and how they work inside. It's [44:39] good to know a little bit, but they need to know how medicine works and be very forthright in expressing that to the powers that be. And believe me, the powers that be developed software, they think they know software. I got a call from India the other day and there was somebody who had developed this software that and they wanted to interface and we're going to talk about interfaces later. I'm not going to get into that right now, but they wanted to interface with our software, which I think is absurd. You can interface and we can talk. I could talk half an hour about why interfacing doesn't work, but basically it was a bot and a patient would call on a Thursday and say, "I have this problem." Would take a clinical history right on the [45:22] telephone and would generate the finding. I mean it would be ambient AI to the next stage because he had the right smart questions and stuff and it would when the patient showed up on Monday would schedule the patient the clinical history would be all set up for the doctor saving a lot of time. I said what if the patient dies over the weekend? Who's responsible? Well, if the system if if we if this AI system recognizes that's an emergency tells the patient to go to the emergency room. So, wait a minute. Are you a doctor? No, I'm not. Well, I was an emergency physician and I don't know what an emergency is and you're telling me that the bot knows. So, we can't let them do this. All right. You have to develop the EMR and [46:06] the EMR has to be developed by having a just like you talk to your ChatGPT and you have this uh dialogue this uh communication is colloquy with a with a with ChatGPT you can have it with your vendors and if they don't listen to you it's time to change vendors and the same thing goes for you can't let administrators tell you how to develop an EMR Sure. Yes, they need it. They need the data to get paid and that's very important. But AI can do that. Of course, that's a major problem because many people who are on the middleman between the patients and you will no longer have a job that that I can believe. All right. As I said, if you own a roulette table, you don't need an insurance company. But that's but the fact that the people who give value add[46:59] it the nurses and doctors not only will they keep their job but I think your income will go up and not down. Anyways, I don't know if I answered your question or if I missed it. Thank you and thank you very much for the work you're doing. I admire the fact that you're a physician and that you know physicians are running the show. This has to happen. Go ahead. Any other question?

[Ann V. Arthur, MD, Director of Alumni Engagement (chat question)] Yes, there's a question in the chat [clears throat] is let me see that if we train AI how we think and how we get to diagnosis what it would it not replace us depending on who owns the AI.

Physician Ownership of Intellectual Property

[47:57 Dr. Low] An AI like a true EMR should be a tabular rasa. It's a process and not content. The content belongs to you is your intellectual property. Uh not two doctors have the same intellectual content. Two

doctors who went to the same medical school, specialized in the same specialty, are in the same town, in the same office, in the same patients, will practice differently. As you all know, it's an art. Well, that's we call it the art of medicine. And the AI learns how you think. And if you've played with ChatGPT, after a while, you begin to realize that it knows you. Okay? It can do that. And I don't think that's negative. I don't see it as a negative thing at all. In the 1950s, in the early 1950s, we're really [48:47] scared of atomic power. There was nothing positive back then about atomic power. There was just an atomic bomb that killed a lot of people. We were really scared and we're still scared today and rightly so. But at it's like anything else. There's a good thing and a bad thing. In this case, the good outweighs the bad. So, the intellectual property belongs to you, you know, uh not to the vendor or to anyone else. Now the vendor may develop the mechanism by which it works but AI belongs to you. Go ahead.

[49:20 Doctor Korevaar] So um you just hit on something that has really troubled me about the interface between AI and medicine. I've noticed that there are people, not physicians, but regular patient kind of people without much medical knowledge who have gotten involved with Chat GPT, who have allowed ChatGpt to learn everything medically possible about them, and who have been encouraged by ChatGpt's permissiveness, inability to say no, inability to stop offering information and answers into a kind of Munchhausen's thinking where these patients actually begin to believe that they have something extraordinarily rare and only they know what it is. And I'm a little bit concerned that in your system in which the AI is getting to know how the doctor thinks, you may end up with a similar kind of result and I'm curious to hear what you think about that.

[50:48 Dr. Low] What what a great question and and I will it reminds me something that happened to here at Yale when I was a medical student that many of you have may have experienced as well. If you recall, the Yale system had doctors exclusively for the medical students. I don't know if they still do today, but you didn't go to the regular doctors. You went to a doctor that only is still to today. Okay. So, I remember going to my doc, I don't remember what the issue was, and he and he says, "I know you believe that you have skin cancer, do don't you?" Say, "No, cuz you're studying it right now. Right? Do you remember that? Okay. So, I don't know if that's the answer, but that that you're wishing. Yes. The patients begin to believe that they have all these crazy things because they're reading about it. Touche. It happened to us as well. We're human. [laughter] So, absolutely, it's not a fear. Knowledge is never knowledge should never cause us fear. Uh we are dealing with the unknown and that obviously causes a little misgivings. You know the you know when when the electricity came out a lot of people who were building candles lost their job you know the what's going to happen you remember the ludites the ludite attitude is understandable you know you work in a certain way you practice medicine a certain way for many many many years is very difficult to change but change we must you know and the beauty as I said what I [52:16] loved about the yellow education was not that he taught us anatomy you know pharmarmac pharmacology. I I was very bad in those fields by the way. Uh but that he taught us how to think out of the box to be creative and they gave us the freedom to do that because it's constantly changing and it's going to change largely. Go ahead. You you I didn't answer your question.

[52:35 Doctor Korevaar] Well, just to follow up, isn't one of the things that we as physicians bring to the table what we used to call uh when I was a medical student the do nothing reflex? the ability to say wait no more questions no more answers let's wait till tomorrow and does your vision of this AI interface include that ability absolutely

[53:06 Doctor Low] Absolutely uh you know uh. Let's say a patient comes in with 10 problems you know the software will bring you the 10 problems but they come in with an urgency and you have a packed clinic you can't take care of the 10 problems, you're going to take care of the emergency and tell them,

"Come back tomorrow. We'll discuss everything else." You know, time is decided by the doctor. Medicine is a human skill. I know it's maybe not as obvious to us as it is for our patients, but I doubt very much that a patient when they're getting into trouble, they're going to go talk to a machine. Yeah, they may talk to the machine, but eventually they want that human. And that's what we were trained to do. And this is what we love to do and that's what we're good at. And that that is not going to be taken away [53:50] from us. It's going to be empowered. I have a colleague I as you know I live in Argentina but Argentina I don't know how many have you how many of you have been there? Oh many. As you know it's a very European country but not completely. In the northwest of Argentina there is a little area called Jujuy which is abuts Bolivia and it's indigenous. It's one of the few indigenous areas 50% indigenous and descending from the Incas and we found brilliant software talent. In fact, eight months ago, the Yale the Yale uh alumni service course sent 41 students up to Humahuaca. It was an amazing experience and they want to come back. And by the way, if you want to join the Yale service corp, you may want to join

[54:31] them. Uh we find that these areas intelligence doesn't stop at the border. It doesn't stop even within regions and people who come from areas that are incredibly underdeveloped whose parents were completely illiterate can do artificial intelligence. So with major reasons you can I mean without any problem at all.

[54:31 Christine Walsh MD'73] Dr. Low, we have a question from Dr. Christine Walsh on Zoom. Okay, Dr. Walsh, go ahead!

Thank you so much Dr. Low for a very thought provoking discussion. Uh, so if doctors are doing all this training, should we not be reimbursed for this and who owns all this training that we're doing?

[55:21 Doctor Low] It's a great question. Well, you know, the word training I think is also part of the previous era. We have to educate ourselves, but we always do that anyways. uh this is going to be a very natural education because what you have noticed I'm sure if you play with ChatGPT is that it speaks your language training means that you have to adapt to something else to another language a good by the way I always say that a perfect software would not require help online there's no such thing not with ours or any other okay and that's a challenge for developers how do you make something so intuitive that it requires no training that's a dream. All right. But [56:06] the fact is that I think AI will cross that barrier. We'll know what you know and will know what you don't know. We're developing another company which is called Zelia from the Latin Auxilium help. And basically, it's a help online and it's going to be used not only in medicine but everywhere else. We think that you can ask questions in normal language in English and it will look for the answer in the database in the knowledge database and will give you three answers like how many of you have ever tried and help online and it didn't help you uh all the time and then the worst part did it help and you say no and it goes away right well that's stupid okay that that's that's a [56:49] previous era what we're developing is one that when you say no it immediately segments you because he knows who you are, right? And it it creates a support ticket and it goes to a human and now the human sees the questions you asked, answers your question but feeds the AI. So if somebody else has the same question in the future, it gives them the previous answer. Did it help you? [57:13] No, it comes back again. But after a while doesn't come back. And how are we planning to sell it? Very simple. We're going to give it away. As long as you add personnel, you don't have to pay. If you reduce personnel, you're going to have to pay a little bit. Okay? Because it's going to improve support and reduce the cost of support. Totally different company. Has nothing to do with the EMR. Go ahead.

[57:39 Richard Kayne MD'76, Past President, Association of Yale Alumni In Medicine] So, as a student, I rounded with Dr. Spyro. Oh, Howard Spyro. Wow. Kept being readmitted to the ICU with serious upper chi bleeds. And he asked that we all wait outside and that he would go in and question her. And he came out after about 10 minutes. He was [58:12] smiling and said, "I have the solution. I think she needs a dishwasher, a washing machine as well." And we looked at him how that was even relevant. He said,

"Well, she's got 12 children. She has all these brothers, sisters, this huge family. She's extremely stressed out. And it just agitates her so much that she ends up coming back." I'm just wondering if AI is going to have the capacity to uh inquire that way. uh my son writes uh uh comedy and he's finding that uh among uh the different AI servers they have no ability to write comedy that would pass muster. I also find with open evidence I question the same question on two different days and I get two different answers. They may only be off my absolutely right. Well, I I didn't just before I get is there the possibility I remember I had a childhood friend who was a Harvard medical student. They had a camel that was the curriculum the syllabus and all knowledge was in there and he would he would give me a copy of that. I'm just wondering if there's a treasury that you will build of trusted sources where we didn't feel safe.

[59:43 Doctor Low] You know, the great Howard Spiro, as you know, he wrote his own book that was very rare even back then an amazing guy, and that brings me to something I tell our clients suppose a patient comes in and they tell you doctor I'm dying my stomach is killing me. Oh yeah how long has it hurt you? Oh god, about 6 weeks I'm dying doctor is horrible. Where does it hurt? They point to the midepigastric region. Well, you made a diagnosis. Peptic ulcer, duodenal ulcers, reflux. It could only be a number of things, but you listen. Yeah, doctor's being horrible. My wife left me and she took the kids, doc. She won't let me see them. I'm losing my job. I can't sleep at night. My body hurts. My back hurts. And it keeps going on and on and on. We've all experienced that. But you listen patiently, empathically, and then you examine the patient. There's a little tenderness in the midepigastric region. You give them maybe a proton pump inhibitor. Tell them if you don't feel better, come back. They **[1:00:18]** don't come back. But four months later, they go see another doctor because their back pain get worse. They do an MRI and they find metastatic cancer from the prostate going up to the upper back. He tells the new doctor and eventually his lawyer, I told my previous doctor about the back pain. That's you. And it's the AI captured it and put it there because errors of omission are far more frequent than errors of commission. Now if you had written patient complaints of back pain, you would have addressed the problem. And if it didn't make sense to you, it wouldn't be there to haunt you. The lawyers call it Res Ipsa Loquitur. That's Latin for the thing speaks for itself.**[1:00:52]** That's when you find the scissors inside the belly of the patient. There's nothing to discuss. And by the way, these companies have to keep the the sound even though they say they won't because doctors are going to say, "Oh, patient never said that." Oh, yeah. You're gonna hear the the sound. And and by the way, the lawyers are the first one to find out cuz you sign a consent. So they give it to the lawyer. Oh, we don't want to see the EMR. We want to see the sound. And so there's a lot of issues there. But the point is that an EMR, well, I don't want to discuss what our software does, but I will. I have to in this case it it learns from each individual physician and then it sends to the AI the ambient we call it **[1:01:34]** reflective ambient intelligence it sense the ambient noise but it also sense your knowledge and it asks the AI privilege my knowledge my text my thinking my words over what you hear and if anything else is that was said by the patient spontaneously or I asked is not part of my knowledge then put in a different color put that back pain in another color and then I decide if I want it or not and if I want it then I enrich the system and if I don't want it it goes away all right so the fact is the authorship matters medicine comes from the mind not from the environment yes you remember the CPC the physician he was wonderful the the presenter wonderful an expert but eventually he **[1:02:20]** had to commit and there was no pathologist at that first stage and we never get the pathologist ologist. So, it's always a productive. The diagnosis is inside the mind. It's an art. I'm sorry. I don't know if I answer your question, Richard. Thank you. Richard Kayne is a classmate.

[1:02:42 Inaudible question] Yeah. Go ahead.

[Doctor Low] Thank you very much. I think Maybe I'm a true just like hey exactly right think about this because okay well I I'll give you another one. Yes. Thank you. in our software and I don't want to talk about our software in addition to what I just said about the reflective ambient intelligence it also captures secretly questions that other doctors would have asked relate because it gets the entire clinical history

sees[1:03:31] it that I forgot to ask and I don't see it unless I want to see it if I want to see it I press a button and it puts the other symptoms in order of frequency of hallucination and yeah the AI hallucinates and so do our patients by the way our patients hallucinate all the time. It's the art of medicine is differentiating the pearls from the hallucinations. So then you pick it and if you think oh my god and then you move your cursor and there's a tool tip telling you why it thinks that that symptom is important and then you decide anyway that I think that's that's the end of me. Thank you. Thank you so much. That was wonderful. I kissed her on the cheek because I I remember I was working in Northern California and and I get called by the[1:04:18] administrator because I abused a nurse and I was shocked and I said, "But I abused the nurse." He said, "Yeah, you kissed her." And then I realized that where I come from, everybody kisses all the time. So to me, it wasn't even he didn't have any other innuendo to that. But boy, I was in the wrong place when in Rome. Thank you.

[1:04:46 Doctor Korevaar] Thank you so much, Dr. Low. That was wonderful. And thank you to everyone in the room and on Zoom for participating and asking great questions.