

RAW FILE

ITU

AI FOR GOOD GLOBAL SUMMIT

GENEVA, SWITZERLAND

DAY 2, 16 MAY 2018

11:00 CET

FOUR AI BREAKTHROUGH TEAMS - THE HEART OF THE SUMMIT

ROOM K

BUILDING TRUST FOR BENEFICIAL AI - DEVELOPER COMMUNITIES

Services provided by:
Caption First, Inc.
P.O. Box 3066
Monument, CO 80132
1-877-825-5234
+001-719-481-9835
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.

>> Right. Good morning, everybody. Good morning. I am glad to hear there's lots of conversations going on, but we are on to the second session. So I am going to keep going. Nice to meet you all. My name is Claire Craig. I am the Director of Science Policy at the Royal Society, which is the British UK national science academy, and I am also a member of the Trust in AI team with the Centre for the Future of Intelligence and others.

So this is the second session. It's B. And this session is about trust and trustworthiness across the boundaries of nations and of cultures. And it matters hugely because, of course, data and applications are not inherently going to observe boundaries, but the perceptions, the stories, and the world views that surround them are often deeply culturally specific and will inform the development and the evolution of the science and the technologies. And we are not technological determinists here. We believe these things go hand in hand.

In order to ensure we get systems that are trustworthy and to have confidence, we need to understand better what these cultural

differences are and what they mean for building trust, and that's the purpose of the projects in this session.

The structure will be the same as the previous one, so we are going to hear about three projects, and each will be introduced by one or more members of their team, and then we'll have time to discuss each, around half an hour for each session, and then we'll break for lunch.

So I'd like to begin by inviting Professor Liu Zhe, a Professor of Philosophy at Peking University, to talk about his project.

(Applause)

>> LIU ZHE: Okay. It's my great honor to be here with you, and I am actually coming here for your help, not for presenting my own ideas. And I really thank Professor Huw Price from Cambridge for this invitation to join this fantastic team and also join in this track, Trust in AI.

But today I am going to talk about something, not merely on AI, but also robots, because I am currently directing the Chinese national white paper for the robot ICT standardization. But today my voice is not representing this government voice, just my own philosophical reflections. So I just contribute something which I can understand with respect to this emerging technology.

Before I start, actually, I would like to stress one point. In the first session we had today, I think the panelists really talked about AI systems, but actually, last week when we had the first review of our document, I did see a kind of serious debate on the relation

or the differentiation between AI systems and robots. So the people seems not really great with the definition of robots, and I also heard ISO try to give a new definition. So today when I talk about this, the beneficial AI, I also talk about beneficial robots, especially social robots. This is very challenging issue for understanding of the problem of trust, trustworthy and autonomy.

The second explanation I would like to make here is I secretly actually bring (?) from Cambridge, her theory on trustworthy autonomy to our discussion, because I know eventually she somehow will, at a certain point, will join this team as well, so I would like to -- it's not a challenge, but try to set a certain platform where people can really understand the differentiation between trust and trustworthy issues in biotechnology and also the same issue in AI and also robotics. This is my ...

The issue of trust in AI must be explained in the context of existing technologies and foreseeable programs in the near future. Actually, I really (?) yesterday is really dangerous to base our discussion and reflection on the basis of fiction and also movies. Here are some technological developments in which so far, as I see, we may have certain consensus. The first thing is the machine learning with Big Data. Yesterday we heard lots about Big Data. That is actually the field of AI technology. The second thing, biomet Ronics, we heard yesterday on the notion of hybrid. No one will deny that is artificial, but that is intelligence. The third thing, the inincrease of automation or autonomy. Here I put the

quotation marks here "autonomy," because autonomy very freely used in cybernetics and in the technological circle, which is really different from the psychological use.

The final one, closer interaction between humans and AI or robots, here one can see the digital twins, cybots, and so on. This is the coming technology right now, scientists and technologists would really like to pursue.

And I would like to particularly stress the importance of increased autonomy in AI and robots. There is converging evidence that the degree of autonomy that a robot exhibits is an important factor in determining the extent to which it will be viewed as human-like. So this is a very interesting empirical observation.

On the basis of this technological views, we can now make a first step to scope the problems of trust in artificial intelligence and also robots. In China, I haven't seen very prevalent distrust or mistrust of artificial intelligence or robots in the entire range of society. People seem very excited and enthusiastic for the emerging technologies. In the truly earliest accidents of Tesla, self-driving car in January 2016, actually, the first reported case in Florida in the United States, it was May, also same year. But actually, the genuine earliest accident was in China. Which is much earlier than the reported one in Florida. The father of this accident said that we have put too much trust -- that's what he used the word -- trust. I understand perhaps this is too much reliance. We will talk about the distinction -- reliance on the emerging

technology, such as auto-pilot system of Tesla. It seems to me that the interview of his father represents a major idea of the Chinese public for artificial intelligence and robots.

In brief, it is misplaced trust or overtrust that really involves certain deceptions in the interaction between humans and AI and robots.

Like in the area of technologies, the laws of trust or mistrust or distrust is often subscribed to the untrustworthiness of (?) and of those holders of public office which may legislate for and regulate their activities. It is often suggested they are keen to pursue their own interests rather than the interests of the public.

Although we have not yet seen any case to demonstrate the consequence and effect of such distrust of holders of office of trust in China, we might suppose that distrust will eventually harm public trust in AI and robots like the topic which actually we design. I understand the background of this topic.

Here, very quickly, some recent issues of personal data breaching by Facebook coming to my mind. I think about these things.

However, the loss of trust or mistrust in AI or robots differs from the misplaced trust and overtrust. Think about the two risks, very different from each other.

I correct my slides a little bit, so it's not exactly -- yeah. No.

Yes. Deception in interaction between humans and AI robots. Here I talk about a couple of issues. First, how to understand the

relation of HRI and difference from the human interactions. Trust, the notion of trust can be differently understood from institutional trust to the trust itself, and also to the trust in government. Now we are only considering trust in the form of interpersonal relations. Another question. Can we really understand human-robot interaction, HRI, or human artificial intelligence relation as a kind of interpersonal relation at all? This is really an issue which we need to think about.

Before any possible answer, let's see some human interaction with robots. This is not from China, but from western countries, from United States, reports made by MIT and also Stanford. The picture I like to show is not from narrow Chinese perspective, but rather the global impact. Patients with ASD tend to have troubles with range of specific social -- ASD -- such as understanding and responding appropriately to facial expressions and social contexts, imitating others' behavior, engaging in cooperative activities such as taking turns, asking questions, and following conversation cues. Some studies have found that patients with ASD actively prefer interacting with robots to interacting with human companions. Other research has also found that patients respond better to smiles and other facial expressions, focus for longer and otherwise engage socially at greater rates with robots than with humans. A worry is then raised that HRI will really be able to replace interpersonal relations for these patients. Is that really beneficial?

Users of healthcare robots, a project traced characteristics onto

robots that they don't genuinely possess. Humans can form very strong emotional ties to robots and other technological artifacts. In the case of artificial friendship or artificial care, the involved deception doesn't concern the incapability of artificial intelligence or robot to do what they are meant to do, but rather, their better performance -- much better performance -- than human beings themselves. It seems that people have their misplaced trust in artificial intelligence and robots in a way that HRI relation is conflated with interpersonal relation. People not only are inclined to misplace the trust, but also to overtrust AI and robots. On June 1, 2009, Air France flight 447 crashed into the ocean, killing all 228 passengers on board. Investigators eventually concluded that the crew's confusion after disengaging the auto pilot and reliance on faulty airspeed measurements doomed the plane. So this is a case of overtrust. Actually, some studies show that people tend to place greater trust in computers, at least as sources of information.

Clearly, the possibility of deception involved in both misplaced trust and overtrust of AI and robots can only as a consequence result from our trust in AI and robots. So it's very different from the mistrust or distrust, but rather on the basis of trust. It might be argued that one should make a distinction between trust in and reliance on AI and robots. But when I think about this issue, I don't see. The conceptual distinction could be made, but I don't see how such distinction can really work here because humans are psychologically inclined to a unilateral dependence on AI and robots.

In east Asia, like China and Japan -- yesterday I also talked to our Japanese colleague -- the popular culture is in favor of artificial intelligence and robots as our companions from the very outset. So it's quite different from the Western countries. I know even the robots, like a forced labor, so there is kind of very prevalent Ideologies like robots, some slave against the master. And then like the figure with increasing capabilities eventually come to be genuine human. But these never exist in China nor, I think, in Japan. So we have the Japanese cartoon when I was a kid, so it's a cartoon Astroboy, so that shows a very friendly relation between humans and also the robots. So here in these kind of societies, I think the overtrust, the risk of overtrust and the misplaced trust is really very high. So I think we need to address this issue today when we think about the trust in AI, whether trust in AI is really good or beneficial for our relation to artificial intelligence, also robots.

But when we try to reflect on this issue of misplaced trust or overtrust of artificial intelligence, it might be good to introduce certain (?) reflection on trust, institutional trust, and also concept of trust, which actually Professor O'Neill made with respect to biotechnology. Both technologies have something right now emerging with each other, but on the other hand, they are both cutting-edge technologies. So I just very briefly introduce this distinction. Professor O'Neill introduced this valuable distinction between individual autonomy and principal autonomy.

Individual autonomy in brief is autonomy in natural picture, so that shows independence in the causal chain of natural states and events, like desires and beliefs. And the principal autonomy, basically borrows from the (?) distinctive formal presence or self-expression, but in the sense of self-legislation or nonderivative legislation. So you never see the actual use this term self, but rather the (?) -- the term which is used. So we can reformulate the moral law, we must act on principles others can follow. This is what actually Professor O'Neill reformulates the categorical imperative.

But if we take a closer look at this distinction between individual autonomy and also principal principle autonomy, both kinds of autonomy, notions of modern values, must be based on we are personal relation. I am just wondering whether this kind of distinction or even the framework of autonomy can be used in the context of HRI or in the context of human relation interaction to artificial intelligence. This is due to mistrust or overtrust in AI and robots. It seems trust can neither be grounded in AI nor principle autonomy. Now the question we have to raise here, whether the concept of trust, really appropriate framework to help us conceptualize our relation and beneficial relation or rational relation to artificial intelligence and robots. If not, what is the alternative? This is my question. When I am thinking through this very messy situation, I see the only thing as a philosopher, what I can contribute is the recent question. Right?

And the rest of my talk -- I have a couple minutes? Just two or

three minutes, because this is really not new, and also not outdated. In China, this is kind of observation. Before closing my talk, I would like to introduce some of my observations with respect to the problem of trustworthy. Today sociologists introduce the notion of risk societies. This notion is not a claim about the levels of risk but about changes in perception of risks, or at least reported perception of risk, so how to view this kind of perception of risk, it depends on the cultural background. It depends on lots of things, not just the technologies themselves.

The current trend in technology or technologically advanced societies, the Chinese public is more and more worried about the risk that the emerging technologies are bringing along with their development and application. Their worries are understood as deep suspect of trustworthy of companies and politicians. Here I have heard some arguments from Chinese artificial intelligence scientists and roboticists to criticize.

For instance, humanity people don't really have sufficient knowledge of technologies because in college, high school, the humanity and natural science part divide. So the second year of high school students must select either one. So that creates trouble.

The second thing, the traditional dichotomy of fact and value, so the scientists claim technology is just value neutral, so how you should blame that? Really, the one should be blamed must be the user, the scientist themselves. They use in incorrect way. So this is kind of the argument which I found.

Another thing is about overseeing -- the overseeing of the new technology. I am Writing this national policy, I also proposed to the government a certain idea of how to do this governance, not so governmental, but governance of the new technologies. But basically, I see the general picture still prevalent, even in China's public, they think there must be a kind of paternalism that should be held down in this kind of thing. So the paternalist mind-set, not merely among the Chinese public, but also AI scientists and roboticists, they think if the public interferes in these issues, the development of new technologies, but -- it will block the development of new technologies. But here still, not just like some observations from the outside, they think it's the Chinese government that doesn't take up too much responsibility for this public issue, but it's just the opposite. So sometimes it's quite complicated picture which I would like to introduce. I think the Chinese public also needs a kind of more education to understand the risk and also the challenging issues in this new technologies.

Now, okay, let's very quickly finish my talk. So here I have some open questions. First, can white lies be allowed to artificial intelligence and robots? Either the robots itself can tell white lie, or as the counterfeit coin, using the human relation to artificial intelligence?

Second, will now the increase of trustworthy endanger the trust in interaction between humans and artificial intelligence, especially when you think about the misplaced trust or overtrust?

Third, is trustworthy of artificial intelligence or robots a sufficient necessary condition of trust between humans and artificial intelligence?

And finally, I didn't list here, shall we insist on trust as an appropriate framework to conceptualize our beneficial relation to artificial intelligence?

Thank you very much.

(Applause)

>> CLAIRE CRAIG: Thank you very much, indeed, and the questions of how the technologists and scientists work with the scholars and philosophers and others is one we are certainly very much involved with in this Trust in AI track.

There's a couple of really interesting questions on the app. I shall start with one of those, and we've just got a few minutes for questions if we are going to keep to time.

There is a question about because the issue of trust and overtrust and misplaced trust has come up several times, a question about how trust can be measured or assessed. In other words, if we have the concept of overtrust, how would we know when we were getting to that point?

>> LIU ZHE: Yeah. Can you hear me? Okay. You know, actually, from my point of view, before any (?) of trust, we should understand the relation between human beings and artificial intelligence. It is not clear whether we can use the trust, this concept at all, in

this. This is my worry. Especially when you think about overtrust, there's certain competence from the technologies, but we might have reliance on technology, but maybe it's wrong. I also have certain argument from the opposite side, why should we trust? Why not just introduce some distrust of the system? And then we can have better life.

>> CLAIRE CRAIG: Okay. Thank you. Super. So the lady over there, please, on the second row.

>> So I have been working on AI systems for a long time. I am a Professor of electronic and computer engineering in Hong Kong University of Science and Technology, and I heard a lot about the arguments about trusting machines. Right? And I think today, all the AI systems have performance metrics that are measurable and quantifiable. You don't have to trust. You just have to demand the numbers, the accuracy, the recognition accuracy, the performance metrics. They exist, and people just have to demand them.

I think one way of establishing trust is to have more of such quantifiable measures and matrices.

And then going from there, because today's AI systems are still reactive to human -- either human likes and dislikes. They are still reactive. Going from there to where we suspect systems are telling white lies on purpose, I think that's a big jump. I think tell you nobody is working on systems like that because we are not that good in making AI systems. But the first type of today's AI system, the performance matrix and trustworthiness matrix, I think scientists

should work together with the standards community, with philosophers, and with business community to establish a set of very clearly verifying, quantifiable performance metrics for every system we provide, every service we provide, such as Facebook, because they do exist. We just don't know about them, and we should demand them.

>> LIU ZHE: This is not an answer, but just a very quick comment. I think for the white lie possibility, if you think about white lie, it broadly exists in human society, in our life. Why not allow these things to our very friendly robots? I think this is really a very big ethical challenge.

The second thing, even though we don't program the white lie into the system, but still you use that into a very intimate emotional relation with human beings, and that risk can now be avoided. It's just like a kind of counterfeit coin, you know, in our currency system. Circulated very well, but it's a kind of deception. Especially those people who really need the help from robots like social robots, they are like -- patients are vulnerable to this. So I think here the risk is really very high. We use this technology, we don't really like to separate those people from the human community but rather to optimize the kind of interhuman relations. So this is the reason why I raise this kind of problem.

>> I agree with you. I think the key there is -- I totally agree. I think the problem there is really training data. Any bot will learn to lie or deceive or even become very abusive from the training data because our systems today are designed end to end based on deep

learning, and they learn everything they define in the data. So I think there we need to mitigate the risk by looking at the training data.

>> CLAIRE CRAIG: Thank you. Can we carry that on, as ever, with the questions, where there are so many, that we are not going to have time to continue with, we can carry them on in the breakouts. I am going to ask you to keep your questions short. There are quite a lot of hands going up, and we are already running a little behind time, but there is a question on the app which has been storming ahead in the lead tables, which is the one with four votes -- I feel like the (?) concontest for those that know it -- how is the different expectation of privacy in China, versus in quotes, "western countries" versus the trust in AI? Can you just quickly say something on that?

>> LIU ZHE: This is a very sensitive issue. I know in China I think the public attitude is really uneven. The educated people have more sensitiveness to this kind of privacy issue. But the general public don't quite understand how dangerous the privacy, especially the personal data breaching, could impair their individual life. I think we need to do more job to really educate and help people understand these technologies, especially in today's era, this is what I observed from my personal viewpoint. Yeah.

>> CLAIRE CRAIG: Thank you. I am going to take -- because there are a lot of hands, I am going to take three questions together, then I am afraid we will have to finish. There's one at the back, then

one there, and there was one over -- I will just do two, one at the back and over there, I will ask Professor Liu to respond to those, then we'll stop. Thank you.

>> Thank you very much. (?) and Professor of ethics. I also teach in China, and it's interesting what you say about the cultural differences. But my question is is the trust not a question of trust in other human beings, not so much the trust in the technology itself? Because the fact of the mistrust, I think, is linked to is the programmer trustworthy? Is the government trustworthy? Do they tell us the truth about accidents or not? Is the company trustworthy? So it's about trust to other human beings which are behind the technology as a framework, the legal framework, the technical, the financial framework, are they telling us what they charge us? So the Facebook issue is not a mistrust in technology; it's a mistrust in the people who run it.

Would you agree with that or not?

>> LIU ZHE: Thank you for this question. Actually, very quick response. It's just like the robot ethics and discipline must be divided on the basis of degree of autonomy of the system. So the moral sensitive -- or sensitivity -- really increases with the higher degree of autonomy. So the current phase, the technologies don't have this full autonomy in cybernetic science in our human moral science. I think right now we really have the problem with the trust within the human circles, so when you talk about the robot ethics, we actually talk about the ethics of the designers, of

manufacturers, of users. But when you think about the autonomy's essential goal for all this scientific development, then we definitely need to think about the space where we should put the HRI relation, human relation to artificial intelligence, especially when the autonomy of the system really increases very quickly.

>> CLAIRE CRAIG: Thank you. One final question towards the front here.

>> The Western ideas about autonomy come very much from (?) philosophers that were circumstances from the particular point in history, in other words, trying to go away from governments at the time that were controlling the situations. Instead in China, if I understand correctly, the confusion idea was to manage chaos in society. So the philosophical approaches that have an impact on policymaking about how we can make society that is more stable. So the two ideas in this circumstance, are probably slightly different. I wonder if that has an impact on how contemporary Chinese philosophers and citizens will see surveillance and the importance of autonomy in everyday life.

>> LIU ZHE: Thank you for this very interesting question. Actually, know, when I designed our document here, we tried to make the framework really inclusive. So we acknowledge a kind of distinction and also conflict between traditional values and modern values. The same conflict, actually, not merely exists in China, but also in Africa, in Japan, and also in Europe. So thinking about the Medieval and also ancient Greek values, also still somehow

inherited in daily life. On the other hand, I think we always, I think everywhere now, people really insist on modern values. So certain really challenging cases, we see this tragic conflict. But how to really develop this integrative system and program to help us really consider the relation. That's why I really like to join this team because it's a way to re-understand, renew our understanding of relations with the world and with not just the human world, but also natural world, and that may be the starting point. At least for me it's the starting point for us to develop an ethical program.

>> CLAIRE CRAIG: Thank you very much.

(Applause)

So we are moving on from that amazingly broad discussion to one that's equally broad, which is about global AI narratives. So this is going to be led by Dr. Kanta Dihal, who is the postdoctoral research associate and research project coordinator at the Leverhulme Center for the future of the intelligence. Professor tokei Takahashi, and we also have up here your former host Stephen Cave, who is the Executive Director of the CFI. But over to you, Kanta.

>> KANTA DIHAL: I am going to first tell you about the AI narratives project to tell you about the ideas that spanned global AI narratives. Then I am going to share some preliminary findings from that AI narratives project before then moving on to the plans that we have for global AI narratives.

So the AI narratives project was originally conceived in late

2016, when Dr. Claire Craig from the Royal Society joined forces with Dr. Stephen Cave from the then brand-new CFI, and this project aimed to examine the stories we tell about AI and the impact they might be having. So the project set out with a few specific goals. First of all, to understand the hopes, fears, and other factors that shape how we perceive AI, and then the imaginings on the one hand and the actual technology on the other. The arrow of causation runs both ways here. The imaginings both inspire and guide the development of the technology, but they, at the same time, strive to reflect that technology. And we are particularly interested in the moment when this relationship between imaginings and the reality of the technology breaks down.

And as well as its descriptive analytical elements, this project has also always had a normative interventionist ambition. It's not that we think there is one right way of talking about AI, nor do we think that we would be in a position to impose it if there was, but we do suspect that there might be limitations in the ways in which we talk about AI now. Those limitations could be that we are distracted from the real problems by fantasies of killer robots. Or they could be that whole sectors of society are put off entering the field because it's framed in a way that seems not to reflect their identity or interests. We hope that by specifying and evidencing these limitations, we can begin to inspire change. As I said, not to one right way of thinking about AI, but to a much richer and more diverse range of ways.

We've been enormously excited by the interest in AI so far. We have had huge interest in those developing, those in industry, and politics, such as the UK House of Lords Committee on AI, which commended our work in their report. But this work is just scratching the surface. The impact of AI will be truly global. Managing it for the benefit of all will require international cross-cultural collaboration. But different cultures see AI through very different lenses. Diverse religious, linguistic, philosophical, literary, and cinematic traditions have led to diverging conceptions of intelligent machines. To build trust across cultures, we must understand these different ways of seeing what AI can and should be. And that's where we need help from international partners. This is why we are setting up the global AI narratives project, which emphasizes the narratives outside Silicon Valley, including narratives from both developing countries and rising AI superpowerers, such as China, Russia, and Japan.

So I will give a brief recap of what we have been doing so far before explaining how we plan to move forward from this. The narratives project has so far held four workshops between May 2017 and May 2018. And we brought together a range of researchers, authors, policymakers, and other interested parties to begin to explore the history of AI narratives and its key issues. From this work emerged the plans for a book that is now in development "AI narratives, a history of imaginative thinking about AI machines," which is due to be published in 2020. But we've also

looked at technologies other than AI. AI is going to be a hugely disruptive technology, but it's not the first to come with such potential for disruption. Thanks to the Royal Society, we learned from the very scientists who were there as part of the development of nuclear power, genetic modification, and the climate change debate both what to avoid, overpromise, and underdelivery, and what to adopt, narratives that anticipate security risks.

(Audio interference)

We notice that AI arrives with millennia of expectations which may or may not include assimilation by the machine; that direct research and through such long history, these narratives have not only influenced the public at large, they've also steered the expectations of those who grew up to become AI researchers in this century.

So one of our project leaders, Dr. Sarah Dillon, and her PhD student have been conducting a research project called What AI Researchers Read, in which they investigate the influence of narratives on those who develop AI technology. And as Stephen and I in our joint research project, AI, a Mythology, are investigating, narratives about AI are often written in tones of wild utopianism or dystopianism. It seems that it's hard to be neutral about the prospect of intelligent machines. They seem to provoke us to imaginative extremes. But at first, these imaginings seem to have been mostly hopeful, fantasies, for instance, of trusty mechanical helpers. So theledest known story of something like AI that we know

of can be found in Homer's Iliad, dating from roughly the 18th century BC. So the God created thinking, speaking handmaidens of gold to assist him in his work. Similarly, in Medieval Europe, there were stories of artificial intelligences in a form of royal retainers, such as copper knights guarding secret gateways.

But as over centuries the technology increased in power, these hopes began to tip into fears. A series of great scholars, including Roger Bacon, were rumored to have created a bronze head that could answer any question. Remind you of Siri, perhaps? But these stories always end badly, with mishaps and the destruction of the Oracle, sometimes by a terrified layperson. And we argue that these fears are inherent in the very idea of autonomous creations. Though it's important to note that that's not the same as saying that these fears are justified. We are talking about our effective and imaginative response.

And it wasn't until the 20th century that that imaginative resonance of Intelligent machines reached its fullest. It was a time of enormous upheaval, rapid industrialization disrupting old ways of life. It was a time of revolution and of murderous mechanized warfare. And against that backdrop, the term "robot" was born, meaning forced laborer, as Professor Liu just mentioned. In a 1920 play, RUR, and famously, in the very work in which the term "robot" was kind, the robots rebel against their creators and destroy them.

The narrative of rebellion has proven to be the most potent of

all our AI fears. We are told repeatedly in terms of the technology of the day from the murderous space ship computer Hal in 2001 A Space Odyssey, to the Terminator films, to the Matrix, in which Intelligent machines farm humans, to humanoid robots once again overthrowing their masters in West World.

We want to create clever tools that can do everything we can do and more, tools that will be the perfect oracles, servants, soldiers, lovers. For these tools to fulfill our hopes, we must give them attributes like intellect, autonomy, and agency. In other words, minds of their own. It's not hard to see the tension in creating beings that are superhuman in capacity but subhuman in status. Our fears of Hal and skynet show that we recognize the deep paradox in our dream of creating independent minds that are enslaved to us. Though our hopes continually threaten to collapse into fears, we hope nonetheless. We continue to have depictions of benevolent AI, such as the digital assistant Samantha from the 2013 film Her. Both kinds of stories, the hopeful and the fearful, reveal to us a complex emotional responses to AI. Managing these responses will be crucial if we are to live with intelligent machines.

One thing will hopefully have been obvious to all of you in this whistle stop tour of AI narratives -- all of my examples are part of the Western narrative tradition. And especially in the 20th century, the narratives that would influence the AI researchers of Silicon Valley -- mostly white men -- was written largely by science fiction authors from the same demographic -- mostly white men. This

is reflected, for instance, in the representation of humanoid robots. They are, by default, gendered and ethnically white, which raises all kinds of stereotypical expectations in the audience, and we look at the ways in which the perception of artificial intelligence is shaped by both fiction and nonfiction and what this means for the impact of AI on those whose narratives are not told or listened to. That creates a value and goal alignment problem. Whose values and whose goals are represented in the development of AI?

In the flowable AI narratives project, we will mobilize scholars around the world to explore how narrative traditions in other regions have shaped both popular hopes and fears for AI and how this influenced the local development and implementation of technology.

So at this Summit, we are seeking to build academic partnerships around the world to bring together experts in a series of workshops in different regions and disseminate their respective findings to each other and to our North American and European partners. The outputs of these workshops would be collected into reports and other media and will eventually culminate in an edited book to be published with a top academic publisher in 2021. We are happy to announce that our partner for the first global AI narratives workshop in September this year will be Professor at that cashi in Tokyo, who will briefly take over for me now. Thank you.

(Applause)

>> Thank you, and good morning, everyone. It is a great honor

to be invited to the prestigious summit at the ITU.

First, a clarification. I speak as a media and communications scholar trying to make sense of the exciting world of AI and robots. I have learned a lot sitting on the Technology Advisory Committee of the Tokyo organizing committee of the Olympic and Paralympic games 2020. After Rio, Japanese government and companies began in earnest to make the Tokyo Olympic Games innovative and sustainable. It is against the background that I shall talk about current issues in terms of new technology in Japan and how can we understand their social impact.

So for today, I would like to give you a glimpse of my ongoing project on global AI narratives within the Japanese context and to invite you to join the conversation. Now we are living in a rich media environment, so the AI narratives are many. They come from all sorts of media, such as newspaper, manga, novels, animation, TV, movies, and social media. So we look at all these different kinds of media and what they tell us about AI. We investigate how many articles, books, movies of AI have been published historically. Then we select some of the more appropriate texts on each medium to analyze.

I am also planning to conduct in-depth interview with audiences because they interpret media texts within their social context. I'd like to understand their perception of AI by using the three modes of interpreting the text -- resistance, appropriation, and adaptation. The analysis of both media text and audiences are needed

to understand AI narratives in Japan.

I will show you an example in terms of differences of media images of AI between Japan and U.S. These are the posters of the same movie in both Japan and U.S. The keywords of Japanese poster are love and friendship, while the U.S. poster focuses on action and battle. Why the difference? And how does it matter in the different culture contexts?

So we'd like to invite to look out for our first workshop on global AI narratives. It is happening on 12 September this year in Tokyo. Please join us to discuss the discussion of how AI narratives differ in different cultural contexts and its consequences. Thank you.

(Applause)

>> CLAIRE CRAIG: Thank you very much, both. So I've got a couple of questions on the app, and we've just got a very few minutes if we are going to -- but we are over about by five minutes at the end of the session because we started late, but we've only got a few minutes for questions. So I am going to start with the app and invite any of the panelists here. There's a question which says many narratives assume consciousness, which is not in our technological grasp. How can we update narratives for modern AI, for example, to show superintelligence but not sentience?

Would somebody like to comment on that?

>> Yes. One of the issues that we often fight with the narratives is that they might quite a big extrapolating leap, often towards

sentient machines, or superintelligence, which I would say is not in our technological grasp currently either. So there is quite a lot of focus on trying to pull out those narratives that work in the near future that are set with the technologies we have right now and extrapolating from them directly. So one of the things that we are working on is working with science fiction writers who are interested in that to connect them to scientists and get them to develop that. And there are several of those plans around the world.

But also in our work, we try to platform existing narratives that do things differently from the standard narrative of a superintelligent or sentient machine, killer robots, et cetera, a very small subset of all the existing narratives around AI and try to give those a wider platform to tell those different stories

>> Thank you. It's a very interesting question and one that comes up a lot because we anthropomorphize a lot with AI.

One of the things we think about a lot is the limits of our imagination. We like to think of our imagination as unlimited, but when we engage mostly with other humans, we imagine the world in terms of human agency to a very large degree and we imagine these machines in terms of human agency, and of course, they are nothing like that. One thing we are all very interested in is thinking about I with as we can try to imagine very narrow kinds of intelligence that are very alien to our own, sort of disembodied, for example, because we believe that if we don't -- if we can't really imagine and conceive of these systems properly, we are not going to understand them properly and

we are, therefore, not going to be able to control them in the way we need to.

>> CLAIRE CRAIG: I am going to take the prerogative and step out of chair for a moment because I have been involved in this project. There are two people essentially complaining about the time scale and saying this is all rather slow. There is a natural tension between the need to be academically and scholarly rigorous and actually feed into a debate which everybody knows is moving incredibly fast. The way that we are in part dealing with in a is by engaging with partners, so for example, the Royal Society, we have been holding workshops with the team here request journalists and policymakers, so it's talking with people as the work is going on as well as developing the insights which are academically acclaimed and take a longer period. We are not being silent until 2021.

Lovely. All right. I think, then, I am going to move on.

Oh, one more then we'll -- one more. Yeah.

>> I have in front of me a book, half a century ago from Jack Baller, and the title is "The Information Society Innovation, Legitimacy, Ethics, and Democracy." The only advice I can give you, it's don't do a special case from AI. Everything is already treated in such a book. Okay? You should be aware that there are people before you thinking more globally, and AI is just a little new case. Okay? And don't build a gas factory with AI and saying a lot of things which are completely dangerous.

>> CLAIRE CRAIG: Yes, Stephen.

>> STEPHEN CAVE: Okay. That's an interesting point. On the one hand, of course, we are very keen to learn from the way in which scholars and others have been engaging in related technologies for many decades, and Kanta mentioned that Claire and the Society hosted a workshop on how other technologies have been communicated. But we have to separate the question from what's happening in the technology from how we conceive of the technology. You can certainly argue what we now call AI is actually a range of technologies, many of which have been with us for a long time. But I think the way we imagine AI is radically different and new because whereas we see most other computing or mechanical devices as basically tools, what we are doing in our imagination with AI is attributing directly to them the attributes of human personhood. And so we conceive of them very differently, even if that bears no relation to the reality of the technology.

>> CLAIRE CRAIG: Okay. There is now a question which I can't resist taking. It's got three votes. Why do we not start a crowdsourced world narrative Web-based repository to which anyone can contribute?

>> So I know of a couple of initiatives that have attempted something like that. I think one of you, two of you are there in the room. So what's the name of the project itself? Bright Mirror? You could just shout and I will repeat it in the microphone.

>> Hello, everyone. Right. With the Future Society and the AI

Initiative, we launched a global civic debate with Bright Mirror and Blue Nove which crowdsourced narratives and story telling by multistakeholder contributors, the public, academics, et cetera, to try to learn some of these perspectives. So potentially, there's opportunity for continuing this in a new iteration because that has recently closed and we are just putting out a report, but we would love to expand on it.

>> KANTA DIHAL: And on the other hand, so aside from the repository, what we want to investigate in the global AI narratives project is what stories are already out there. What is the history of AI narratives in different parts of the world, and how has literary history -- so years of reading stories about this, how has that influenced the the people who implement AI and develop it.

>> CLAIRE CRAIG: Thank you. There are more questions coming in, but I think these are now for the breakout sessions. Let's thank the panel again.

(Applause)

And thank you very much. And we move on to the final project. So this is Professor David Danks, who is Professor of Philosophy and Psychology at Carnegie Mellon University, and Aimee van Wynsberghe, who is the Co-Founder of the Foundation for responsible robotics. I would like to invite them.

>> DAVID DANKS: Thank you. So we -- so to start, I wanted to make an observation based on several of the questions that have come up and actually it feeds off of some of the things that were just talked

about with the last project, and that is I think sometimes when we think about trust and trustworthiness, especially with AI, there is a bit of this feeling that it's a vague, amorphous notion, and we are really not sure what we are talking about. I would encourage people to think about and perhaps look into the over 50 years of research that's been done in social psychology and organizational behavior on the nature of trust. It's actually a fairly well-understood phenomena, a fairly well understood relation. This isn't something that we as AI researchers, as roboticists have to reinvent the wheel. One of the key insights that comes out of a lot of that work is, of course, as you might guess, the fact that there are similarities. The trust that I have that my car is going to start in the morning is very different from the trust that I have that a radiologist is going to correctly diagnose something from an x-ray, is very different from the trust I have in my wife. But what binds us all together is the willingness to make ourselves appropriately vulnerable in response or on the basis of justified expectations about the trustee. So I, as the trustor, make myself vulnerable in order to gain certain kinds of benefits on the basis of appropriate justified expectations about the trustee.

So one of the things that comes up in a lot of these -- a lot of the debates are exactly these questions of what are we justified in expecting of the technology? What are the ways in which we make ourselves vulnerable? And do those fit together in appropriate ways?

So let me illustrate this with an example. I am at Carnegie Mellon University, which is based in Pittsburgh, Pennsylvania, in the United States, which is the home of the Uber advanced technology group, the people building the self-driving cars, that up until about seven weeks ago, were close to ubiquitous in the neighborhood in which I live. So it's always interesting to talk to people who say self-driving cars are these very strange creatures roaming the streets because if you live in Pittsburgh in certain neighborhoods, they are boring. You don't even notice them any more because you see them five to ten times a day every day, obviously not, again, for the last seven weeks.

One of the things that was interesting to watch as somebody who lived in these neighborhoods was the change in pedestrian behavior from when the vehicles were first introduced until right before the tragic accident in Arizona. And in particular to watch the ways in which people started to Jaywalk more and more frequently directly in front of the Uber cars, disproportionately more than they would ever jaywalk in front of a human-driven car. It's the U.S. People text all the time when they are driving. So I am that person who when I am walking to the office will stop somebody and say hey, I noticed you jaywalked. Did you know that that was an Uber car you were doing it in front of? Across the board, with unanimity, they would say oh, yeah, I knew it was an Uber car. I was sure it was going to stop. Why were they sure it was going to stop in they were sure it was going to stop because they'd watched it stop when other

people walked in front. They'd watched it stop when cars would cut in front of the Uber cars. Plus some other unreasonable expectations they had about the quality of the technology.

But so what this points towards is the ways in which trust builds over time in these technologies, sometimes appropriately, sometimes inappropriately, and also the role that's being played more broadly by the cultural and national regulatory context.

So what we want to look at in this project is exploring the ways in which different nation states regulate AI technologies, look at the impacts of cultures interacting with AI technologies, and all of these coming together to either promote or inhibit the development of trust, all understood against the background of the substantial amount of research that's out there in terms of how people actually develop trust in a technology over time.

And I gave the example of the Uber cars and pedestrians jaywalking for a very particular reason, which is that's obviously -- you can't do all of this for all of AI. That's not a one-year project. That's not even a one-lifetime project. Instead, we thought we wanted to focus in on a particular instance, a particular case of technology and interaction with technology that might be tractable in the 12-month timeframe where we would need to be able to make real progress and then come back and report back at next year's AI for Good Summit. So we thought we would focus on the autonomy of self-driving vehicles. There is tremendous variability across nations in terms of how autonomous vehicles are being regulated.

There's variability even within nations, for example, in the United States. There is also a great deal of cultural variability in terms of vehicle-pedestrian interactions. If you just look at the variability between, for example, India and the United States and the UK and Singapore, you see wide variability in terms of how pedestrians and vehicles interact with one another. So it seemed to us to be a really interesting target exactly because we have enough variability both in terms of the regulatory dimension and in terms of the cultural dimension that we might be able to find some interesting feedback.

So what we would like to do in this project is have essentially a systematic survey. What are the differences between the regulation and the cultural norms across diverse countries? We already have partnerships and commitments from individuals and groups in the U.S., China, Singapore, Netherlands. We would love to have many, many more of them. Those were a somewhat skewed population of countries and cultures, though starting to get a little bit of diversity. And the hope -- and Aimee will say a bit more about this -- but the hope is to be able to use this as a starting point that can actually lead to more systematic understandings about sort of best practices, whether in terms of regulation, technology development, social norms, and these sorts of things. With the goal of leading to solutions -- and here I will hand it over to Aimee to talk a bit more about the broader framing.

>> AIMEE VAN WYNSBERGHE: Just really quickly. And also one point

about pedestrians, we wanted to broaden the scope of what pedestrian means. So not just the individuals walking by on the street, but also coming from a Netherlands context, there's a lot of bike traffic; right? So how can cyclists understand when a self-driving car is actually in autonomous mode? So the kind of scope, the very specific case that we were looking at that allows us to make these broader connections was facilitating some kind of communication between a self-driving car when it's in self-driving mode and the pedestrian. So how can pedestrians understand what the car is doing, what's happening?

And so coming from a Canadian context -- and maybe this is also in the United States as well -- when you have a driver who is learning, there is a light on top of the car, almost like a taxi; right? In the Netherlands context, there is an L, just a sign at the back of the car. So that provides information for the other traffic on the street or the pedestrians also on the street. It gives a way of communicating, which is not verbal at all. So we thought perhaps there should be something similar in a self-driving car, some sort of light or some sort of way of indicating to other traffic, other pedestrians that the car was in autonomous mode. So that's the small scope that we were looking at. And then, of course, this is talking about a kind of, you know, if we can get to this level where we create a signal, this becomes a kind of system requirement, this then facilitates the trust on behalf of the developers and the actual bystanders, I guess, not the users but the individuals who become

a part of the social experiment of implementing self-driving cars.

So that's the small scale or the small, you know, pilot project that we were looking at, and then tying all of these national, global thoughts and reflections together. If that makes sense. Short and sweet right before lunch.

>> DAVID DANKS: Yeah, we are keeping you from lunch, so we are not going to drone on.

(Applause)

>> CLAIRE CRAIG: But you are not escaping without a few questions, I am afraid. Thank you.

>> DAVID DANKS: Oh, well. Tried. Sorry.

>> CLAIRE CRAIG: So there's one over there, but also just to get things going, I wanted to ask if you could say a little bit more about the different types of trust that you you gave at the start. Your trust in the radiographer is different from the trust in your wife, and talk about how that projects to the different contexts that you would be looking for.

>> DAVID DANKS: Sure. So again, a lifetime could be done on trust in humans. So at a high level, you can think of trust as dividing up into two primary sorts. One sort is driven very much by -- is grounded in reliability and predictability. So the fact that my car starts every morning, I have the trust that it will start, I make myself vulnerable, I don't go down and test it early in the day even though that might make me late, it's based on an expectation grounded in my past history. You can think about this as a behavioral trust.

The system has worked every time the way it was supposed to work in the past.

Now, this is good in many ways. It can be relatively fragile. So for example, I have no idea whether my car would start if it were sitting in two feet of water. I don't know how cars work particularly, so I can't extrapolate to that kind of a novel circumstance. But as long as you stay well a well-defined context where you have past experience and interactions, it's a very stable kind of trust that tends to be fairly robust against violations, so if the system doesn't behave as you expect it on one occasion, people tend to be fairly tolerant of that.

The other kind of trust is grounded in an understanding of more generally how this system works. You can think of this sort of based in a kind of whether it's a mechanistic understanding or general understanding. And this is often the kind of trust that we have in one another. We know and are able to expect how somebody else is going to behave in a novel circumstance because we know something about the values, their beliefs, their interests, their goals, their desires. And in this case, it is particularly helpful kind of trust because it can generalize to novel situations. So part of what we have in mind when we say things like "I trust my wife," is that I can have reasonable expectations about how she would behave in entirely new circumstances. She's never been to Geneva, but I know basically how she would respond if she were just dropped into the middle of the city right now. And so that sort of understanding trust

really is important when you want to deal with the system in a novel context.

So in the case of the pedestrians with Uber, let me say one other thing. Sometimes people talk about role-based trust. You trust why you are doctor because she is a doctor. That's actually just a route to build one or the other kinds of trust. It isn't a distinct kind of trust.

So we think about interacting with autonomous vehicles, pedestrians, you could imagine trust developing because of repeated interactions. That's the kind of behavioral trust that I think most of the people in my neighborhood in Pittsburgh had. They don't understand how the autonomous vehicles work, they have no idea how to predict what it will do in a novel circumstance, but they have seen it functioning enough times. And I think part of the way behaviorally that you can see that this is the nature of the trust is at least anecdotally, people were a lot less inclined to jaywalk if it were snowing out. So when the weather was bad, that's a new circumstance. People actually, I think, justifiably were concerned about how the vehicles would behave in snow. And so they adjusted their behavior until they had more interactions.

An alternative way of building trust with these vehicles -- and I think this goes to some of the things that Aimee was pointing to -- would be for people to understand something about how the vehicles work. Does the vehicle see me as a pedestrian? Of course, they don't see in the way that we do, but does it recognize that I am there

in the way that if you had a human driver you would make eye contact with them to try and ensure that they see you? So I think that there are these different ways of building different kinds of trust depending on the kinds of interactions we have. So longer answer, but hopefully helpful to some people in the audience.

>> Picking up on the last point, when you have never had an experience with a robot or self-driving car, you can't expect people to have this behavioral trust. We are also focusing on a temporal dimension. What do you do at the very beginning as a way of establishing the longevity of the trusting relationship?

>> CLAIRE CRAIG: Okay. Thank you. So there was a question -- yes, the question there, the lady.

>> Okay. So I think it's really interesting what you are doing. I would like to challenge the last part of your outline, which is coming up with best practices. We are in a human context where the idea of best practices has been well utilized and misutilized as well because we do know that when we talk about best practices, we do make an informed decision what is best and what isn't. And in this sense, what is really interesting is to kind of mimic other AI behavior, which will be to come up with most (?). So instead of looking at best practices, why not look at most appearing patterns within the research so that you can keep the context really alive because this is what makes the big difference.

>> DAVID DANKS: So I think that's great feedback. Approximate I will say, I won't speak for Aimee here. My own view is what I had

in mind -- and maybe I used the word inappropriately, I apologize -- when we were saying best practices was something more like you could think of it almost as a set of if, thens. If this is what you want to achieve with the technology, or if this is the kind of context in which you find yourself, then here's something that seems to have worked well, here's something that didn't work well. So it wasn't -- I didn't -- I certainly didn't have in mind some sort of universal recipe that would be usable everywhere. I think that that would be -- not end well if that were our goal.

>> AIMEE VAN WYNSBERGHE: Which is why we are hoping to have multiple countries participating -- hoping to have multiple countries participating so we also get an understanding of what's going on. So are pedestrians in the Netherlands reacting in the same way as the pedestrians in the United States? Who knows? We don't know right now. So it's also a case of understanding how are people reacting to what's going on and then suggesting modes of best practice.

>> CLAIRE CRAIG: Thank you. There's a question from the app, which is the one -- is there not any higher bias towards AI than to other car braking systems or traffic lights? This issue about when -- is common to a lot of the discussions going on. Could you say something about how you would propose to tackle that in your work?

>> AIMEE VAN WYNSBERGHE: Well, so I don't know if we are placing a higher bias. I think we are reacting to a situation where a technology has been introduced that the majority of individuals don't

understand what's happening when the technology is functioning. So with the traffic lights, you have lessons -- or at least, sorry, from my own perspective, you write an exam, you take in-car driving lessons, it's something that's built into learning how to function on the road. Even if you don't have a license, you understand how the traffic lights work just because virtue of using the streets if you are walking as a pedestrian or on a bike. But with the self-driving car, it's something that we don't have an understanding of, so this is an attempt to try and create an understanding to facilitate not a reciprocal -- I guess a one-way interaction to provide information to the pedestrian.

So I guess we are placing a bias in that we are focusing on this technology, but we are focusing on it because there's just such a lack of understanding.

>> DAVID DANKS: I think this also connects directly to something that Stephen was pointing out about the anthropomorphization of AI. I think it would be interesting to try to have these interact in some way. Undoubtedly the way the pedestrians and bicyclists think about self-driving vehicles is going to be based on how -- you see things like people thinking that self-driving cars distinguish between an elderly person and a young person if they are thinking about trolley problems, which they don't. They just don't understand the world in the same way we do. So I think there are a lot of these questions.

>> AIMEE VAN WYNSBERGHE: Another question. You know the duplex that came out, and now there is a question of whether the AI should

be interacting with the human, whether or not it's an AI doing that. So this conversation is happening now in multiple applications of AI, that there needs to be some sort of facilitation of knowledge of when you are interacting with AI, what it means to interact with AI, so this is one small-scale project to bring this into light. And there are also questions when you come up, you interact with a robot. Should there be a light on the robot that goes off when it's collecting data about you? These are conversations that are happening. This is just one, I guess, spotlight on the issue.

>> CLAIRE CRAIG: Thank you.

Francesca.

>> FRANCESCA ROSSI: Yeah, so this is not a question. I mean, I really like this project. It's very -- it tackles very general ideas, but in a very specific context, so it's, you know, I think it's a recipe for success.

I wonder if you have been looking also at this recent announcement by Andrew Ng, who is going to deploy self-driving cars starting July of this year in a small city in Texas, U.S., with special attention of the interaction with pedestrians. So for example, is planning to put some displays in the car that say to the pedestrians "I see you. You can cross the street." Or whatever. So that -- and I think that I mean, it's not necessarily that those solutions are going to be the solutions for the problem, but I think it's an interesting experiment to explore the dimension that you have in mind, and also with real cars and real pedestrians, and you may want to interact

with him to brainstorm about how they came up with certain kind of plausible solutions.

>> AIMEE VAN WYNSBERGHE: So do you have a contact for us?

>> FRANCESCA ROSSI: Yeah, yes.

>> AIMEE VAN WYNSBERGHE: That would be a wonderful starting point for us.

>> FRANCESCA ROSSI: Yeah, because that gives you a place which is very restricted, very, you know, not just -- not like Uber cars. It's very restricted itinerary where the cars will go, and the pedestrians will interact in that way.

>> DAVID DANKS: There have been a couple of -- I mean, you know, the one that probably is best known is was at Stanford. I think Andrew was involved in that at one point. But looking, you know, starting to try to look at these questions. I think one of the things that I think is interesting is to the best of my knowledge, that work has tended to focus on -- has tended to have somewhat a universalist idea that if we can figure out a solution to it, it will generalize to other nations and other cultures, and I think that that, you know, the ways that pedestrians move around the Stanford campus might not be the same as in others. So I think that it's a data point, but I think it would be great to be able to situate it in this larger discussion, so I think that's wonderful. Yeah, thank you.

>> CLAIRE CRAIG: Thank you. One final question, then we'll break for lunch.

>> I also agree with Francesca. I think it's a really good project

because it's very well scoped. What I think it might be interesting is to look into the two types of narratives that could come up. One is one about the the technology, the vehicle is one thing, so people might trust a vehicle in a particular way, but they might trust or not trust Uber in a particular -- because technology is just mediating your relationship with the company behind, so it's not really autonomous. It's just following the goals of the company behind, and they will have their own business models, they will have an impact on the socioeconomical environment, so people might start questioning, well, if I took the Uber autonomous car, all the taxi drivers end up without jobs. In some countries that might be more important than others. Or they might be, oh, they don't pay taxes in this country. You know? So there might be different narratives that come at different levels. Particularly those two levels I think might be interesting to distinguish when you are collecting the data because otherwise you have different expectations on what to talk about.

>> AIMEE VAN WYNSBERGHE: Yeah, I really like that idea. I think, then, when exploring those narratives, it would be relevant or important to point out that you can still dislike or distrust Uber -- the Uber -- but you know that when there is an Uber car in front of you you can at least rely on it that it will indicate when it is in this mode. So I like the idea of when do the two meet and when can you still have separation between the two. That's a really nice --

>> CLAIRE CRAIG: This is great. We are building and developing

the project as we go along. I am going to stop it because we need to have lunch and then reconvene for the third session at 1:30. But now I just wanted to thank you for your questions and for listening and all of the panelists for these really three fascinating projects.

(Applause)

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.
