

Ethical Artificial Intelligence: Oxymoron or Possibility?

Elliot Gerson: I'm just thinking, having been,

Elliot Gerson: this is our 15th Aspen Ideas Festival. I think this is perhaps the coldest day in the history of the Aspen ideas festival. So a, a special message to all of you who are watching this, uh, online. Then if you see the panelists shivering, it's not because of fear of the implications of artificial intelligence, um, but, uh, absolutely delighted that you're braving the weather to be with us today. And again, welcome to everyone watching, uh, the topic for this panel. Is ethical, artificial intelligence, Oxymoron, uh, or is it actually a possibility for a tremendous, uh, innovation improvement and actually improved equity in health care? Uh, essentially our questions today relate to the fact that the algorithms used in various forms of artificial intelligence are extraordinarily powerful, but as they inherently capture historical data that reflect exclusion, privilege by us and the like, uh, how can they be used without perpetuating the same, uh, will they, will they advance equity?

Elliot Gerson: How can they advance equity without exacerbating it? And particularly, I think, and there've been many conversations here over the last couple of days, what are the special challenges? At a time when there are so many people in the world, so many sectors that are questioning a fundamental trust in the scientific enterprise. To start out with our, to discuss these unrelated questions, we have a wonderful panel. Uh, you all have their biographies, so I'll be very brief and to allow questions also to all of you, uh, I hope not too long. A to my immediate left is Christopher Gibson, who's a co founder and CEO of recursion pharmaceuticals are clinical stage by a Pharma company leveraging the latest in automation, artificial intelligence, biology and chemistry, uh, to discover new metaphor, to discover new medicines and bring them to scale. Uh, then we have Mildred Solomon, president of the Hastings Center and a professor at Harvard medical school where she directs the fellowship and bio ethics.

Elliot Gerson: Millie of course, is a bioethicist and a social scientist to sought who studies these very same ethical implications of emerging technologies. And finally, of course, uh, Eric Topol is executive vice president of Scripps Research, a professor of molecular medicine and founder and director at the Scripps Research Translational Institute. Uh, he previously of course led the Cleveland Clinic and his new book is deep medicine, how artificial intelligence can make health care human again. So let me just start with sort of as this, that the notion in a title about there being an oxymoron here. Uh, you know, ethics as we all think of it is fundamentally an exploration of moral philosophy of a code, of what is good, what is bad, which means it's inherently a pursuit rooted in humanities. Uh, it's subjective, it's nuanced and it's, you know, fundamentally fundamentally human. So you know, the very words artificial intelligence in their own rights seem somewhat oxymoronic.

Elliot Gerson: Then you add the add the word ethical to artificial intelligence. And an almost seems like, you know, oxymoronic squared, yet artificial intelligence, uh, in its various forms offers absolutely extraordinary opportunities to make healthcare more inclusive, more equitable. And as we've been discussing in many of our discussions at ideas health, so how can this circle really be squared? How can machine learning, uh, guide, guided presumably by ethical and human principles make for more humane healthcare? So let me start, if I may, with, with Millie and in the middle, uh, what principles do you think should guide artificial intelligence, uh, w with respect to health care so it can best assure trustworthiness and, and reduce the kind of bias that I taught.

Mildred Solomon: Yeah. And as you said, Elliot, there, there is, there are many similarities across a lot of these emerging technologies. None of them are all good or all bad. It's in how they're used and deployed and how much risk we're willing to take as we develop them during the research phase and what purposes we put them to. So to answer your question, I think the overriding principle is human welfare. Now I might want to also say planetary welfare. If we usually talk about human welfare in moral philosophy. But increasingly I think we should be talking about planetary welfare too, but today we're probably going to stick to human welfare. So that's the number one umbrella for the conversation. And that means developing applications that are going to help all of us. So there's, you know, artificial intelligence is likely to reduce costs, um, increased profits. And for those people who develop successful applications and platforms, there's going to be a generation of enormous wealth.

Mildred Solomon: So we need to be sure that we're focusing on benefits to people, not just to profits. So I'd put that as number one and then I'll leave you with three key words and I'm happy to come back to them as we talk further human agency. We need to keep humans in the loop. Ai has transformed the air travel industry to make it one of the safest ways to travel. But we also know that we're kind of building with people out and that's created some of the problems that we've seen. Human Agency privacy. Slash, liberty. I'd like to come back and talk about that and fairness and that has multiple facets to it. So I put those three out first and then I can bring some in as others. Join the conversation. Yeah,

Elliot Gerson: please do. Jump in. And then again, we only have 50 minutes, but I want to turn to the audience, uh, as well. Uh, Eric [inaudible], your new book declares that artificial intelligence can make healthcare human again. So I mean, you really take this oxymoron on directly and I love that. So explain that intentional irony and why you are so confident. Right. Well, it is an unparalleled opportunity. Uh, that is because here technology, which usually is thought to be depersonalizing dehumanizing it can flip. And, and that's because so much of this technology not only can make the lives of clinicians, doctors, nurses, all types of clinicians, a much more efficient, more accurate, um, getting rid of the data clerk functions. But the biggest thing that it can do is to give the gift of time. That gift of time is what we've lost. And so that humanity, uh, the bond between a doctor and patient has suffered greatly over multiple decades and it's been steady erosion

when here we have a very unique chance if we do this right, to transcend just the, the early phase of getting this accuracy and speed and efficiency, um, and elimination of keyboards, that sort of thing.

Eric Topol: Getting more advanced, uh, efficient drug discovery like Chris is working on it. But the bigger picture, the, the ultimate goal would be to bring back, uh, the trust, the, there's precious, deep empathy relationship that we used to have. But I just, let me just follow up quickly with you. Uh, uh, you know, just the very phrase deep medicine, um, just of define that for us. And that sort of implies that there's something called shallow medicine too. Yeah. So, so what are they, well that's good cause I do have a chapter in the book called Shallow Medicine. Uh, that is referring to what we have now, which is very little time, uh, very little context and presence. So that we were just talking before we started that when, uh, when you go see a doctor, they're not even looking at you and that little bit of time you have, um, and it's error Leighton and a, we do 10 tests are done just because there's not enough time we made, we'll just order more tests. And so we have lots of waste in the system as well. So anyway, you looked at it, the shallow medicine that we have today is unacceptable, but moreover it's led to terrible burnout rates. So there is another hit on humanity. So clinicians across the board, it's the highest in the medical profession, doctors, nurses, and across the board for, uh, not just, uh, uh, Bernie out, but doubling of errors with that, but also depression and a very high rate of suicides. Hm.

Elliot Gerson: And I will come back, I hope to too many of those things. And, and Chris, um, you know, talk to us a little bit about, about designing data for health care as opposed to just sort of using the data that is already there, which may be part of the answer about just perpetuating bias that may be in that data.

Christopher G: Yeah, that's exactly right. So one of the big challenges with, with machine learning and AI is that typically you give data scientists a set of data. And I guess usually it's not a stack of paper, it's electronic, but you give them some set of data. You ask them to answer a question. And at least during the last decade or so, most of the time they probably don't have the right data. It's not been collected in the right way to actually answer that question really well. And so speaking from sort of the earlier side of things on discovery, one of the things I'm most excited about, I think a lot of people are, is trying to design data sets so that they have less bias in them. Right? And so if you can actually have perfect data, which I know is just to hope and not actually possible, machine learning would be even much more powerful than it is today.

Christopher G: And a lot of the bias would be absent, uh, from it. So, and you know, the last point I'll make is everything in moderation. Um, and AI has places where it's exceptionally powerful and humans have places that they're exceptionally powerful. And we talk a lot about the bias of AI, but I think it's also really important to remember how biased we are as people. Um, and so in some ways actually if we can find the right balance and have both of these things together in the loop, it actually feels like we could, we could move, move humanity

forward pretty profoundly if we're, if we're integrating AI and human systems very strategically and appropriately. I would hope that we would not only bring in more humanity, but we'd also reduce the terrible medical error rate. So, you know, physicians bring cognitive biases all the time. It's not just machines that we have to worry about in terms of bias. And I think there's real opportunities here to, to get to another level of him.

Elliot Gerson: No, let me just stay, stay with you for a minute. You talked earlier about privacy and they are, lots of things talked about. I like to come back to Andy and Eric mentioned issues relating to sort of depression and mental health met. Let, let's just talk a little bit about these issues in the context of mental health for a minute. Uh, and uh, what, what are the most promising aspects of artificial intelligence in your mind with respect to mental health in particular? And also perhaps what are the special challenges, some of which may actually relate to privacy, uh, in the context of applications of artificial intelligence to mental health?

Mildred Solomon: That's a great question. And it speaks right to the oxymoron in the title of our session because there's both very promising mental health applications and very concerning ones. So I'm going to talk about two different areas. One would be clinical interviewing. You know, we have a shortage of psychologists, psychiatrists, psychologists and mental health workers. And so there are some promising applications that are, um, focused on clinical interviewing. For example, after the Iraq war. And since for quite a long time, the Department of Defense has been trying pretty earnestly to provide mental health services to vets and not with not that much success. And part of it is a lack of, of clinicians, but part of it is also cultural, um, a feeling of that it's a stigmatized thing to seek. And so they've developed an APP called Sim coach. It's not meant to replace the therapist ultimately, but it's meant to address the issue of stigma and help people feel more comfortable.

Mildred Solomon: So it's a kind of exploratory app where you can talk to an Avatar in a sense. People seem to prefer to tell their secrets and fears to, uh, to a non human entity. So it has some advantages in that as well and they're showing some progress. There's another application that I think is more illustrates the more worrisome side, which is I think called Sim, um, Sim Sensi and it uses sensory data to, um, infer the mental state and emotional state of the user. So it, it has, it gets voiced information, body gestures and, and um, and face facial recognition and makes a lot of inferences about the emotional state of the user and the end documents them again, on the path towards a referral to a provider. But you as the user don't really know how you're being interpreted. And many of these effect, uh, recognition systems have been designed on the assumption that we can characterize our complex emotional range in terms of six emotions. And I don't know about you, but I know that I sometimes feel multiple things at once. Happy and sad, curious and afraid. And this system doesn't allow for any of that kind of nuance and it infers about you without you

knowing how it's labeling you. So I, it's a privacy issue in a very deep, in a very deep sense.

Elliot Gerson: All right, Eric, let me just ask you directly. Is, is the black box problem for artificial intelligence in healthcare? You know, a deal breaker? I mean, you know, we're, we're all used to principals, you know, like peer review and validation of, uh, you know, by very rigorous methodology. I mean, if we don't know what's in the black box, how in the world are we ever get to trust it?

Eric Topol: Sure. Well, this is really a big controversy as we bring artificial intelligence and into medicine. The first thing is to recognize that there are so many things today in medicine where we have no clue how it works, what we're doing. We just do it. And it's just the standard of care. Uh, and that list is long. So the question now is, are we going to hold a machine to a higher standard? That is if we, if it's validated to the hill, you know, prospective randomized trials and we still don't fully understand it, how does it really work? Uh, then the question is, are we going to hold it up and not deploy it until we deconstruct it and have it fully explainable? And essentially the field breaks down to two different camps. One, the computer scientist say, if it works, I'm using it. Uh, and a lot of physicians and medical folks say, well, we really need to understand this better.

Eric Topol: We haven't really gotten to that point, Elliot, where we've seen things fully validated. So we don't know. Um, because there's very little so far of this prospective validation, it's limited. But soon enough, over the next year or two, we'll start to see that be challenged and we don't know exactly how it's going to be playing play out. One other thing about this is in Europe, the new standards about privacy and the Gdpr have said, unless the algorithm is explainable, we can't, it can't be used. And uh, one other thing that's really exciting in the field of AI is deconstructing the algorithms. So now, so much effort is being done in to this so called ablation analysis where basically it's, it's going backwards to find out what the features are. So there's some light here that by the time we get there with the validation, we'll also have gotten there with the ability to understand algorithms

Elliot Gerson: better. So have you thought about advising our Congress and development of an equivalent of equivalent by [inaudible] in this country or even an in California? I mean, is there a principle like that in the California privacy? There

Eric Topol: be, there isn't yet. A Europe is way ahead of us and I echo Millie's concerns that we are really behind in standing up for privacy. It's medical privacy's different than, you know, just day to day, uh, privacy and our data, our medical data is at least five fold more valuable on the dark web. And it's being the cyber thievery and hacking that's going on now is rampant. Wow.

Elliot Gerson: Yeah, please Chris. And then I'll follow it. It's

Christopher G: one of the things that comes to mind for me is this, maybe this implicit assumption that we have as humans that we can understand. And you know, I think one of the great prospects of Ai is the potential to generate inferences and predictions that are really useful to humanity, but that are far too complex for us to be able to actually understand. And so this is one of the things that I find challenging around this idea. To your point, Eric, you know, statins do a lot more than we thought they did when they were approved, right? This happens all the time in medicine. And yet if you imagine the thousands of diseases today that we truly don't understand at all, we don't have treatments for. And if you study biology and medicine, basically your conclusion typically is that we don't know very much about it as humans, right?

Christopher G: With a tiny, tiny fraction. What that, what that leads me to believe is, is that this human assumption that we must be able to understand the algorithm might be one of the things that holds us back more than anything else. Show Chris you, you are the resident technologist on this panel. So, or one of them on the business, on the business side. Um, so, uh, and obviously, you know, technologist, business technologist have lost a lot of trust recently. And, uh, as someone really working at the intersection of healthcare and technology, particularly at a time where more and more people are questioning technology and its ethics, morality, it's fundamental. Goodness. How do you, you know, earn and maintain trust? I mean, do you, how do, how do you think of of that and, and it's important in, in, in, in, in what you do. It's a hard question.

Christopher G: When you're trying to build a company that's doing something that hasn't been done, oftentimes you find yourself pretty deep and sort of siloed in the, in the sorts of problems they're working on. And I was at a conference, um, in near London, uh, about six months ago. And essentially we talked about the whole host of people from lots of different disciplines, psychology and data science and economists. And we talked a lot about this dystopian future where AI meets biology. And it was sort of a wake up call for me in the, in the context of what we're seeing in the tech industry. And so, um, it's triggered me to try and ask a lot what could we be doing now that's actually setting us up to, um, be in a tough spot down the road to be doing things that we don't feel really good about.

Christopher G: And I don't know the answers yet. We just have to be aware of it. That's the first part. And we have to ask ourselves constantly and hold ourselves to a high standard. Um, but the things that I know we do now, you know, there's this idea in the tech world that you move fast and break things right and maybe it'd be a little bit controversial. I actually subscribed to that in a part of our business and that's the part of our business where we're developing algorithms around human cells that we grow in, dishes that we do experiments on in the lab. And we iteratively work on around that science. And I actually think it's really good to move fast and break things there. At the same time, we currently have two drugs that are in human clinical trials where we're actually testing a medicine in real humans. And our approach has been as about as conservative as it gets in

terms of running those trials because that's the one place you really don't want to move fast and break things. And so to me, I think it just comes down to being practical to try and look to the future, to reading books by real futurists and technologists and talking to really smart people to try and think of all the ways that you, all the ways you're not thinking about,

Elliot Gerson: hi Millie. Okay, cool. Who owns the algorithms? I mean, I mean, what, what is the importance of ownership and all the discussions? I mean is, do you know, do, are they privately owned? Should they be privately owned? Should the government own them? Uh, you know, and how, what's the relationship between trust and ownership here?

Mildred Solomon: Absolutely. I mean, ownership is really occur. These applications are likely in our country anyway to be commercialized. Right? And so one of the biggest issues around trust is what is the relationship between the owners of these inventions really and the users of them. And Are we going to make, take steps to make sure that low cost applications are marketed and to low to low cost to low paying, um, low income audiences. So I think they're going to be owned by both the public sector and the private sector. But increasingly obviously venture capital is very, very interested, um, in a wide variety of applications even in the use of AI in, in finding new drugs, which was an application pretty recent in terms of what you're doing Chris and also in terms of identifying people who are appropriate candidates for Onco oncological clinical trials. So there's, it's, it's going to be multi ownership and the question is what responsibilities can I also respond to the, to the privacy concern.

Mildred Solomon: I think it's worth noticing that we're in a paradigm shift. We've talked about personal health information is really sacrosanct. We have the HIPAA regulations and we have a commitment, a sense that this is ours, it's mine. It's an individual thing that's I should have control over. I think we're making a shift here that we might not even be aware of to seeing data, big data and artificial intelligence is use of it as a, as a social good. And I think that's why we're intention because we're trying to get the community benefit from something that we've identified as a singular, a singular, right. So we're figuring this out so we know quite know how to, how to negotiate the tension between that. And the other piece of privacy is rarely talked about in the privacy sector, but it really is another meaning of privacy, which is the right to be left alone, freedom and liberty.

Mildred Solomon: And one of the concerns about AI as real and the trustworthiness and who owns it is, is the extent to which, um, this can lead to surveillance of all kinds. And then your question about who owns the algorithms is really important, but I think the problems can come both from the government, you know, ice, just a modified their algorithms to make sure that everybody spit out as 100% detainable. This has been reported by AI now in their annual report. So, in that case, the government owned the algorithm and decided to massage it to, to fit political purposes with tremendous consequences. And then we're, it could also be owned privately and used for marketing purposes to better understand

consumer behaviors and influence our choices in ways that were not even aware of through choice. You know, the way we frame choices for people since that's kind of a long answer, but I think the good, the liberty issue is really one another side of the privacy issue. Can we should talk about them separately. So Eric, how do you,

Eric Topol: how do you see those issues, the issues relating to Liberty Privacy, surveillance? Yes. Uh, this is complex stuff. Um, you know, just to go back first to the ownership, you know, this is software algorithms, so that's really not very expensive. What's expensive is this validation, doing large, developing these large and annotated datasets and then putting them through tests and uh, you know, start ups just can't do that. They, this is a tech titan thing or you know, this is a, as Emily mentioned, investors. So then the question is you, how we as an editor of a journal, this comes up every week, now we're going to publish your, your algorithm, but we want it to be open source. And they say, no, no, we can't make it open source. You know, we put so much into this. Got To get a return. Yeah. So this is a real thorny issue that is unsettled.

Eric Topol: Um, what'd you don't want to, you don't want to de incentivize the big efforts that it takes because what we're seeing now is a remarkable ability to take medicine, um, you know, to the hinterlands, you know, to do things that we never would conceive of to level the playing field because you can have, uh, algorithms and sensors and data that is really cheap, cheap chips. And this is really not, um, holding us back from, from reducing inequities now. But to this point about the, the liberty and privacy, my biggest concern is that we don't own our medical data. And because we don't own it, that's why it sits on servers and it's a target for the cybers thievery, uh, uh, issue. And it's not protected. So if you're going to do artificial intelligence, and this is coming pretty quickly, where you have your virtual medical coach and we'll be giving you feedback as to how you can prevent this condition or better manage it. And if you don't have all your data, which nobody has in this country, nobody has all their data, uh, it sits in different places and it's very incomplete. So your inputs to the, uh, neuro network are very incomplete and it's just garbage in, garbage out story. So we're compromised. And so that's part of the situation we're in now. We were having assert our rights. It should be a civil right to own and control your medical data in Estonia. That's the way it works. And if Astonia can do it, why can't we do it?

Elliot Gerson: All right.

Christopher G: I think that's really important because if you talk to a lot of data scientists, they'll tell you, at least the ones that I've talked to that, um, algorithms themselves are pretty likely to become a commodity over the next decade or even less. And so the fundamental value actually lies in the data itself. And so owning an algorithm may actually not be the important part. It may be the access or ownership specifically of the data because 10 years from now you'll be able to retrain the perfect algorithm on whatever data set you have in the staff

of a blink of an eye. And so I think that ownership is a really, really, really important piece.

Elliot Gerson: Yep, totally agree. So Milly, let's, let's come back to some of just the practical problems of bias. I mean, you know, when in terms of big data, we, you know, the and, and inclusiveness, I mean, in many cases are, you know, the data significantly underrepresented women, are they underrepresented people of color? Are they under represent people from lower socioeconomic statuses and then their medical ma many medical conditions that we, that we know relate differently in the, and the best treatments are, are different depending on those kinds of characteristics. H how, how do we deal with those problems? So the, you know, so future medicine can be more inclusive and level the playing field as I as, as opposed to, you know, doubling down on the bias of the historical data based on who has the privilege to receive care. Right.

Mildred Solomon: Let's distinguish between three kinds of problems that are all fairness. They're all problems that arise from lack of inclusivity. There's sort of, there's bias in reflecting what we learned. So if we use a clinical historical clinical research for the machine to learn on, we know women were excluded from clinical research for decades and IHS corrected that. But if this is older data, it's not going to be corrected. So we get skewed knowledge. That's one form. Another form is that what we build doesn't work for women or for other, for other people. So an example would be, um, um, it was the kinds of games yes. And Yeah, and in the non health area there's been some gaming games that have been designed by men. And then when women took them home, they couldn't, you know, they couldn't run them. It didn't recognize their voice or their body language or something like that.

Mildred Solomon: And, and then the third are sort of deeper fairness issues about what's the use to which the application is going to be applied. So right now for example, um, you know, I gave you the, the ice example around migrants, but one that's more squarely in the healthcare sector is that there are decisions systems that the government owns. Again, back to your ownership question that are making decisions about how many hours of care somebody with disabilities is allowed to have. And they've, they've used to be assessments that were made by social workers or nurses who visited the home. They're now being made by decision systems and they are being designed to save money and therefore drastically reducing the hours of care that somebody gets it home and without recourse. So it's a human system they could have built, maybe built those algorithms but had a human in the, in the mix. So I don't know if that's answering your question.

Elliot Gerson: I'm not sure there's any simple answers to it, but it's very helpful. Do you have any thoughts on, on that, Eric? Or actually to me, either way, I'm going to open it up to the audience. So many more questions and we get to have time to answer, but I just thinking of your book and you made reference to it earlier,

another definitional questions. What is deep empathy, right? And how can we all get more? Yeah.

Eric Topol: Oh, scientists here. So, um, that's what we're missing now, which is that human human bond. So as this machine, you know, you're human intelligence probably not going to change over the next century. Uh, Denny, significant expense hopes on, well yeah, maybe it's a little bit, but you know, to, to a substantive way you won't quit. Machines are just getting smarter, taking on broader tasks and uh, well we have to do is getting more human and what we, what we have done in the medical sphere is get less because we have let these, this erosion occur. Um, so deep empathy refers to the, all the aspects of, uh, that human bond, the presence, the, uh, you know, listening. Typically a patient is interrupted within seconds. They never get to tell their life story or any story actually. And uh, that's really rich and it'll never be digitized. You're never going to be an algorithm to know what the patient's life story's all about.

Eric Topol: So to get us back to this, uh, time, which by the way was, if you go back in history, uh, in the 80s and before there was a deep relationship, the patient doctor, uh, was a precious bond and it was trust. It was big and it just fell apart. So deep empathy is basically getting that back. Uh, it's basically an outgrowth once you start to have time to spend. Why did doctors and nurses and all the people that work in healthcare, why do they go into this profession in the first place and why are they so dejected?

Elliot Gerson: It wasn't a fill out insurance forms, I don't think.

Eric Topol: No. And they did that. And why they're so dejected is because they can't execute their charge, what they, their mission. If we can start to restore that, then we can start to see deep empathy. Okay.

Elliot Gerson: Uh, we have just under 20 minutes left. So I would like to turn things to the audience. I see already. Lots of hands coming up, lady in the second relative. Please wait for the microphone for people who will be watching this later.

Eric Topol: Okay.

Audience Member: Hi. Thank you. I'm so I represent um, part of the community that has lost a lot of trust in doctors and scientists, which is a problem because I am also a doctor and a scientist so I don't have a high bar time records. I like that he psychiatrist and then I the audience. Um, and I spend all my time now with patient advocacy groups and I represent about 40 of them. Um, small rare diseases, small groups that are literally going to revolt soon if we don't do something, I spend a lot of time saying, please do not bomb NIH. Please do not lay on the floor of Nih. Please do not do these crazy things at the aids activists did back in the day, which they hold up and value. And if and, and they love your second book, Eric. Just so you know. Um, anyway, so I'm wondering, one of the things that I fight fiercely for is, is whoever the end user it is, and it's, in my case, it's the patient

advocates to have a seat in coproduction ones that are intelligent. I mean I'm trained as a scientist and a doctor that, that understand economics as well to do this along with you and then go back to their communities. And they also serve as a sanity filter if you're trying to do something that's just not gonna work for the patients. Have you guys tried incorporating that boys at your various stages to help bring back deep empathy or to help solve some of the problems of how do we make this usable?

Eric Topol: Hmm. I think that's a question for you. Oh, well, you know, I think your point is really important with respect to the medical organizations are not advocating for patients, their trade guild that are protecting the reimbursement largely or other matters for their doctors and specialists and whatnot. And so we have a big gap in what any, any patient advocacy group is relatively small and doesn't have enough power to take on this. We really do need doctors to organize to get this on track. And that doesn't mean to organize for themselves, but to stand up for patients. We haven't actually seen that doctors being activists until recent times. Like when the NRA said, stay in your lane, Austin doctors activated. And so there is a chance now, just like you alluded to as the, uh, the, uh, in the eight zero when patients, uh, took a charge, we need to see that throughout the medical profession.

Eric Topol: So that's what's lacking right now. But there's a really important opportunity and this is something that I'm hoping that we'll see. Take, take, hold. I just had one other point about bias and fairness about Millie. Uh, brought up and you did Elliot it not just the human bias that is really the, the, the localizing the culprit for, it's not the algorithm there are neutral what goes in, but the thing is medical studies are not doing diverse patient inputs. So when a d, a big dermatology app that could diagnose skin cancers with 99% accuracy, who is only in white people, there were no people of color that were put into the study. And so this has to change. We've already seen this, we blew it with genomics, all European ancestry. So that's another big issue about fairness that has to be, uh, brought it to the, uh, for uh, the gentleman in the third row and we just wait for the microphone and then I'll go to someone over on this side.

Audience Member: I thank you all at my name Jonathan Cohen. I direct the public health program at the open society foundations. Um, there is increasing recognition, um, including in a really powerful set of case studies in the New England Journal of Medicine and the need for clinicians to be able to diagnose the social and structural determinants of health. So for example, to be able to trace respiratory symptoms to substandard housing and indoor pollution or to trace unwanted pregnancy to intimate partner violence and a thousand other examples of that. What would you see as both the opportunities and the, the limitations? And risks of AI in assisting that kind of diagnosis. Millie,

Mildred Solomon: I think that's a wonderful place where an artificial intelligence app like application x exists because it's getting missed. People are, physicians are don't have the time. It's a hurried, harried context in clinical care. I could imagine an

APP that would routinely ask people questions about their symptoms and be designed to know that if, if the person has asthma to ask them questions about, you know, the quality do they have, do they have an air conditioner at home and all these other questions. I think that could be designed and would be a great use of, of ensuring that we, that this happens. But again, it would have to be built with the, with the clinical team and the machine so that we use that information and there was a protocol and a pathway for it in terms of how it could be badly used. You know, a lot the social determinants of health have a lot of proxy measures.

Mildred Solomon: Like you could, you could, you could build algorithms that, um, look at smoking, let's say, but smoking isn't just smoking. Smoking is almost a surrogate marker these days, at least in the United States for ses and and so it's a, it's an a way a surrogate marker for poverty. And if we started to build algorithms that looked at social determinant variables in the wrong way, we could penalize poor people and end up charging them higher insurance premiums for example, because they're showing higher levels of risks. So I think it's really wonderful to bring social determinants into the conversation, be less biomedical and more social. Um, but it's a two edged sword, like most things and we just have to be sensitive to how does design something

Mildred Solomon: like that.

Elliot Gerson: I don't want us to come to sampling bias certainly as people on this side of that. So, so in, in the third row right here.

Audience Member: Thank you. My name is [inaudible] and I'm with the Mashantucket Pequot tribal nation. My question is, while we try to use AI to free up time and provide for more trust and empathy, why are we not asking the question about redesigning payment models and actually incentives to free up the time and unshackle the physicians and the clinicians from being tied to note bloat and Emr dudes,

Eric Topol: right. So, uh, this is a problem that's specific to the United States. Unfortunately in the UK, which I had the real privilege to work with our last couple of years. They're going after this and they don't have any barriers like you're alluding to. So we have to do this. What you're seeing, we also have a big F. I have my biggest fear about the whole AI era in medicine is that the squeeze that's on right now by the administrators and the people that are overseeing the financial matters, the business, the big business of medicine is to make things worse because you've got all this productivity. Why don't you have doctor's see more patients and radiologists read more scans and on and on. So to turn this around and we'd go back to this organized effort, uh, of the, of the medical community to not let that happen to, to do what you're saying, but right now the default mode is to make things worse, if that's at all possible.

Elliot Gerson: All right.

Elliot Gerson: I in the back and the orange jacket. Yeah.

Audience Member: Uh, hi. I'm a data scientist by trade. Um, and I was wondering, um, open source AI technology in natural language and image processing, video processing as sort of revolutionized a lot of things, but I've only seen that flow be one directional into medicine, not so much out of. And I'm wondering, have there been any good, either opensource developments in medicine that are making it to other medical communities or is it just really Belkin ISED and not just balkanized but not transferable?

Mildred Solomon: I don't know.

Eric Topol: I'm not sure. Right. I'm not sure if I understand what you're driving at. Where would it go? Go. Just wait for the microphone if you went there.

Audience Member: Uh, like, well for radiology for example, like you can find, have a algorithm that could find a tumor in an image that a human couldn't notice, then would that not be shared with hospital a over here because hospital a has different radiology machine and everything else is different. I'm saying the spread of Ai is,

Eric Topol: yeah, so like you're, while you're bringing up a point that's already been shown in certain studies that the findings, the algorithm, our venue specific, like they work really well, not pneumonia diagnosis in this hospital, but not in that hospital. So that is another thing a t has to be kept in mind, but as far as, that's where the whole commercialization of algorithms and transferability, you know, we're, we're, we're, it's already happening, uh, particularly in radiology in the u s much more in China. But no, there's no barriers. Once the problem is how good is it really because of this specificity, whether it's, you know, the types of patients who went in or the, or the, or the hot site, the venue, that sort of thing.

Elliot Gerson: Uh, David,

Audience Member: I thank you very much. I'm curious about the global implications of this. And um, you mentioned a little bit about sort of garbage in, garbage out data and I'm curious for all the panelists, um, thinking about the global scheme and if the data that we're getting for these algorithms is all going to be coming from places where we capture large, large says did pieces of big data Europe and North America, how do we make sure that we're inclusive of places where most of the world is Asia and Africa? So how do you, do we capture that data so it's inclusive?

Eric Topol: Yeah. Uh, well, I, I, I'm putting forth with k through we of China, uh, proposal, which you'll see coming out soon that we work together to develop a planetary, this is back to Millie planetary, uh, infrastructure. Now that's probably not likely given the current state of affairs, but, but if you think it's true, there's this thing called federated I AI, which is really exciting. And what it means is you basically, you can tap the data locally without ever pulling out of a country or a site or

whatever. And now that's getting momentum. Th Th th this is what's so exciting about Ai. Longer term is a planetary a information resource. We lit the whole species, learn from each other. David, to take your point a few steps further that someone comes in ill, let's say they have cancer, did you could go through this mill, billions of people of all types and you could match up that person as closely as possible. So called digital twin and then have a, um, a treatment and outcome. But you could predict rather than from some randomized clinical trial that doesn't apply to that person. You could actually match the person with other people. And so in order to do that, you need lots of people of diverse ancestry with all their data. And that means sharing. And that means if we can't do it with China, we can't do it alone in the u s as you're alluding to, we what? We do need to do that.

Mildred Solomon: Yeah.

Mildred Solomon: Oh yes, please. I was just going to jump in and say, you know, obviously trust is the assessor is the essential thing. We have a history of having done clinical research, traditional clinical research trials in the developing world in exploitative ways that's been being corrected in some ways. But there's that legacy. And so not only have we been exploited even how we've done the trials, but we've also neglected ensuring that the benefits of the research are going to go back to those populations. Kifle Lee is talked about how the wealth it's going to be generated is most likely going to be in China, the United States in Europe. And he's been starting to think about how there can be a global economic landscape in which there's benefit brought to countries veteran not generating the wealth and that are actually helping to, you know, that are, that are on the losing end of that equation. And talking about what our responsibilities would be to, to make sure that the benefits are available. And I think only that way can we build that kind of trust like you were in the front row.

Audience Member: And the question of time with patients, what I've seen in Europe with national health and Britain and France, uh, Eh, the budget is continually collapsing because of a promising more than you can deliver. And as a result the system is going to say congenitally broke. And so any, uh, efficiency you would get from a AI system I think would be used to reduce the number increased patient load or their existing doctors rather than give you more time for existing patients. Um, so I think there's kind of a wild goose chase.

Eric Topol: Well that's what will happen if we don't, if we don't use this, is to turn inward and make things better. But I just want to point out how broken things are here in case you haven't kept up with this. Um, you know, we, we have, we pay in this country \$11,000 per person for healthcare each year and we have the worst metrics, life expectancy, three years in a row going down the only place in the world. And the only time in American history, we also have the worst maternal mortality, infant mortality, childhood mortality. I could go on. And the UK, which you mentioned, I also touched on it, is at \$4,000 a year instead of 11,000 with far superior outcomes. So that's not the only country. Of course, the whole Oec

d has 36 countries, 35 of them above the u s uh, for all metrics. So we, this system is so bad. We say that to say that you can add, you can hardly make it worse.

Elliot Gerson: Well, on that note, we have time for one more question. Actually an,

Audience Member: there's a question in the corner and this'll be our last question. And then we try to end these sessions on time. My name is Erin [inaudible]. I lead the Aspen Institute Science and Society Program. I'm just waiting. When does something cease to be AI and witnesses to become a standardized technology? And how much of that frame of AI elicits fear and these kinds of discussions appropriately and possibly even unnecessarily truths? Yeah. And one of the things that I think is actually quite interesting is that if you put a lot of pretty believable people in a room, data scientists, software engineers, et Cetera, and you ask them to define Ai, you actually, we'll often see a lot of arguments. Um, so it's, it's, it's self, a pretty nebulous term. And in many ways it's sort of become marketing. And I think as we start to get used to seeing it in society, as we realize that AI is being deployed and it is in, in many settings, um, it's being deployed all the time when you're clicking around on websites and optimizing ads for you. And, and so I think it will become normalized just like every other technology over, over, uh, you know, a decade or so. I think you had a no, that's fine. Yeah. Well, it's appropriate to say that the sun is coming out again.

Elliot Gerson: All right. Let's see. Thank these three. Wonderful.