TECHONOMYHEALTH

Artificially Intelligent Healthcare

Speakers:

Walter De Brouwer, CEO, doc.ai

Ron Gutman, Founder and CEO, HealthTap

Maria Luisa Pineda, CEO and Cofounder, Envisagenics

John Mattison, Chief Medical Information Officer, Kaiser Permanente

Bud Mishra, Professor, New York University

Moderator:

David Ewing Duncan, Arc Fusion

(Transcription by RA Fisher Ink)

David Duncan: Everybody here represents some aspect of AI, of healthcare, entrepreneurism, you know, lots of different points of view. Starting with Walter, just say who you are, and I'd like you to answer, in a fairly short, concise way, your greatest hopes and greatest fears about AI and health.

Walter De Brouwer: Walter De Brouwer, I'm a computational linguist. My greatest hope for AI is that it comes very soon and I hope my greatest fear doesn't come soon enough, because we are actually losing neurons, sextillions of neurons by the second because of, you know, the fertility rate going down and life expectancy going up. And I think we now have this great technology that we don't have to use organic substrates but can put actually intelligence straight into silicon and that intelligence has to go somewhere.

So my company is a deep language company. We are a business-to-business company.

Duncan: Walter goes way back with the Internet and coming out of IT, and he's actually a linguist by training and started the company Scanadu with his wife Sam, who's here, and they have this new venture, which I want to hear a little bit more about. So, Ron?

Ron Gutman: I'm the founder and CEO of HealthTap, now a network of more than 108,000 physicians here in the United States, another 2700 healthcare professionals in New Zealand, and expanding this year all over the world. You know, we connect doctors with hundreds of millions of people all over the world to help them access better care on their terms. And you know, we're very excited about the opportunity to work not only with individuals in healthcare, but HealthTap has a consumer application that you can access by downloading the app or going to our website, or were actually working now with healthcare systems, with large insurance companies, providers, healthcare systems, and governments to help them manage the health of their population.

You asked what I'm most excited about in AI and what I'm concerned about. I'll start with what I'm most excited about. What I most excited about in AI is finally making healthcare smart and helping us personalize the right kind of programs to people so they can actually live healthier, happier, longer lives, and I'm happy to elaborate on that and what we're doing with Dr. AI and some of the other things that we are doing.

What I'm concerned about with AI is that we don't turn into the wild west. And you know, HealthTap builds its AI applications with knowledge that was collected over the past seven years from a network of more than 107,000 physicians. And also, everything that we are doing that is related to artificial intelligence, machine learning is always, always being scrutinized and being tested by thousands and thousands of physicians before we actually release anything to the market and we open what we're doing to them, we get their advice, we get their vote of confidence and only then release it to the market. So I'm concerned that anybody that comes to this doesn't take this kind of approach, because we are dealing with people's lives, right? So it is very important.

Duncan: No, in fact, that entrepreneurial issue, that's a huge issue. A lot of people come from IT and build these wonderful apps and devices and when they take them to doctors and patients, they go "We don't know what to do with this." So you know, there is a lot of just, you know, how do you position these. So that's our entrepreneurial voice, and Walter is a bit too. Maria?

Maria Luisa Pineda: I am the cofounder and CEO of Envisagenics. Envisagenics is one of the latest startups out of Cold Spring Harbor Laboratory. We have developed a platform for drug discovery using RNA sequencing and machine learning algorithms to find and discover brand-new targets and therapies for patients with diseases like cancer and other genetic disorders.

I'm most excited about with AI—which I guess I'm part of, because we use machine learning algorithms to just get better targets. It's the fact that we can now ask big, big questions because of the amount of data that's being generated, at least in sequencing data. And for patients, it's the only kind of emotional part is that we are able to deliver new therapies hopefully sooner and faster for them and their families.

What I'm scared for the AI movement is the limits that were going to be facing soon because of the computing and some of the technologies and the advances of them coming in.

Duncan: Okay.

John Mattison: I was a marine biologist and evolutionary biologist before I decided to go to medical school, practiced in the fee for service and academic world for a while and then joined a small startup called Kaiser Permanente. [LAUGHTER] And because of incentive model, because you get what you incentivize, and I think, back to David's question about what is our greatest hope and our greatest fear, my greatest hope is that we follow Einstein's advice that if he had an hour to solve a problem, he'd spend 55 minutes deciding what the right question is first and then five minutes solving the problem.

I think a lot of what we're seeing in AI today is not understanding what the right questions are, and I think the right questions were alluded to earlier by both Esther Dyson and Arianna and others in the morning around, we already know that in adverse childhood experiences, that our patterns are deeply set for the health span and our lifespan. And so how do we begin to focus some of our work more on children in a pan-generational and the transgenerational approach to solving some of the health problems. And I think we can do that, and I think we have enough knowledge today—so we're data-rich in wisdom-poor, as Arianna said.

We have enough data today to change how we practice medicine for the next 20 years but we need a lot more wisdom about how we do that, and where I believe AI can play a role is helping us to realize what works for what individual and link it with a very mature field of motivational science, which was also spoken about this morning, in ways that lies not just what we tell people, but the narrative and how we put together in a multisensory toolkit in what I call a motivational library of motivicons, which are most broader than what most people think of as an emoticon. So I think AI is really going to help us in personalized medicine to customize how we deliver messages in a motivational framework, especially the younger, the better.

What I worry about the most the current lack of sight into what's going on. If you haven't read *Weapons of Math Destruction* by Cathy O'Neil, I highly recommend it. It talks about the inherent bias in machine learning, because after all, humans do the supervised learning that machine learning is. And so my concern is what I called the "dyadarity." So we all know what the singularity is. The dyadarity is a term I coined to refer to having a transparency into what's going on in the black box and being able to ask specific questions about it. My fear is the dyadarity will not be instrumented soon enough, although there is promise in both the software and the hardware side. There are people working on transparency into the black box with some very creative, clever tools today.

Duncan: Fantastic. So, Bud?

Mishra: I'm a professor of mathematics at the Courant Institute of Mathematical Sciences. I also have a bunch of other affiliations with Cold Spring Harbor Lab, Mt. Sinai—I forget the other ones.

Duncan: He's one of these people, he has 17 titles.

Mishra: Yeah, I am a tenured professor. I sleep for eight hours. [LAUGHTER] I go to the gym for two hours. I follow one kind of hearty rule of never working for more than four hours. Well, I don't go watch cricket, but I could do that. [LAUGHTER]

So AI has been sort of part of my life. When I started my PhD at Carnegie Mellon, we had to take a qualifier on AI. I passed it with good grade without learning a thing, by using what I call weak methods, same thing that Siri and Alexa use. Didn't understand a word they said, but aced the exam.

But there is actually a—historically, there is an interesting question. In mathematics, Aristotle started what I call inductive logic. So you start with self-evident truth and derive other truths, theorems and things like that. But there is another side and inductive logic that never really took off. And the reasons are two. One is something called Goodman's paradox, or Grue paradox, because a lot of things are time-dependent. And now there's Simpson's paradox, that not all samples are the same. So if you take patients and do not stratify, you could get contradictory results. Or if you're looking over time, you could get results, causalities that are not correct.

So a particularly important is Ibn Sina, who wrote a book called "al-Qanun," "The Canon of Medicine." And he actually did not just things like clinical trials, case-control, experimental studies, but he also talked about philosophy of inductive logic. One reason was there was another, Al-Ghazali, a philosopher who wrote a book called *The Incoherence of Philosophers*, which is actually the core of Sufi philosophy—influencing ISIS, probably. But a big breakthrough came through another Andalusian Arab philosopher, Ibn Rushd, who wrote a book called *Tahafut Al-Tahafut*. And that influenced the Renaissance philosophers. The primary one is Pietro Pomponazzi, who wrote about why miracles don't work, 100 monks chanting for 10 days will not cause rain. And they burned his books in Venice. Another is Giordano Bruno. They burned him. Another is Niccolo Machiavelli. Galileo Galilei, who figured out the difference between chronological time and causal time. So there's a long history of Renaissance. But then there is David Hume, Adam Smith. And ultimately, there are people like Sufis, and that kind of causality analysis, with respect to all these paradoxes, all of these, need to be at the core of what we do. And it can't be, you know, just because you can tag some videos does not mean that you can actually understand the kinds of mathematical and deep questions that come up.

Cancer is a good example. It's highly temporal. It's highly heterogeneous. There are 50,000 different types. Single molecule makes a difference, each mutation makes a difference. And if you're coming up at the drug, you have to find the target, which patient is ideal, and not give it to off target. You have to deliver it to the right cell, the tumor cell, not the off target normal cells. You have to target the right genes, not the off target.

Duncan: Bud, I'm going to take that as your hopes and fears.

Mishra: Okay. [LAUGHTER]

Duncan: And Bud, by the way, I met him fairly recently and we spent, I don't know, several hours—it's amazing, I mean this whole stream of history. And as we talk about technology, as a historian especially, it is important all of this in context, because especially the philosophical side—and you know, Bud is intriguing. Everybody on the panel is intriguing. Bud is working on kind of these philosophical constructs along with the mathematics he does and moving that into oncology and some other fields.

So basically, we have an extraordinary group of people here. And as I said, we represent different facets. We have an Internet pioneer, we have an entrepreneur—and I've used Dr. AI. He demoed it for me, which is pretty fascinating, and boy, do I wish we had that—when I had patient interactions, I had the old, you know, fill out the form and what the hell is going on. You know, we have a biologist who is starting a brand-new company around molecular biology and getting into AI. We have a provider and a physician working in a very interesting healthcare system, very different from most in the country, a kind of closed system. And you guys have been collecting digital data on patients for years, right?

Mattison: Decades. Yeah, decades.

Duncan: Yeah. And designing some AI systems. And then Bud, who is more on the real AI side in terms of being a computer scientist and a designer.

So let me ask you guys, when you say artificial intelligence, what actually does that mean to you? And I'm going to preface that by saying that I think we all use this term and I don't think most of us really mean literally artificial intelligence. We're really talking about kind of advanced analytics, as I said earlier, all these different technologies. But what does it mean to you from your different perspectives?

Gutman: Sure. I mean, to simplify—

Duncan: Or is there a better term? That might be another—

Gutman: Yeah. The simplified is basically the power, the real power of machine learning or artificial intelligence is in identifying patterns in human knowledge that we can codify into systems that can create for us a better understanding of the world or better motivation for people to do certain things. It's basically taking human knowledge as a whole and finding the excellence and patterns that work for certain things and then putting them to work in the context where they're most valuable.

So I'll give you an example. What Dr. AI does is basically taking a knowledge base that we built over almost seven years of serving answers to people that ask health questions, a lot of them about symptoms, you know, I have this and this and this symptom, and sending it to the doctor to ask, "What do I have? What's going on with me? What should I do about it?" So the problem that people have is, I have some pain or I have discomfort and I want to figure out what's going on with me and what should I do about it, and they ask it to doctors. And we have more than 100,000 doctors answering these questions, you know, serving them more than five billion times to people all over the world and then taking this knowledge base is starting to look for patterns on when people ask on a certain combination of symptoms in this particular context, what did the doctor answer, what it might be, and what is the best course of action. So this is what AI is really very good at.

If you have a training set of the best knowledge in the world--not of one physician or a small number of physicians, but a huge number of physicians and many, many use cases of people asking it and doctors going back and forth and asking the right questions—because when you put a question or a symptom into a search engine, you know, the search engine cannot go back and forth with you like a doctor. If you put a headache in a search engine, you know, how many types of headaches can be there? About 300 different manifestations of headache. So having a doctor to ask the question what kind of a headache you have is very important. Putting it in context like a doctor would to understand your background, your health history, etcetera, and then start computing the most likely condition that you may have and the best course of action. So this is a really good use of AI.

Duncan: So you're talking about really, it's a huge connector.

Gutman: It's a pattern recognizer. It's finding the patterns of how doctors think about a particular issue and serving it in context to people.

Duncan: And that's basically, it's almost a machine-human interface too, because you've got real people—and I've seen this. I mean, part of it is automated, you know, certain questions, but then you do end up interacting with doctors, so you've got that human-machine—

Gutman: Yeah, because the doctors continue training the machine all the time. As people asked the question, the doctor answers it, they see what happens at the end of the day, and they keep refining the machine all the time. So there is basically a partnership between human and machine, or expert and machine, that make the machine smarter and smarter all the time.

Duncan: Yeah. Okay, who else wants to jump in? Maria?

Pineda: So at Envisagenics, the way we use AI is particularly, for instance, to do hypotheses-driven experimentation. So scientists throughout history have actually used scientific method to experimentally formulate their hypotheses and test them. Most of the time, it was done one at a time in the laboratory and so forth. What we use AI for at Envisagenics is actually using machine learning algorithms to allow scientists to actually be able to test many, or maybe thousands of hypotheses at once using high computing. And what this allows us to do is actually give a particular information to the scientists to make better decisions at the end of the day on what to follow on, what to spend resources on. So for instance, for drug discovery, right now it's around an average of 10 years and $2.1 billion dollars to develop one therapy. So how do we accelerate this process? If each of those is one hypothesis, if we could do this using AI and machine learning algorithms, we could accelerate this process and actually hypotheses-driven algorithms that allow the scientists—not to take away the scientists, actually just give them the tool to hypothesize faster in an easier way so they can actually choose what to follow and where to put resources and why they should add that investment and those resources so we could get drugs faster to the community.

Duncan: And you're dealing actually in your own work in the molecular level—and we have all these omics—you know, we heard about some of that earlier on the cancer panel and some others. I mean, you've got genomics to start with. Which actually was the first thing people looked at because it was the easiest, but we got proteomics, we've got the microbio—

Pineda: Or transcriptomics.

Duncan: Right. And I mean, just saying those words with all the syllables makes my head spin. But the amount of data being produced even now—and some of those are in the early phases. We aren't even close. And they're dynamic processes too. I mean, they're not static like the genome mostly is. So how do you see this—well, actually, just to follow up on my own question there, how does this play out in the developing—

Pineda: In the AI world?

Duncan: Yeah.

Pineda: So actually, as Ron was putting it, the more data we get, the more training sets algorithms have to train on so the better they become and the bigger the database is. So for instance, at Envisagenics, we focus on RNA therapeutics and RNA, which in the -omics world is called transcriptomics. So in the transcriptomic world, in order for us to actually use the RNA and all the sequencing that is being developed through all the sequencing efforts and all the patients is not only to actually use these data, but actually use other types of data as features that will allow us to do hypotheses-driven AI and actually formulate those to give us a better—for instance, I think Bud was referring to, with our platform, we are able not only to find new targets, but actually then model that target for an RNA therapeutic, and on top of modeling that little RNA that will become a drug that could be delivered to patients, we also can model where are the off target effects, what gene, how the protein is going to be done, what patients could it be delivered to, what age, what sex, all of those things that are extremely important in at least oncology moving forward personalized medicine.

Duncan: And RNA, by the way, is what reads DNA, and we only have 20,000 genes, which actually seems like just a few compared to what people thought. But RNA goes into different parts of the genes and its, I don't know, millions of—

Pineda: Yeah. So that's what my company does.

Duncan: Yes. So anyway, that's why it's so complicated.

Pineda: It's so complicated. So for patients, we could go from four teras, if we want to talk about AI, to 40 teras, because we don't only have the whole transcriptome, or the copy of your DNA, but we also have the metadata that goes along with that. Which means what type of patient, what tissue was the tumor extracted from, what stage was the patient in, all those things—you know, how is the protein being structured, what type of drugs are available right now. So all of that data is what you fit into your AI in order to actually create these hypotheses.

Duncan: By the way, these two people aren't—you know, what they're doing is not connected up yet, but that's just two elements. So John, tell us about the provider side.

Mattison: Yeah, I want to key off what Maria was describing, because there's been a couple of absolutely fascinating discoveries in the past couple of months. One comes back to RNA. Someone said earlier today that 95% of those of us living today will not have actionable information coming out of our genome, and that's because our genome is up and down regulated daily on a diurnal pattern, massive variations in our transcriptome throughout the course of the day. And oh, guess what? Our microbiome, those bacteria that live in our gut, have diurnal patterns as well, very substantial diurnal patterns, and the RNA that they leak across the bowel is what has a profound effect on regulating our genome expression. So our microbiome expression and its diurnal pattern has a huge impact on our genomic expression throughout the day. So the complexity of that, you can imagine is beyond the capacity of the human brain, so this is where machine learning becomes essential, and Maria's reference to generating thousands of hypotheses, I like to think of serendipity in silico. So serendipity is often when you recognize that something is really interesting out of those thousands of hypotheses.

So two studies that were recently, one is a study of cavernous venous malformations in the brain in a relation that is severely affected by these. It's a genetic basis, but what they serendipitously discovered is there is a single bacterium, a type of bacterium, gram-negative rod, that causes the malformation to progress, and then in the mice model, they've been able to do fecal transplants to eliminate that bacterium in those mice affected by the gene that causes the malformations and completely arrest further development of those malformations. So the fecal transplant actually is one of the earliest demonstrations—there are others—of a therapeutic intervention based upon the site.

The last thing I'll mention is, there was another study published recently about a very common genetic variant, which—about a third of us are homozygous, about a third of us are heterozygous and a third of us don't carry it, which is highly associated with accelerated cognitive impairment beginning in the sixties—in the age of the sixties, not the 1960s none of us remember. [LAUGHTER] And so what they found is that there is a viral infection that activates that genome and upregulates its activity in a way that leads to that cognitive impairment.

So if you start thinking about the interaction of the immune system, the microbiome, the genome, the transcriptome, and all of these factors, it's so far beyond the human capacity. So it is this hypothesis generation in silico, or serendipity in silico, as I like to call it, that must be recognized by what I like to refer to as a meta tropical brain forest. What I mean by that is a tropical rain forest is one of nature's greatest examples of how genetic diversity can take a very nutrient-poor environment and maximize the capture of solar energy and conversion into carbon-based life. The same is true for understanding what to do with thousands of hypotheses and to create serendipity in silico, and that is you must have a team of people who cross different disciplines working together to look at these thousands of hypotheses. If all you have is data scientists or all you have is people who are great at programming algorithms and you don't have people with clinical and operational knowledge in genomic and transcriptomic in the room trying to understand where the serendipity resides within those thousands of hypotheses generated, we're going to be stuck in a loop that's very costly and not very productive.

Mishra: So just to curb our enthusiasm, part of the problem is multiple hypothesis testing. So what that means is you can generate more and more hypotheses. Some of them serendipitously falsely can turn out to be true. So if you torture the data, they will confess. If you water board them, they'll confess more. [LAUGHTER]

So one big problem is how do you generate a hypothesis. Here is an example. If we look at the patient visit data, the patient who's visited a physician three times without a diagnosis is likely to be in a critical condition within a month or so. Why is that? What's the difference between a physician returning the patient with diagnosis and without diagnosis? About $22.63. He has an incentive to do that. So the real cause is not that somebody is—the pattern is predicting something. There is something else that's causally related to that hypothesis and we need to see that.

So the data by itself is not going to tell you that. There are plenty of examples of that. So there are challenges. There are challenges for AI. There's also challenges for biotechnology, biomedicine. There are tons of data that we do not see. We do not see translocations, gene fusions just on DNA. We don't see all the splice variance. We don't see all the isoforms. We can't count the small copy number transcriptions. There are real barriers. We don't know how to read epigenetic changes.

Duncan: What do we need to do to protect ourselves from those fears?

De Brouwer: Well, if you take any DNA framework and paid under a heat lamp and you have human DNA, eventually it will evolve structure. Once it has evolved structure, it will speak.

Once it will speak—sorry, once it will evolve structure, it will see. Once it will see, it will speak. When it will speak, it will think, and once it will think, it will predict. So that's intelligence. So that's why language and vision are the most important parts of our intelligence. Probably intelligence is vision raised to the power of language.

There are two schools in AI now. There's what I call the old school, and the old school is the school who does statistics, programs in R, and uses MATLAB. And they use discriminate means, like support vector machines, decision trees, things you can sort of explain. Then there is the new school, which has enormous success to connectionists, deep neural networks, unsupervised, very deep, with all sorts of dimensions that we can reduce. These schools in the last three years have 37% error correction. So they are now surpassing humans, basically. The only problem with that school is—and that school doesn't use, or hardly uses statistics. More like multivariable calculus and Python and, you know, and deep neural networks.

That school has one barrier against it, that in healthcare, doctors want to be able to understand why these black boxes tell us what to do, while actually, medicine is the biggest black box of all. It's completely qualitative, very, very small data, symptoms—you know, we are no longer in the symptom age. We are in the age of algorithms. So that's going to be the big thing, how can we convince doctors that—as in mathematics we have different fields we don't understand. We just bridge it over, you know? Ramanujan, we still don't understand how he came up with the infinity. But we have to do the same with these black boxes. And we now have great tools like stochastic gradient descent, where can sort of control it. But it's a leap of faith.

Duncan: Yeah. Ron, yes?

Gutman: Yeah, so I want to go back to—I mean, there's a lot of big things that are very, very important, fascinating and how we tie them all together to get a great understanding of the human and creates a big level of interest, but also complexity. But it's kind of like scary, right? So instead of that, like I'm thinking to myself, maybe because I'm an entrepreneur, I keep thinking about MVPs all the time, what's the minimally viable product that we can do today to solve a human problem. You talked about the tool; You talked about Einstein. Like talking about what question we're asking and what problem we're trying to solve today with the tools that we already have in hand to help people accomplish something in the next 30 days, 60 days, 90 days with fantastic tools that we already have today.

So let's think about this for just a second, right? What can we do with AI, right? So we used AI in the past 20 years in Silicon Valley to be able to do behavior change in commerce. Where taking people now with ads—we use ads, words, pictures, videos and making people that used to buy one product buy the other product, or take people that used to buy no product and buy a product or a service, right? So we are able to understand human behavior in commerce in a way that causes behavior change in commerce.

Hold on one second. We perfected that in the last 20 years in Silicon Valley to an art, right? We're doing it really well. Google does it really well, Amazon does it extremely well, Netflix is doing it amazingly well. Why not apply exactly the same framework to behavior change in healthcare? Rather than making you buy a ticket to a show or this laundry detergent, why don't we help you take your medication using exactly the same tools that we've perfected in the past 20 years using words, pictures, videos to help a patient take a medication that they don't take more than 50% of the time after prescription, use a medical device, get another doctor appointment or a test done and interpreted.

These are all small simple actions that they should take to get to a better outcome that we know will happen. So we have the toolset in our hand. These companies are using AI every single day, so we are hiring now these data scientists that are doing it in ads, and instead of serving ads to people—HealthTap could have made exactly $120 million dollars in the past two years in just serving ads on our pages. We are not doing any ads. But we are hiring the same machine learning people and data scientists to start using exactly the same tools to convince him to take a medication or use a medical device. This is a good use of AI.

Duncan: So I'm not getting anyone here to bite on where we're really going with this, like some of the darker things that might happen. [LAUGHTER] I mean, you know, there are people that talk about privacy—we talked about that today. We willingly give up our data on the one hand and we freaked out about our medical data. I think that was a very good comment, that it just kind of depends on if we get some use out of it or we think we can get use out of it. So there's the privacy issues, a lot of ethical issues. I mean, we're going to be turning over—you know, one of the nice things about AI is it processes its data, it connects it, it sees patterns, but it also hands it off in ways, and we are used to doing this with banking and some other things.

There's also issues about—you know, John, I might ask you about, you know, is your job going to be replaced by a robot equipped with AI at some point? I mean, you're already seeing like—

Mattison: But a lot of people—I'm not worried. I'm not worried. [LAUGHTER] A lot of people are inverting the AI and coming up with IA, intelligence augmentation. And I think that's much more realistic about what we're going to see. And the hope is that, rather than disintermediate everybody in what they currently do, is to augment what they do with intelligent augmentation of all sorts, not just machine learning, but robotics and the whole spectrum. And so we can upskill everybody.

The problem that I see—and I blogged about this yesterday—is that this intelligent augmentation may be disproportionately affected for those at higher levels of access to education and socioeconomic status and that we may see an acceleration of the social inequities, and that's part of what Cathy O'Neil talked about in "Weapons of Math Destruction" is there's inherent biases toward social inequity in the machine learning black box. So that's why I've been very interested in the dyadarity and how do we provide a level of transparency.

So there's a fellow in Silicon Valley with a startup—biomimicry, by the way, has been phenomenal in terms of helping us understand how to tune what machine learning does. Machine learning basically operates on a pyramid of algorithms, and there's a reason that our brain cortex is called the pyramidal cortex. It's very similar in how it works. The big difference to date is that almost everything in a traditional machine learning stack is feedforward. There's very little feedback, and the brain has lots of feedback loops. And so transparency may be collateral benefit of providing more of a feedback process within machine learning than exists in most models today, and there's people, as I say, working on that.

There's also people at Intel who are working on—not to call them out specifically, because other people are working on it, but Intel is looking at how to instrument what goes on inside the processor to provide more transparency into the black box.

So how we provide that transparency will not only help us accelerate how to learn using machine learning and large datasets, but will hopefully give us some sort of transparency into what's actually going on. The FDA is profoundly worried about this. The FDA is incredibly interested in reinventing themselves, recognizing this new reality of the black box world that we're entering and you put data in and you get an answer out and how do you know. And if overall everybody's better but occasionally you kill somebody with a really stupid mistake, what does that mean in terms of policy? What does that mean in terms of FDA approval? We need to have a whole different way of looking at the transparency and what goes on inside that black box.

Duncan: No, in fact, that's a whole other area which is fascinating, you know, how is this regulated. I mean, the FDA is having a hard time in a lot of other new technologies, keeping up with technologies from 15, 20 years ago, and this will be a big change.

But let me, in our last few minutes here, open it up to you all. Obviously this is another enormous topic we could spend a lot of time on, but you all hopefully have been pondering this. Any questions out there?

Audience: My name is John Meyer. I'm a pharmacist and a scientist, both working as a director of pharmacy here in New York City for many years and also for a pharmaceutical company for many years as well. I was especially intrigued by Dr. De Brouwer, and of course the whole panel, but talking about old school and new school. And I recently finished a book called "The Undoing Project" and I think you would all probably be familiar with the work of Kahneman and Tversky, and I'd like to see what you think about whether they're considered old school, new school or some sort of hybrid—if you're familiar, but I'm assuming you're familiar with them.

De Brouwer: It's a hybrid because, you know, a lot of people from the old school are my friends. [LAUGHTER] But you know, also the dominant university has been—you know, it used to be MIT. Now it's clearly Stanford for language and vision. So everything is changing, and especially with—so with Python coming into the mix and things becoming—you know, and TensorFlow as a result of AI.

But you know, coming back to that black box thing, so is a machine, where you put things in and things come out and you don't know exactly what's going in between, that should be—if that should be actually regulated by the FDA, then every doctor should be a 510(k). [LAUGHTER] Because a doctor is the ultimate black box and makes a lot more mistakes than our black boxes.

Duncan: Okay, let's grab another question.

Audience: You said before it was unclear how to get doctors to adopt the outcomes of black boxes, but it seems to me that data and outcomes would be very persuasive, so if you were able to show a doctor that the output from a black box is able to diagnose or lead to a better treatment than would otherwise occur, that that would persuade doctors. Do you think otherwise?

De Brouwer: So I don't know—you know, in the 1990s, we made word nets, you know, for computational linguists. Then image net came. Now last week Stanford open sourced medical image net, 0.5 petabytes of data, 1.2 billion radiology images, 4.4 narrative radiology reports, 23,000 conditions. There, radiologists are convinced because basically the machine does a better deal than a carbon-based unit. Although, it's a black box, because this vision—you know, Inception 3, which we use, is basically a deep neural network.

Gutman: But it doesn't need to be a black box. I mean, everyone is talking about the black box. I think that if we can actually leverage the fact that we have a lot of doctors out there—not an individual doctor or an expert, but actually take the physician community and help them open the black box to them, tell them how the method works, you know, what are the algorithms that we're using, and more than that, when the machine starts spitting back results—we've done it with Dr. AI.

What we did is we trained algorithms, we got the results, we broke it down into individual results, we start sending it to physicians one by one, so a set of symptoms, context, and here's the outcome, and let them vote yes or no. Not one physician, not a panel of ten physicians, but more than 10,000 physicians in 141 specialties, allow them to say for individual results, yes or no, wait until a statistically significant number of doctors answer either yes or no to either retrain the machine when it's a no—what the physician community thinks—or say yes, enough physicians, that each and every one individually trained think that at least what the machine is doing is right. And then you can retrain the machine again and again and again. It's an organic thing that actually works together to make the machine better all the time. When you open the black box to the physician to tell them how it's computed, now you have visibility and nobody's concerned about that.

Pineda: For us, for molecular biology and bioinformatics, our black box does need predictive analytics and it has to be experimentally tested, and actual percentage of accuracy is not based on the black box predicting, it's actually the experimental validation done by scientists of our predictive analytics, and then from there not only experimental validation, publish it, but also going into for instance using the Cancer Genome Atlas, the TCGA, and looking into 10,000 or more, almost 20,000 cancer patients and finding the same correlations that your predictive algorithm found and then texting it back—in mice and then in further patients. So unless we do that with our black boxes, then our scientists, at least for us, are then saying, okay, your predictive power I guess is true.

Duncan: Well, this is back to trust. I mean, we have a lot of black boxes we deal with in other industries we use all the time. Anyway, let me get one more question here.

Kofinas: Hi, thanks. Demetri Kofinas. Two things, one a sort of observation on something that Dr. De Brouwer said about the black box. I think it's true that each physician is a black box and that they're qualitative. However, a patient can ask for an explanation from a physician and derive some sort of answer in response. I think the concern with the algorithms is that we don't understand that process in the soup and there are so many variables that are coming in—

Gutman: You can derive the same thing from a machine.

De Brouwer: So we now have—you know, so Stanford has SQuAD, which is the big dataset with inference. It's pretty—you know, it's now possible actually with like 72% of accuracy that the machine unsupervised goes to look for things that it can't explain. You know, basically it's a bit like when I sit before a student and I ask him a question, he's first going to give me a definition from a textbook. I'm going to find out if he really understands it in your own words, and then I'm going to make him or her reason to see if these steps are there. The same thing we are now working on for machines, and I think the paradigm shift in language and reasoning will come in, you know, probably next year or the year afterwards.

Duncan: Yeah. Unfortunately, we're going to have to leave it—

Kofinas: I think just about the foundational mathematics point you were making and the paradoxes—

Duncan: Yeah, you guys, unfortunately, we're going to have to leave it here. And we can pick it up later. I was under strict instructions to end on time. [LAUGHTER]

Listen, these are all amazing topics and I think this fall at Techonomy in California, we'll talk more.

Mishra: Can I just have one second? So there are other things we do, clinical trials, cell lines, organoids, we're not going to eliminate those. That would be a very bad idea. We also have mechanistic explanations. We look at pathways, we look at gene regulation. We also do that kind of causal structures. I don't think that black box eliminates any of those. We're not going to shut down clinical trials. FDA is, they're part of—they need to integrate it together.

De Brouwer: Absolutely.

Duncan: Well, I want to thank our five black boxes here up on the panel.