

FINISHED FILE

BOSTON UNIVERSITY
SCHOOL OF PUBLIC HEALTH
APRIL 24, 2024
12 P.M. CT

Services provided by:
Caption First, Inc.
P.O. Box 3066
Monument, CO 80132
719-941-9557
www.captionfirst.com

This text, document, or file is based on live transcription. Communication Access Realtime Translation (CART), captioning, and/or live transcription are provided in order to facilitate communication accessibility and may not be a totally verbatim record of the proceedings. This text, document, or file is not to be distributed or used in any way that may violate copyright law.

>> SANDRO GALEA: Good afternoon, good evening, good morning, wherever you are. My name is Sandro Galea. I have the privilege of serving as Dean of the Boston University School of Public Health. On behalf of our school, welcome to today's public health conversation. These events are meant as spaces where we come together to discuss the ideas that shape a healthier world. Through a process of free speech, open debate and generative exchange of ideas, we aim to sharpen our approach to understanding. Guided by expert speakers, we work towards to develop a deeper understanding of what matters most.

Thank you for joining today's conversation. In particular, thank you to the Dean's office and the SPH communications team, without whose efforts these conversations would not take place.

In the past couple of years artificial intelligence has moved to the center of discussion. The emergence of AI echoes the early years of the Internet, as we come to grips with a new technology's potential for both good and bad. Today we will discuss AI's implications for public health. Guided by our speakers, we will talk about how we can use AI to advance the health of populations and how we can address some of the challenges posed by AI. I look forward to learning from the conversation.

Today's event is being moderated by Jennifer Strong. Jennifer Strong is the host and creator of the SHIFT podcast.

She is a journalist covering the impact of frontier technologies on the way we live and work. She is the creator of top science and tech podcasts for newsrooms that include ProPublica, The Wall Street Journal and MIT Technology Review.

Her reporting has been recognized by awards juries dozens of times, including six Webby and three Ambie or Podcast Academy Award nominations. She has been a keynote stage host and moderator at the United Nations General Assembly, SXSW, Web Summit, Bio, Emtech, AI For Good Global Summit, The Future of Everything Festival and many others. Jennifer, it is great to have you. Thank you for being with us.

>> JENNIFER STRONG: Thank you for that introduction. It is my pleasure to be moderating this discussion, which in the interest of preserving as much time as possible for, I would like to dive right in and introduce our speakers.

First we are going to hear from David Cutler. Dr. Cutler is currently the Otto Eckstein Professor of Applied Economics in the Department of Economics at Harvard University. Professor Cutler holds secondary appointments at the Kennedy School of Government and the School of Public Health. He is also a research associate at the National Bureau of Economic Research and a member of the National Academy of Medicine.

Next we will turn to Alondra Nelson. Dr. Nelson holds the Harold F. Linder Chair in the School of Social Science. Between 2021 and 2023 she served as deputy assistant to President Joe Biden and acting director and principal deputy director for science and society of the White House Office of Science and Technology Policy. Last year Dr. Nelson was included in the inaugural TIME100 list of the most influential people in AI. And in October she was appointed to serve on the UN High-Level Advisory Body on Artificial Intelligence. She is also a member of the National Academy of Medicine.

Finally we will hear from Mr. Greg Singleton. Mr. Singleton serves as Chief Artificial Intelligence Officer at the U.S. Department of Health and Human Services, where he is leading the department's efforts to supercharge government and private industry health efforts through safe and responsible use of artificial intelligence. In his prior roles at HHS Greg led the national COVID data collection, that's across 6,000 hospitals. He played critical roles in tracking and distributing COVID vaccines and led them to the availability and distribution of millions of doses of COVID therapeutic treatments. We will be coming to you first. The floor is now yours.

>> DAVID CUTLER: Thank you very much for the very kind introduction for having me on the panel. Thank you to the Dean as well. I have been able to watch a couple of the programs and they are just terrific. I'm honored to be able to be a part of it. I want to start by I guess sharing some economic thoughts about how do I think about AI and its relationship to public health. So I have a couple of slides although. There is going

to be nothing too bad there but nothing too much there. Let me sort of start off with a question that I think is in people's minds or at least in my mind which is AI going to be the next killer app in public health. Is it going to help us reshape the public health system for the 21st Century.

So let me -- and so let me give you a couple of thoughts. I'm not going to give you a definitive yes or no. But I will try to sort of highlight areas that I think are possibilities and then a couple of challenges as well. Let me first start by defining terms. So AI I'm going to think of as having two components. First one is machine learning which is a way to ask a computer to learn something rather than just give it wrote instruction. Whenever Jennifer says Dr. Cutler please turn on the microphone. Under this circumstance you should turn on the microphone. So machine learning is the first component of it. And then the second component which is important is natural language processing which is the ability to take the spoken word and turn our written word and turn it into something that's machine reader and structured so you can do something with it.

So the example of that would be taking physician notes and turning that into something structured that one can work with. So what do we want to think about AI as doing? So in my mind, what I really want AI to do is to help us with the Triple Aim. Sometimes the quadruple aim. We want better health. Better medical system and we want lower costs. And, you know, in the U.S. we obviously fail on all three of these. Our population health is not very good. Our per capita spending is atrociously high. Even if other countries that do much better there is still debates about can we do better still.

That is can we have a medical system that's still better and in terms of cost and the experience of care is better. So I would say the problems are more acute in some countries but I don't know anywhere where this is not an acute issue.

So let me talk about the possibilities. So there are three things that I'm going to highlight and highlight three challenges as well. The first one, since this is a public health crowd, I can't help but talking about the possibility of early warnings of pandemics. And that is something that we should absolutely be able to take advantage of AI to learn about. The story that I keep in my mind if you want to know where the next pandemic is likely to come from, keep track of chicken prices in Thailand. Why is it a chicken market in Thailand is going to do anything for you? And the answer is as best I can understand it that it is possible that the next pandemic will start by killing wild fowl in Asia. And that chicken prices will go up in the market. One should be monitoring chicken prices around the globe to get some indication about is there a problem with wildlife.

So that's just one example of where you can train a machine to say, you know, look at these -- here is a bunch of indicators. See if you can use them to pick out when you are

having problems and then alert them, alert us to them.

So that's the first possibility that I would like to see. The second one that I would like to see is simplified administrative processes. When I think about medical care systems in AI I want to differentiate the administrative side and clinical side. And the administrative side is where the first big impact is going to be. U.S. health care's distinguished by enormously high administrative sense. Higher administrative costs.

That contributes as well more images and things like that. But it is the administrative side and that a lot of that is people. So a very common thing that a person does is takes computerized information from an electronic medical record. The minute I tell you that one person is taking information from one record and making it be in another record then you think gosh, why isn't my computer doing that. That's what some of the natural language processing is doing. I would like to see fairly big improvements in administrative costs, the most -- many health systems have entire buildings devoted to people. And I would like to see a situation where we don't need those people. Electrons tend to be a bit cheaper than people are these days.

So that's the second area of possibilities on the administrative side. The third area on the clinical side and what I want to think about here is the possibility of using AI to do diagnosis cheaper. So we have all sorts of situations in rich and poor countries. This is in rural India where we are trying to diagnose something. We are trying to diagnose the chest pain or diagnose diabetic retinopathy. Or from a brain scan or blood test. It is not possible either for financial reasons or access reasons to do that for everyone.

On the other hand, there are often cheaper things to do. We can have a one lead EKG attached to a chest or you can have cell phone attachments that can get a guide to heart and lung sounds and those are not as good as what you can get in a hospital or doctor's office. Can you develop AI to turn those into something that's much better. But the question is can you do well enough so that in areas where access is poor here in rural parts of the world but even in richer countries there are areas where access is poor, can you find ways to do better, cheaper through the use of AI. All of those are on my possibilities and would have phenomenal impacts on public health and save money as well.

On my pitfall end I will give you three things. The first one is hallucinations and I don't refer to myself in this, all I do sometimes is hallucinate, but I'm thinking about the AI programs. They do have this problem is that many of them hallucinate. If you go back to the billing form to take the very simplest example, given what was actually done then nobody is going to want to pay attention to what the AI says. We have to discipline the AI that way.

The second one I call avoiding embedded isms, what do I mean

by that? The medical system is full of racism, sexism, ageism, discrimination against people with different sexual orientations and gender identities and so that's all going on. And sometimes what we say is well, I want the AI to mimic what a doctor would do. So I want the AI to mimic what a radiologist would come up with. In fact, doctors have all sorts of Isms in what they do as well. They are more likely to diagnose things in upper income people than lower income people. You want to train AI to ignore that. But the key in that is not to say I'm going to ignore race to the patient. I'm going to train AI to predict what doctor would do. So you need to train to the ground truth. How much pain would a person have a given pain image. That latter one has a lot of bias seeped into it that the former one does not. I think of this to remind me if you train an image to figure out who has pain and who should get a knee replacement. If you have knee pain, it will allocate many more knee replacements with people are lower incomes and minority groups of

Finally, we want AI to be neither overused or underused. How we pay for the AI is going to be important. So if the clinician has to pay for the AI, then somehow we have to pay to cover the cost of the AI or it won't be used. So I think about AI that's used in an embedded in a particular technology. On the other hand, if it is or the -- if the AI comes along for free, you buy the retinal scanner and the reading of the image comes along for free. You don't want to pay extra for the AI component because the docs will do way too much of it. Just like with any medical technology when we come up with a new drug for patients, we want to use it for the people and whom is appropriate and not use it for the people who is not appropriate. We use it in more people than we should use. We don't use it in some cases where we should use it. And typically that's accompanied by reimbursed incentives.

People who are overused are paid well and people who are underused with paid poorly. Unless we get the economics of it right there is very little hope that we will get the clinical don't components of it right. I want to stress some of the economic components here. Can it be done and the answer is well, I don't know. You look at this graph and tell me. What this graph tells is the share of the population that received the latest vaccine. Like the world did really, really well because there are vast, vast part of the world. And then on the other hand, the world did very poorly because there are vast, vast parts of the world that it is lower. We can get anti-retroviral medication to people who are HIV positive. We can do lots of things well. Some things, we don't finish it up and don't do it right. I'm hopeful that we will figure out some of the economics here so that we can get more of the technology to do what we want. And I think I will just stop there and thank you again so much.

>> JENNIFER STRONG: Thank you. Dr. Nelson, now it is

over to you. I believe you are muted.

>> ALONDRA NELSON: Sorry. The particular use case is really matter, domain expertise matters for how we benchmark the possibilities. I'm delighted to be speaking with a group of people with a shared interest in public health. So we were asked to consider the role of Artificial Intelligence and the health of populations and how AI can be used in public health for good or bad consequences and to think about those potential use cases. Delighted to be in a conversation with David and my former colleague, Greg, about these really important issues. I approach these questions as a recent political appointee in science and technology policy in the Biden/Harris Administration.

And also as a sociologist of science and technology and medicine. Applications of artificial intelligence have been in most people's lives for many years. That the world took notice of AI in a new way. We know by the start of 2023, more than 100 million people have reported used ChatGPT and millions of others had engaged in some way, some form of generative AI. A headline of our conversations today is that the public because we have had these tools that were really enterprise tools that were tools coming in the background of our lives, allowing us to open our iPhones with our face, we weren't thinking about AI in public discourse in a major way.

After the fall of 2022, the public became a major kind of stakeholders with active interest in thinking about the uses of AI. I want to suggest the public health took on a new residence and new relevance for how we think about its uses different from medical technologies that we have again introducing in public health for decades. Now the public was being asked to look at this, to think about it in a new way. So this moment that started in the fall of 2022, occasioned I think excitement on one hand, grave concern on the other. A lot of it documented by the media, some media that had been writing about AI and technology for a long time. Some in the business space. Some in the culture space. So we also got a lot of the media report in trying to create a narrative out of something fairly extract and nerdy.

So we were warned about rival camps in the AI community and existing harms and foreseeable risks. Some people called for course projection while others are spurring on AI environment. It is in this backdrop in which the public has been engaged that we think to think about today's conversation. I have tremendous optimism. Some of the things that we hear about AI, you know, are highest aspirations for it are about curing cancer or about as Dr. David Cutler, expanding radically access to care. It is important that we know that none of this is in the inevitable outcome of AI technology. Whether you are talking about deep learning or predictive technology, none of those tools inherently is going to lead to the outcomes, the beneficial

prosocial outcomes we want.

So we have got to steward their use to beneficial outcomes and policy is a key aspect of this and I'm proud to have worked in that space. Greg will have lots to say about that. I want to offer just four additional observations for our conversation about AI and public health. The first is don't believe the hype. The second, consider workers. Third consider patients and a little in the end about professional responsibility. So opportunities to expand health care I think will multiply if we are not caught up in the hype or only thinking about the technology as an accomplishing the social outcomes, public health outcomes that we want. It is the case that these technologies can help tremendously, Dr. David Cutler mentioned COVID-19. I have been -- I was curious to read and have been tracking research using AI and deep Neural Networks to analyze one based ultrasound images to see if you can diagnose COVID-19 early on.

I think that, you know, this again early days there is a lot to see what's going to happen with these tools but we could have imagined in the early days of the pandemic being able to try to use a tool like this to help us get more traction and understand what was happening at a moment of a high watermark of a pandemic. So that's immensely promising but there is much for to be done. We have examples in which I think we get caught up in the hype. 2016 the prominent scientist we should stop training radiologists right now. We are now in 2024. It was the case that by the end of 2022 we had more than 200 FDA approved radiology algorithms. There was a lot of technological change happening in the space of medical technologies to be used in public health. But we also at the same time had historic shortages of radiologists for training issues. And so I think that we can't just take prognostications about technology as being the outcome that we think might happen or fear is going to happen.

It is clear now that what is likely the case is that AI is going to be a tool to augment the work of radiologists. It will add value and save lives. Suggests that using AI to augment the work of radiologists might improve the performance of some but it might worsen the performance of others. Instead of coming out of the gate with new technologies and being caught up in the hype cycle, let's learn to speak with nuance about what these tools might do and possibly do because as I was trying to flag at the beginning, patients are listening. Then are paying attention now to the very first time how we talk about AI and this is shaping their expectations.

So with this study from nature medicine suggests to us is that we shouldn't look at radiologists as a uniform population. If we want to maximize benefits and minimize harm and use of assistive technologies, we have to think about the people. We have got to go back to things like human interaction. So

we -- don't believe the hype, but I think have optimism and I think -- but talk about these technologies and their potential in a much more nuanced and complicated way. And that doesn't mean that one is not optimistic about what AI might promise but it means that one is more forthright and honest and that matters.

Second, is to consider the workers. So I think one of the -- it wasn't implied I think in the framing questions that we got. Part of what's happening in public health is how the medical workforce is changing. President Biden did an Executive Order on artificial intelligence last October. There was a lot of callout for what HHS can do.

And it said that workers shouldn't be subject to undue surveillance. That there shouldn't be new safety risks or health references introduced into their work. We need to consider labor unions and employers and educators need to be in the conversation together. So I want to sort of raise up the public health workforce to be included in how we are thinking about this AI transition. And as we are thinking about the potential perils and possibilities of what it might do. There has been -- I'm thinking about Wall Street Journal reporting from June of 2023. That reported on nurses at UC Davis Medical Center and at Brigham Young women's hospital who felt paralyzed. They didn't feel like these were assistive technologies but technologies that were put in place to make them question their judgment.

And they also felt that they had implications with regards to liability and their own personal responsibility in being penalized if they felt like they made the wrong decision. If we think that these tools can be beneficial in one case one of the nurses was working in oncology, we might not have cured cancer but if we want to help people live longer lives we need to figure out what the relationship should be between nurses and AI assistive technologies and to take that very seriously and not just introduce them into the workplace.

Thirdly, prioritize and consider patients. I want to talk here some of my work has been writing about histories of mistrust and distrust and public health and medicine, from a perspective of communities of color and African American communities. So these new AI tools and systems entered this world for Isms. And we know that these -- the mistrust that already exists already has implication that we struggle with every day. When the introduction of new forms of AI to health care, and ways that, you know, AI had been in health care for a long time. But as I said at the beginning it is -- the public is well aware that AI is being introduced in every facet of their life. They are more aware now.

So we face the risk of a compounding of mistrust and skepticism. Because as we know from some recent polling the American public, you know, in some instances has mistrust of

medicine and starting to learn from the new polling that they are not sold on the benefits of AI. So in February, KPMG did some survey that suggested that most people are weary about trusting AI systems or have low moderate acceptance of AI. As I said at the top it depended on certain context. There is less trust for hiring if you are using it in resume screening or for hiring, the American public is less trustful. They are a bit more trustful for medicine. People recognize the benefits of AI but only half of the people who have been polled in this KPMG poll believe the benefits outweigh the risks. We have to think about how we do that without compounding and exacerbating the mistrust.

Just a little bit more data. The peer research center has done some polling around American's attitudes about AI. They found that six in ten Americans say they would feel uncomfortable with their care provider relying on AI to help care for them. Many people poll expressed that AI will take patient's personal relationships with their providers worse. And a general concern about general data privacy. I want to say a word about professional responsibility and ethics.

As Dr. David Cutler said generative AI systems and tools are brittle. They don't always work. He mentioned so-called hallucinations which is another way of saying fabrications. And, you know, it is the case that some of these tools are being released without what I think many of us would think is appropriate testing, safety testing for use in clinical care public health settings. So one I think can be reasonable and perfectly fine with the introduction of a generative AI tool for, you know, leisure time or to write a cover letter or give you a sonnet that sounds like Shakespeare based on your favorite ice cream or the fun things that we use chatbots for. When we are talking about clinical care, the bar has got to be higher. It has always been higher. It is why we have in health the only federal privacy law that we have which is HIPAA. We don't have a general law. We do have a law in health. Hopefully we will have a general law soon.

So I think that the introduction of these tools that are not tested and designed for the public health setting is confusing the question of who has a duty of care. It is turning over to the public health system to figure out how to test these tools and systems and that's fine what I think colleagues and organizations want to do.

So I have been tracking again since last month the emergence of coalitions, like trustworthy and responsible AI network which is engaging in developing evaluation standards for medical technologies and for the effective and responsibility applications of AI. I think it is an extraordinary effort. It is a partnership of many major health care providers in partnership with Microsoft. I think we want to pose the question like is it their job. Before the technology is released was it

the job of the particular company or the developer to make sure that those tools were responsible for this particular public health application. But again I think what the work of train is doing is extraordinary and important.

But I also want us to think about how the ecosystem about the duty of care who is responsible for making sure tools are safe and effective, who is responsible to the patient. Because I think getting a little bit scrambled in this new AI public health ecosystem I don't have any answers but I think I would love to talk about that and we are thinking about.

So in sum I think that advanced AI helps a wholesale change to the ecosystem of public health from the role of medical professionals, to the expectations of patients to the responsibility of workers to the very way that we think about professional ethics. Is it my responsibility and ethically responsible or culpable for the use of these tools at my workplace. I think that in addition to do this and Dr. David Cutler mentioned some of this we are seeing AI be introduced into how we do billing. There are things like electronic health records, clinical decision support, remote patient monitoring some of the work that FeFe Lee has been working on and robotics to lots of different people.

So I think to the bigger question of AI in public health, we need to be thinking about all of these ecosystems pieces and not forgetting the things we need to remember about how important it is to include workers to include the perspectives of patients. And lastly, to never assume that the tool alone, the system, the chatbot itself is the answer to the challenges that we face.

Thank you.

>> JENNIFER STRONG: A number of incredibly important points there as well as just very interesting conversation. I think we are going to have ahead. Thank you. Let's now turn things over to Mr. Greg Singleton.

>> GREG SINGLETON: All right. Good afternoon. Thank you. So glad to be joining you all. Thank you to Dean Sandro Galea, Dr. David Cutler, Dr. Nelson. So great to be interacting with you again. It is great to be here as in another link I have to BU in 2007 I got to sit in the audience and watch my wife go across the stage from BU to get her MBA from the School of Public Health. It is so great to engage with the community once again.

What I'm going to talk about, I want to talk about, you know, some of the dilemmas we face as looking at artificial intelligence in public health, I want to talk about how we think about artificial intelligence. Some of the things that we are doing as veteran government and some things we are excited about. Because, you know, ultimately we are looking at these tools to improve that Trifecta of care for the American public.

But, you know, going back to the theme of the talk, we caught between the promise and peril of artificial intelligence. Caught

between our imagination of what the March future of AI might be versus the possibilities of today or near tomorrow. One way to illustrate that, I am talking to my son and he is a teenager. He watches a fair bit of Marvel movies. Dad, you are working on AI for the health, for government. You know in all the movies the AI ends up to be evil, right? Hold on. And, you know, I think Dr. Nelson accurately said it the academic research reflecting some of those attitudes. There is this common perception that AI could do great harms. You know, I think one of the important messages we have, you know, out of the Biden Administration, out of the framework for Bill of Rights of AI, out of the Executive Order 1410 on safe trustworthy artificial intelligence, my role is to be on the field. As again Dr. Nelson said. There is no guarantee that AI causes either benefit or harm. It is up to us as society to figure out how we engage with these things and how we shape them for good.

But I do want to talk about why AI. If it is a sensitive or we are concerned about it. Why leverage AI at all. I ask people to look back where are we in society and where we are today versus where we were 20, 30 years. In the last 20, 30 years the volume of data of information of communication that's computed, stored, transmitted has just increased exponentially. And literally exponentially and it will continue to do so.

But when we look at our institutions our organizations and the workforce, that hasn't increased that much. That hasn't kept pace with that volume of stuff out there. So how do you handle that? And this is true all across the society. It is true in health and media and everywhere.

If you have the same workforce you have got more stuff out there. You either take longer to do it. You do a worse job to do it. Or you find improved methods and that's where we are turning to AI. We are turning to AI to find improved methods to find improved outcomes in this case health care. We are a little disserved as treating AI as one bucket of stuff. Because AI is in truth as Dr. David Cutler started off with talking about the different methods and the different uses. Dr. Nelson talked about we have to be specific about the use cases. The best definition I think of is they provide human insights at machine speeds.

Can we do them faster? Can we do them more efficiently? Dr. David Cutler said that electrons are cheaper than people or helpful to keep in mind as how a reason why we go forward with this.

So in practice, rather than being a bucket of artificial intelligence we use AI to get the right piece of information at the right time and put it in the right place. We use it to get information from a patient's medical record in front of the doctor or when it is important.

We use it to identify obscure research or help develop new compounds, use it to generate communications with patients. So

rather than just kind of -- it is all AI or one thing of AI, it is all these different techniques and the use case very much matters. We care a lot about the use cases when we design and leverage AI. Because the same basic technology that we can use to predict what movie you might want to watch, can help break the onset of disease and forecast the cost of care and recommend treatment options. In theory we can use AI to predict pretty much anything that we are able to collect data on. And that's what an image looks like, a tumor, whether a patient it will likely to develop a specific disease. Even a patient is going to show up for their next appointment. These are all things that would help advance care and health outcomes across the nation.

I want to talk about a few things we are particularly excited about at the department. You know, we are excited about AI enabled technologies in a number of near term areas. So if I think of that far future versus near term we are very interested in using AI to improve clinical outcomes. Knowing that we are working and augmenting human systems that also have errors in them already. We are excited about the opportunity to accelerate scientific break-throughs that increase treatment options, increase quality and quantity of life. We are interested and excited about reducing clinician burnout and allow greater time to be spent on care. Getting the information in front of the clinician they don't have to dig for it. Can help a lot with routine patient communications and engagement, reminder about your appointment or received your message about being late and getting back and forth with patients on that.

And then really important empowering patients through improved health and human services benefits delivery and plain language explanations for what is this and can be complicated human services and health care system that we have. Being able to translate what these things mean and make decisions for themselves.

So for the department priorities and how we are approaching this, is through the Executive Order on safe secure and trustworthy development and use of artificial intelligence that the administration has in October, that's following on the great work that Dr. Nelson did on framework for AI Bill of Rights we are prioritizing five activities over the common year. Working on developing policies to enable safe and responsible use and management of AI. Quality of AI in health care and processes. Working to leverage our grant making and contracting efforts to advance the development of responsible AI in health. And come up with the answers for how do we do this responsibility and how do we do it well. Working on public education, across the health care ecosystem to help constituents, doctors, patients, understand the engagement of AI in that system. And we are working to evaluate and deploy AI across our enterprise. Across the organization to help the government and HHS drive process, innovation and improve services and mission delivery for the

public. In taking together this this will enable the department to mobilize all the components of the department to improve our services and get people in the nation into a healthy space.

So with that, I just want to, you know, come back to I think the premise, the title, we are looking at the promise and peril of artificial intelligence and there is a lot of promise. But we have to be engaged in the discussion. We have to be engaged in doing it right and properly. And ensuring that we are taking care of the equities and working to adjust the isms that bias and coming up with tools and approaches to improve our health systems, our health organizations and the health outcomes for the nation.

So with that I will turn it back over to Jennifer. And look forward to the discussion.

>> JENNIFER STRONG: Thank you so much for that. Another fascinating way to start the consideration. I appreciated the near term look. Those are helpful to those of us who try to work about these things. Let's move on to our discussion now. Everyone, we are going to turn to audience questions when there is about 20 minutes left in the program. You can start submitting those any time use the Zoom QA function at the bottom of the screen. That's where I will be looking to see what you want us to talk about. So building off what everyone has said so far, AI already plays an important part in public health. So much more. But the list of reasons to be cautiously optimistic or just plain cautious pretty long. How are you interacting with AI in your most recent recent roles.

>> ALONDRA NELSON: Those were great presentations. So I am back in my academic role and here at the Institute for Advanced Study. One thing I have done is to start a multi-sector working group that we call the AI policy and governance working group and part of what we are trying to do is bring people together from industry, Civil Society and academia to think about what we agree upon. I sort of mentioned at the top the reporting about all the things that we don't agree upon and sort of what area's consensus from the expert practitioner and research community that we can sort of lean into. I think that's important. I think in the -- in the scrum of all of this is I think still a lot of confusion on the scientific side that we are still -- there is a lot of the science. Much like we have invoked COVID-19 a few times. Sort of all those preprints, draft papers there is still early production taking place. Working together to try to get a handle on that is really important. And then as part of the work of the working group we also do a lot of public engagements. So we do public events around that as well.

So that's -- and then I use some chatbots in my day-to-day work as a scholar. That's my engagement there. Trying to think about across sectors, certain use cases what are the guardrails, what are the new tools, pilot studies, research that we need to

be doing to advance the field.

>> JENNIFER STRONG: Dr. David Cutler.

>> DAVID CUTLER: Yes. So I will pick up on a couple of themes. One is partly in research, what I have been doing is I have been trying to talk with various health care organizations as they contemplate diving into AI for different things and then to try and understand how it goes. Implementation is not as easy as starting a patient on the drug where you go to the pharmacy and get the pill and take the pill. It is much more than that. One of the things that I can sympathize with health care, we are facing the exact same challenges about how do you use AI for beneficial things and how do you use it for not so beneficial things.

So the beneficial thing is gosh, you have motivated a student to be interested in a particular problem and now they can use AI to help learn about the problem and much easier than they could do it. And the pitfall if you are not careful they will start their term paper the night before it is due. And they haven't learned anything and it is just an easier way of cheating than it used to be. That's not -- we don't worry about exactly the same thing in medicine but we do worry about things like that. Like when is the AI just going to substitute for what you -- for what is the real hard work and the thing that only people can do. And how to combine the ChatGPT with the individual's knowledge so they can do well. And how do -- how much are we sort of saying go for it, do all that, do all this great stuff and how much are we there with the broom saying I better clean up this mess that you left behind and now you have given this problem. They have come up in the chats, about liability issues or workers who are displaced and you just got to bring your big broom behind. I have maybe a little bit more of a sense of how the struggles that are going on because I mean in a related industry to the medical care one.

>> JENNIFER STRONG: Makes a lot of sense. Mr. Greg Singleton, did you want to chime in?

>> GREG SINGLETON: Sure. Yeah. Looking at how we are gaging -- engaging on AI and focused on. It is quite a lot as a department. I think at the top level, it is kind of iterative, deliberate, risk managed approach to AI looking at the challenges in front of us today. Address those and learn from them and kind of grow. But overall our activities are guided by the Executive Order 14110 on safe and trustworthy use of artificial intelligence where across the administration they assign some 150 deliverables and assignments out to different departments. We as a department were fortunate enough to have I believe 16 assignments come our way. As we are really focused on those, in delivering those over the coming year.

And those are things like developing strategy and framework for leveraging AI in public health and health benefit systems. Looking at quality assurance plans, how do we integrate and

think about AI when we look at, you know, drugs or device development. We are looking at elements of art first intelligence and enhancing or improving nondiscrimination and nonbiased aspects of the law. So they still apply when we look at artificial intelligence. We are broadly looking at our department as a strategy, both internally and over the common year developing our approach to helping shape these sectors approach and sectors engagement of AI.

So we are doing a lot through that strategy. Overall working to understand what's the sector doing now. What's the sector focusing on tomorrow. And then engaging with the policy community so we can help shape and steer that in productive ways. Quite a lot. But again the short version is, you know, what's in front of us deliberate, cautious judicial engagement and continuing learning.

>> JENNIFER STRONG: Sticking with generative AI is having its moment. How do you see it being used responsibly for public health and that question is for anyone who wants to take it. Or perhaps --

>> GREG SINGLETON: I can take a crack at that. I mean one of the -- one key of the areas of public health is engaging with the patients and population. And helping them feel empowered and confident in their -- in their care. In their level of knowledge, in their being part of the discussion about how they are feeling. There are a number of areas where, you know, at least, even pretty soon in the near term we can leverage generative AI to translate documents, to engage increased populations, to provide information to people in an accessible form. There are the opportunities to structure generative AI to digest and provide information to people in forms that they can understand it. They can get access to it and provide kind of steering to people so they can engage with the system. In maybe a little further out there are opportunities for the AI to be a patient navigator. There are a lot of opportunities here. The question as always is how do we do this safely and responsibly and cautiously and proceed in a step wise manner so that we are managing risk.

>> DAVID CUTLER: I was part of a survey. We recently did a survey on uses of AI, potential uses of AI and barriers to AI use in health care.

The group in health care that's most eager to use AI is insurers, payers. And they want to use it to streamline some of their things like prior approval processes. So they get a request from a doc to do something and they want to have the AI look at it and decide what to do instead of always having a human look at it. That's sort of the first use. They are gung ho about that and prepared to do that and finding the right systems.

And on the medical provider end there is a lot more skittishness. So they would like the AI to be helpful. They are

worried about the things my colleagues were saying, confidentiality, making up data and patients and reaction to it. And so what I think the provider community is sort of looking for some like, if you will, guidelines about here's how you should do this. Here's the safe things to do. Here is how you should interact with it. There should be some leading edge. There are a lot of folks waiting for either the Federal Government or someone to come on and say -- we believe this is fine and you should come ahead and do this.

>> JENNIFER STRONG: Yeah.

>> ALONDRA NELSON: I don't -- I don't know that we know yet that generative AI if you are talking about large language models and talking about the use of multimodal, you know, foundation models, can be used responsibly in health care actually. I don't -- I think that's an open empirical question and right now the answer may be no. I think we have to see both in the New England Journal of Medicine and also in Jama articles that shows there was racial bias and large language models that were used to ask questions.

Greg is right, that translation is probably the best case scenario, but I think even sort of ingesting large documents, you know, we face the challenge of on the one hand you are trying to solve the problem of what do I do with all these records, tens of thousands of pages but we got the challenge of the fabrication. You have not been able to read the 20,000 pages or whatever it is distilled. So we find ourselves -- I think -- I do think there is -- there needs to be a higher duty of care and a higher barrier in health care. It is just too important and the stakes are high. If you are thinking about physicians, how do we think about using a tool that we know can be inaccurate. I don't know how as a physician you are looking for better outcomes. You hope to make improvement for your patients. But you also know that the tools don't always work or work at the level that you would feel comfortable with.

And so I do think again not providing any answers, it poses for medical professionals a real ethical quandary, particularly for generative AI. If you are talking about deep learning or Neural Networks or things that are -- that are not intended to generate new things as part of the work they do, you know, you are talking about radiological imaging or imaging of lung diagnoses, or COVID-19 maybe that's a little bit different. Generative AI is a struggle for public.

>> JENNIFER STRONG: Sure. I have heard plenty of interesting things using generative won drug discovery. I live in New York City and there was a very carefully trained chatbot that was providing advice that's against the law. It is hard for me to imagine, I wanted to ask if people were working on things or knew of things.

Getting back to those guidelines that Dr. David Cutler just raised. As is often the case technology is vastly outpacing the

frameworks that might shape development. How concerned should the public be about that right now? How concerned are you about that right now?

>> GREG SINGLETON: It is something that we are keeping a close eye on as we look at the evolution of these systems and look at the evolution of practices in the sector.

And I think there are a few sides to that question. One is certainly the intentional legitimate use of AI in delivery of care and provision of trustworthy services to clients. There is also the possibility for, you know, others are outside actors to use generative AI in harmful ways to cause problems that we are, you know, also taking a look at.

It is something that right now with how the tools are being used again, as people are looking and focusing on these specific use cases, what are they doing with the AI, what's the risk if it goes right. What's the risk if it goes wrong. But then also how does this compare to the existing system. There are systems that okay, if it gets a citation wrong or it hallucinates one element of a conversation, but your existing clinician might have a higher error rate than that. We have seen some studies where since the electrons are free in chatting with patients, some of these chatbots can be more empathetic and leave the patients with a greater feeling of care and engagement than you can from a normal clinician. You have to be comparing to the existing system that we have that is not perfect and looking at the existing system as if it is perfect and flawless. So I think the perspective is do these tools right now and by the way these are kind of Gen 0 tools. There are a lot of trials going on trying to work on that.

If these systems can be used to improve the overall system we have today, great, let's go forward with it. If we are not able to get them to work, if we go through these, you know, frameworks training system, the group, the triad group and we find we can't use them comfortably and reliably then no, it is not something that we can proceed with. It is part of our everyday discussions with the technology sector, with the medical sector and clinicians and the public.

>> DAVID CUTLER: In my mind a very good guide is to follow the money. Whichever way the money goes, think about that. So like these things will not just appear randomly. The financial use case will be high on the list, hopefully will have at least some resources to think about clinical and nonclinical use cases where there is not much money. But where we could do a lot of good. I would look a lot at what kind of incentives get put in place and I think folks are still trying to work their way through that. Helping that I think will really help shape which way where we see the biggest issues and both possibilities and pitfalls with AI.

>> JENNIFER STRONG: Yeah.

>> ALONDRA NELSON: Can I follow up on that? Part of the

work that I was doing when I was working as a science advisor in the Biden Administration was helping to usher through ARPH which is this endeavor. And the incentive structure piece is so important. Public health, there are lots of things that we want medical technologies to do that might not necessarily -- effectively market failures. They are not going to make money for a company but they might do a lot of good for a lot of people. And so, you know, I'm glad that David raised that because I think it is that particular incentive structures for these particular technologies that still remains to be thought out. We are living with a major technological transformation driven almost solely by the private sector. So we have got Mariana and others have written about the investment of Federal Government through DARPA and other agencies that got us to GPS and iPhone.

And so the instinctive structure is vastly different. It is, you know, you have got to keep companies in business. It is a different proposition for something like public health. There is a misalignment that it is -- it is fine. We have got a lot of market misalignments, but it is worth thinking about that specifically. You asked about keeping up with the technology. I wanted to raise two other points. About this technological transformation being driven by the private sector or by industry which is fantastic. Innovation at its best is fantastic. But it also -- but these are also disruptive models. In some way the whole purpose, many models in Silicon Valley are disruptive.

If you think about Uber, it was getting around the taxi medallion system. But it is also the case that the industries are set up explicitly to get around regulation. So you can't be surprised and say they are behind with regulation when the whole sort of raise on a particular business model or strategy is disruption.

The other thing I would remind colleagues, people of here is as a former government employee that wanted to defend the great work, there has been legislation to help us put guardrails exactly like we are talking about for years in Washington. There have been people advocating on the hill, staffers, legislates for data privacy. Longer. Like in this more immediate kind of AI conversation. And so it is not that legislators don't actually know what are the kinds of regulations that need to help to keep up with the technology. It is like we don't have the political will to get these things done. There have been very good bills around algorithmic accountability, around data privacy for several, several years now. That would have anticipated and could have been in place well before generative AI emerged in the fall of 2022. I want to keep that piece in mind.

>> JENNIFER STRONG: Absolutely. There is another element, AI -- I often say we needed to drive before we decided that seat belts was a great idea. AI has a potential to help

break down tremendous barriers and help many more people gain access to care. Also has the potential to make this process less fair than it has been. Dr. Nelson, with your work, are you -- you have specific suggestions for how we get that first outcome. Things that people can begin working on now?

>> ALONDRA NELSON: We have got to think about the full technology stack. Thinking about these issues around risk and harm, I mean the way we get to benefits is by mitigating the risks and harms. And we need to be thinking about them literally in the design and development of technologies. And to anticipate that they are likely to happen. So training data. Whether or not it is scraped with claims of violate IP or not. We will leave it to other people to litigate that. But datasets come with various kinds of constraints. They have been demonstrated to have biases around race and gender and other issues that matter for public health in ways that they might not.

The same training dataset used for a recommend your algorithm on social media, like maybe we don't care. But I think when you are dealing with the stakes of public health we have to ask those questions. How are we thinking about whether and how we are mitigating bias. Are tools being risk assessed or auditing them before they are released to make sure can't foresee everything but good faith due diligence been done. And then after they are released, particularly when you are talking about more general purpose AI that can have a multiplicity of use cases, then you have got to dig out and sort of be watching to make sure that there is not risk and harm in all those different use cases.

We are talking about what we have large language models, things in public health and education. The risk calculates for all of those is completely different. But I think we have to move out of a space of, you know, like about the FDA, this is a medical device and we have approved it for one time for a dynamic and nimble. We are keeping more AI systems.

>> JENNIFER STRONG: One more quick question though. I don't know how it is quick. Let's talk about jobs. Kind of frame that as quick. But many of us here are familiar with the example that Dr. Nelson gave earlier. They were predicted to be largely replaced. How do we deal with such uncertainty? And what should health systems do to prepare given this?

>> DAVID CUTLER: On the clinical side, AI technology will compliment and not substitute for clinical workers. You are going to need the clinical for liability reasons because the doctor knows more things than the computer knows because the doctor has seen particular circumstances or can recognize particular constraints. And now as I think what Dr. Nelson was talking about in terms of the recent studies, the computer and doctor bring different things to the table. I'm not worried about mass unemployment. There may be more substitution. And I do think that there is some administrative personnel who may be

duplicative of things that AI can do. I also think that the administrative side are areas where we need a lot more people to be doing stuff other than just administration. So we need more outreach to people. We need more people to monitor chronic disease stuff. And a lot of people who are involved in medical records, coding and billing have skills that would be quite transferrable to where we need people but we haven't been able to have them and stuff.

So I think that we are not going to have mass unemployment of people because the phase-ins are going to be gradual and so on. We could have a redirection of some of the administrative work that's going into work that's more patient facing and focused more on how do you deal with the patient in need so that you are not there answering the telephone which doctor should you talk to or how should you submit this bill. These are what the history of the test results are. So let's -- let's figure out clinically what to do. Make sure that things are going the way we should or the plan that we have in place is working appropriately.

>> JENNIFER STRONG: That's something I have been waiting to see more broadly in the economy. Instead of chasing that idea that we have had for five years. We have budget cutting. Wouldn't it be nice. A lot of people are offering this variation on a similar one. I'm going to mesh them together and offer this one up.

How is an AI system created to be without Isms when the creators have their own set of Isms?

>> GREG SINGLETON: I will offer a perspective and that's to say we are at the early phase of these technologies. We have goals in sight and the goals that are laid out in the Executive Order fair, trustworthy, safe, artificial intelligence. And what we are working through and developing are the processes to get there. And it will take experimentation. It will take learning. It will take development communication. I think a little bit like baking a cake. You have ingredients. You have your process and outcome. And we are working across those dimensions of data, data reliability, data collection. We are working on model develop processes, other things. We are working on the measurement and understanding the outcomes and how they match to our goals.

But ultimately we are working through that process with those goals in mind as in ways that maybe we aren't with the existing systems or existing organizations aren't able to fulfill. So being able to measure things through AI allows us to constantly work to improve. With the goal of getting to AI that reflects American values.

>> JENNIFER STRONG: Anybody else want to jump in before I go to the next?

>> DAVID CUTLER: I will add one thing, outside of government health care is a very sort of prove it kind of

industry. You don't get to just make a claim. Your claim has to be published and done in a clinical trial and so on. So I think that's going to help us with this. Because I think clinicians are rightly skeptical about someone coming in and announcing I can replace this task that you have been doing with this program. Then you are going to generate the evidence for that and stuff. So my hope is that as a community we can hold the ground firm in terms of these when we expect. We expect it to be accurate and we expect it not to hallucinate. If we find that we are going to fix that. And I think the Federal Government with its recent report has laid some good goals that we can try and translate into how to make that happen there can be kind of a standard setting. A little bit like the FDA. This trial did this and what that trial did that. Sometimes you get it wrong but at least the goal is to make a scientific determination about what's right and wrong.

>> ALONDRA NELSON: The challenge we face with AI and generative AI in particular it is dynamic. You can't do it once. The FDA model has often been you approve the medical device once and off you go. We need to figure out an ecosystem that allows for the empirical science based checks throughout that whole stack. Throughout that whole lifecycle which is the design, development and deployment. So to the question that you synthesize from all the other questions coming in, you know, any designer of any tool is going to bring their own sort of isms. Other sort of perspective to the work. It is a case that we can do a better job before we release these tools knowing that's the case. We can be -- we might think about the use of vision based AI. Trying to use, you know, sort of AI powered, you know, faucets to watch their hands. Because they had not been tested. That's just a basic design failure that the companies -- I mean, you know, we should expect much more than that from products. I do think that there needs to be to David's point a real transformation and the responsibility that I think the designers, developers, employers feel is theirs. And thank goodness we have Greg and others working in this new important role. Putting a check on that.

>> JENNIFER STRONG: Uh-huh. Another question from the chat, we have who should be held accountable in cases of errors resulting in the use of AI in health care, particularly in scenarios involving autonomous decision making? Anybody want to take a stab?

>> ALONDRA NELSON: Go ahead, Greg.

>> GREG SINGLETON: No. Go first.

>> ALONDRA NELSON: This is your job.

>> GREG SINGLETON: I'm hesitant to say who it should be. I will observe there are very few cases in the economy where we absolve folks from liability for things that they participate in. And that the, you know, responsibility tends to lie in staying with designers, developers. It is a shared

responsibility chain based on roles and, you know, I don't expect that principle of U.S. law to really change too much there.

>> JENNIFER STRONG: That was my question about autonomous decision making and how we define that.

>> ALONDRA NELSON: The only thing I would add is someone who can speak more freely than Greg, released from the strictures of government. The challenge we face we have no visibility into them. If something goes wrong on something that's built on top of a foundation model, how do we think of third part actors in this space. We have got a lot of fresh unanswered questions still to be answered around liability. We had similar ones with predictive algorithms. Some of the ones that are -- were used more generally came out of the meteorology or seismology, the large chaotic systems. And so you could sort of track those a little bit more, but you still didn't -- we often just don't know what's inside the black box. And I think it makes the liability question a bit more different than it does with some other medical technologies.

>> JENNIFER STRONG: Scrolling down here, we have folks who would love to hear the panelist's thoughts on the regulation used to train the models. I was thinking about Colorado if anyone saw it in the New York Times on brain data. If anybody has something they would like to share by all means take the floor.

>> DAVID CUTLER: Just to start off, I want to come back to one of the issues that we came by briefly. We wanted data to be trained on what is the truth, not what is the physician's perception of the truth. Because the physician's perception of the truth has a lot more stuff built into it that may not be the actual truth. I think about doing things on the basis of ground truth. So, for example, predicting who is going to have a cardiac event. I would rather base it on who has a cardiac event not on -- not can I reproduce the physician's perception of it. I want to be cautious, sometimes the ground truth itself can be biased. So, for example, you know, there is the example that was written up in an article in science about a company that was trying to reduce readmission rates to the hospital or post acute spending. Want to reduce post acute spending.

So the algorithm -- they wrote an algorithm to try and predict. The algorithm figured out that white patients use more care than black patients. Therefore you should clearly intervene if you are interested in saving money in follow-up care, intervene relatively more for white patients than black patients.

But it is still in that case the actual honest to goodness truth was that white patients use more care. Not because they needed it because that's the way the -- that's the way the system was. So there I really want to think about the truth as not being I want to save money. I want to improve the health of

patients. Then I will think about the various health problems. Not the physicians telling me what the health problems are but the actual health of the individual. It is going to be very difficult, but my brain keeps coming to this, what is the honest to goodness truth about the world. Not what is it that someone is reporting to be true.

>> JENNIFER STRONG: And how you capture that, I think about this in terms of women as well or with synthetic data we can have an entire conversation just about that. We have one here that says question to Dr. David Cutler. You said you will be happy to see people do bills, outcome of the financial, medical, paramedical, the universal health institutional level?

>> DAVID CUTLER: Will likely replace some people over time. I think and I believe that there are plenty of things that people need to be doing in health care. Where people with skills like those that are doing medical report keeping and coding can work that would be extremely valuable for folks.

And having them -- having very, very dedicated loyal workers doing that is I think a waste of many people's talents relative to what they could be doing.

>> ALONDRA NELSON: Of course, we can also -- there is probably a litany of jobs that we can't even imagine I think that will be produced out of this transition. So, you know, we can think about now the sort of emergence of prompt engineering as a kind of role. If decisions are made in large health centers to use large language models with the understanding that they have got, you know, a challenge with (inaudible), sometimes accurate information, then who are going to be the people that sort of help to suss that out in the system? If we stick with the systems and we introduce them with public health, there is going to have to be lots of different guardrails. You could imagine sort of medical prompt engineering being a whole new kind of role that takes a particular kind of expertise because you would want people with subject matter expertise to prompt the tools for the things that would be prompted in a public health space as opposed to general conversation for the broad public.

So and we saw this, of course, also with automated cars. Or getting rid of all cars and drivers. No. There were humans helping these automated cars be on the road. I think we can -- we shouldn't only predict that jobs will be lost as a function of this transition. And I would also point back to President Biden's Executive Order on work in particular that really asked the Department of Labor and others to look at ways in which, you know, job loss could be mitigated.

Like how can the Federal Government be supporting rescaling, upscaling and sort of creating a job ecosystem that's not just tremendous job loss.

>> JENNIFER STRONG: A number of people on this call would like to know if you have any favorite resources, if they are looking to further educate themselves in the space of AI

within public health. There is also a number of people wanting to talk about recommender systems and predictive analytics, ones that could be used to reduce cultural disparities access to or levels of care.

>> ALONDRA NELSON: I would offer the journal of medical association has done a great job. There is a whole series that started last year. There is videos and interviews and the like. I participated in that. That's a good resource.

>> DAVID CUTLER: NGEN has started a NGEN AI which I find to be particularly helpful as well.

>> JENNIFER STRONG: Interesting comment here as well, from a woman in South Africa who says on the one hand we have challenges with electricity supply in our health care facilities, but we are using AI to assist with our TBI diagnosis. If anyone.

>> ALONDRA NELSON: I don't know. Greg, you might know, you might be tracking this closer. I think that train is American only. I imagine many of these health systems are multi-national.

>> GREG SINGLETON: It is early days for these efforts. There is a lot of work being done I think with OECD, with the European Union, with the UN around kind of AI development. I will say I know that we also have programs at NIH and others that work to engage diverse institutions, diverse researchers and help facilitate kind of that engagement with broader communities. I know we have programs there. I don't know on the specific question on the African based institutions.

>> DAVID CUTLER: One of my hopes without answering on those specific institutions, one of my hopes that will evolve a system for AI like we have for essential medicines. Which is that we -- the world basically agrees we will sell them in the poorest countries for cost. We will give them away. As long as someone will pay the cost of physically producing the drug and getting it there. We will just give it away. The world is not going to make money off selling drugs, excellent drugs in poor countries. We should have the same view about diagnosis, about AI operations, TB testing, all the rest of it in low income countries that that's not a market that we want to make money in. We want folks to -- we just got to cover the cost of getting it from there. If you are in a higher income country, fine, you have to pay for it.

>> JENNIFER STRONG: We are at time. Thank you, everyone, for joining us. Thank you to our esteemed panelists. This was my great pleasure to moderate this discussion today. Very much appreciate your time.

>> SANDRO GALEA: First of all, thank you, Jennifer, for your outstanding moderation. Thank you to Alondra, Greg and David for an interesting conversation. Thank you to the audience. I thought the stream of questions and interesting topics in the chat were great. I feel like I leave this

conversation both much better informed and also have more questions than I started, which is how it should be for a topic like this. Thank you for engaging with us. Thank you to our panelists and thank you to the moderator and everyone in the audience who engaged in this topic. Have a great afternoon. Take care.

(Session concluded at 1:28 p.m. CT)

This text, document, or file is based on live transcription. Communication Access Realtime Translation (CART), captioning, and/or live transcription are provided in order to facilitate communication accessibility and may not be a totally verbatim record of the proceedings. This text, document, or file is not to be distributed or used in any way that may violate copyright law.
