

FINISHED COPY

ITU AI FOR GOOD GLOBAL SUMMIT
POPOV ROOM
PLENARY 12
CLOSING SESSION - APPLYING AI FOR GOOD

GENEVA, SWITZERLAND
9 JUNE 2017
1630 CET

Services provided by:

Caption First, Inc.
P.O. Box 3066
Monument, CO 80132
1-877-825-5234
☎719-481-9835
www.captionfirst.com

This is being provided in a rough draft format. Communication Access Realtime Translation (CART) or captioning are provided in order to facilitate communication accessibility and may not be a totally verbatim record of the proceedings.

>> REINHARD SCHOLL: We are going to start in a couple of minutes. When everything is going to be finished, we are going to finish before 6-ish, is there interest -- let me ask, is there interest for some of you to go for a drink just next door to get a feel for if someone would like to participate? Then we can reserve that. Okay! That's sizable.

Then let's meet at the reception -- yeah, right, that's where you picked up the registration. Registration, where you picked up the badge and where we assembled yes yet to go to the event.

Let's meet there at 6:30 and it's just across the street, 30 meters to walk. Okay? 6:30. We'll go for a drink, whoever would like to go for a drink. Thank you.

>> DOREEN BOGDEN MARTIN: Okay. We are going to get started. Good afternoon, everyone. Almost, I could say, good evening. My name is Doreen Bogden Martin, Chief of Strategy here at the ITU.

I will be moderating this closing session. And I am amazed at how many of you are still with us at 5:00 o'clock on a Friday.

(Applause.)

>> DOREEN BOGDEN MARTIN: So wow, is what I have to say. Wow, it's really been an amazing three days. It has been such an incredible event for me and for all of my colleagues here at the ITU. Like all of you, we have been inspired by what we heard. We have been intellectually stimulated, often challenged and sometimes, as we are human, even a little bit scared. The energy was contagious here in this room, outside this room, and also online.

I wish that there was an algorithm that could capture the mood of the participants in this room over the past three days. But as that's not possible, I will be happy with the satisfied faces that I see on all of you here, with the new partners that we have engaged. Many of you have come to the ITU for the first time. We welcomed all of you on Wednesday, and we bid farewell to all of you today, but we hope you'll come back. It really has been great to have this sharing of experiences with all of you and to really engage in all of these important discussions over the past three days.

I would like to, before I turn over to our speakers, to thank again our co-organizer, the XPRIZE Foundation and, of course, the 20 U.N. agencies that have partnered with us in this initiative.

I would also like to give a special thanks to our sponsors, the Kay Family Foundation, Wiki Omni, Word4App and, of course, our corporate sponsor, PWC.

During these days, we have heard amazing talks, moon shot inspirations for the future, AI state of play, transformations on the horizon. We have done lots of work in our breakthrough sessions, some even with colored post its. We have been brainstorming on several ideas from privacy to ethics to societal challenges, sustainable living, capacity building, poverty reduction, ending hunger, AI for prosperity, investment, economic aspects, and designing the future.

And as we come to an end of these three inspiring, forward-looking days, I am much honored to have with us this afternoon an outstanding set of speakers once again -- because we have had outstanding speakers for the past three days. So we have some outstanding speakers for this closing session that will leave you with some thoughts on continuing collaboration and working towards the achievement of the 2030 sustainable development goals.

And with that, ladies and gentlemen, I'm going to turn over to our first speaker. We are going to hear from three speakers.

Then we will open up very briefly to a few questions. And then we will hear some final remarks from Marcus as well as from our Secretary General.

First we will hear from Peter Norvig. Peter, as many of you know, is the Director of research at Google. Especially is joining us remotely this afternoon. Prior to joining Google, Peter was the head of computational sciences division at the NASA Ames research center. He was also a Professor at the University of southern California and the University of California at Berkeley. So Peter, over to you, please.

Please, go ahead, Peter. Can you hear us, Peter?

I heard something. No?

>> Remember the saying, AI is easy, AV is hard! That's one of Brad Templeton's jokes.

>> DOREEN BOGDEN MARTIN: I'm looking to the technicians. Should we pause for a second? Or we can go ahead to our next speaker?

Okay, going to old tech!

(Pause.)

>> DOREEN BOGDEN MARTIN: Maybe we let our technicians sort things out. Perhaps if we can turn over to you, Stuart. Many of you may have seen Stuart Russell over the past three days. He is a Schmidt Professor in engineering at the University of California, Berkeley. Stuart, over to you.

>> STUART RUSSELL: Thank you very much and I should also mention I'm Peter Norvig's coauthor on the book. Peter and I have a long standing connection.

So I would like to begin by thanking the ITU for pulling off a near miracle in bringing so many U.N. agencies together and so many representatives from the AI community and other organisations around the world. I have tried and failed to do this kind of thing in the past. I'm very impressed how well it's worked. And I think this is obviously the first step on a long road. And we have a pretty successful foundation in place. I think it's very important that the participants understand the ground rules for collaboration. And for the AI community, I know that you are all kind of raring to get your feet into the problems and get hold of the data and start measuring your results and announcing success, and so on.

But unless we understand, we the technical community understand the ground rules for how to collaborate, to understand what the U.N. does and doesn't do, to listen to the people on the ground who are working on the real problems, and to understand that this is not just an opportunity for us. This is actually a responsibility for us. Because as we develop capabilities to solve problems, we have a responsibility to use those capabilities and not just an opportunity.

With that responsibility comes the kind of humility that other branches of science, for example medicine, have understood for centuries; that in medicine, the patient comes first and not the solution and not the medical expertise. But it is the only thing that matters, is the patient.

So we have to develop a similar kind of humility. It is not that we have a fancy algorithm or not that we can get better results than someone else. The only thing that matters is the end result for the people who are what matter.

So I'll talk a little bit about my own experience working with the U.N. in the past. So I worked with the CTBTO, the comprehensive test ban treaty organisation which is headquartered in Vienna. They had a challenge which is to monitor the entire planet for seismic events which could be clandestine nuclear explosions.

There I think we were in a good position because they already understood this as a data processing problem. They had a very large data collection network in place, satellite communications sending terabytes of data every year to Vienna for processing. In some sense the input and output was already well defined. That gives you a leg up.

But still the key step was listening to and understanding what exactly is the problem and what are the constraints under which we operate.

So those constraints include rules in the treaty about what kind of data is allowed to be used, what data can't be used, considerations of security and verification of the data itself to make sure that incorrect data is not used to attribute nuclear explosions to a country that is perfectly innocent, and so on.

And then there are some administrative barriers. It took us two years to negotiate a zero-dollar contract. So that was a bit of a surprise to me. So patience is a virtue. So the United Nations is a kind of sovereign organisation, but the University of California is also a sovereign organisation. Each of which believes that they are entitled to set their own contract terms. So patience is a real virtue in dealing with this.

The kinds of problems that we've talked about in this meeting are often much less well defined. There isn't already a computational problem that is set up. There isn't already a data feed that is set up. When you think about the problem of hunger, it is not even clear that AI is a fraction of, a large fraction of the solution to that problem. And so it may take quite awhile, a long set of conversations, and really listening -- maybe even going out to places where hunger is a problem until you understand what role might AI play.

And I would also say to the agencies: Don't believe everything you read about AI in the newspapers. That it's just a matter of collecting a lot of data and then deep learning will do something magic and your problem will be solved. I think many real problems simply don't work that way at all. Systems in AI, there are lots of other kind of systems besides machine learning systems. Often data is a useful adjunct but it is by far not the most important component in building successful solutions. And so setting up an ongoing platform for a wide ranging conversation between the agencies on the ground and the technical communities that may be able to contribute to the solution is going to be, if this conference is to be a real success, it will only be because those conversations take place in the future.

So maintaining the momentum from this meeting and setting up the area specific Working Group, involving the NGOs, researchers, and really the industry partners will be important because in setting up global scale communication networks, or logistical capabilities, often industry is a crucial part of that. I was at the Red Cross this morning and they use DHL to solve a lot of their logistic problems on the grounds that DHL is in the logistics business. They are efficient, because they have to be. Otherwise they would go out of business.

And why not use DHL for that capability? The idea that we are going to come along and produce a better logistics capability than DHL might be a built presumtuous.

That kind of public-private partnership is something that, the World Economic Forum is specifically designed to facilitate. They are the international organisation for public-private partnerships. They are just across the lake. And there are people here, Rigas was on the stage half an hour ago, and he's from the Forum. Their resources include the Global Futures Council on AI and robotics which has 25 people, many of whom now have a reasonable amount of experience of interfacing with policymakers and thinking about these large scale challenges.

Also there are 900 experts in all areas, everything from the drug trade to food to climate to the Arctic, to the oceans, all the industry sectors, telecommunications, computing, mining and metals, chemicals, et cetera, et cetera. There's an enormous reservoir of expertise. I think I can speak on behalf of the AI and robot particulars Council in saying that the Council is ready to participate in these Working Groups to move progress forward on the SDGs.

The idea of a CERN for AI was brought up and I think that was maybe Gary's suggestion. And I don't know if that's quite the right metaphor. CERN was set up because they needed a \$65

billion piece of equipment and it's hard for a university to have such a thing.

For AI we don't need a \$65 billion piece of equipment. We don't need a new institution built around that. But I think what we do need to do is be aware that if we are successful in the Working Groups, we identify problems, we figure out how to solve them using AI. Then setting up an organisation to maintain, to roll out and maintain that in operation is going to be very important. And there are many, many successful applications of AI that then evaporate because that follow-up organisation and effort is not in place. So I think that's a very important thing to plan for.

Lastly, on Wednesday in my Plenary presentation I spent a lot of time talking about some of the risks and down sides because I am sure, I was sure and I was right that everyone else was talking about the opportunities and the up sides. And that's been a very positive part of this meeting. How many opportunities there are. But I want to repeat something I said in one of the breakout sessions that the area I see where AI could make the biggest impact that would have beneficial effects across many of the SDGs, and many other of the United Nations' responsibilities as well, is the provision of high quality one to one education at relatively low cost. This is a grand challenge for AI. It is not a simple, not as simple as putting a few apps on cell phones. The reason why AI hasn't been successful in education so far is two fold. One is the AI system has been unable to communicate with the pupil. And because the AI systems have not had up to now the ability to communicate in natural language. So they have been restricted to clicking, selecting buttons on menus and canned text, which is not a very good communication medium.

The second thing, the AI systems have not known anything about the material they are trying to teach. They are teaching physics without understanding anything about physics. Or they might be teaching about the mechanics of ink jet printing without knowing anything about ink jet printing whatsoever.

Imagine yourself trying to teach a pupil something about which you know nothing in a language that you don't speak. Right? That's sort of the situation that AI for education has been in for the last 40 years. There is a chance that this is starting to change fairly quickly, both speech communication interfaces, the beginnings of real language understanding, which will still take time to achieve, and then the capability for fairly robust representation, reasoning about the content knowledge that the system is trying to communicate. These are not solved problems but I think they are solvable problems in

the medium term. That will make a real sea change in our ability to deliver high quality education.

I have to say this is not just a matter of in the developing world or the LDCs. This is in the Developed Countries as well where the standards of education are far below what they could be. The potential of almost every human being is wasted under our current schemes for educating young people.

Thank you.

(Applause.)

>> DOREEN BOGDEN MARTIN: Okay, thank you very much, Stuart. I think we are going to turn back to Peter. I understand that it is working. So Peter, apologies for the technical problem before. I believe that we can now hear you, I hope.

>> PETER NORVIG: Yes.

>> DOREEN BOGDEN MARTIN: Okay.

>> PETER NORVIG: Hello, everybody. And thank you for letting me speak to you. I'm sorry I can't be there in person. I love the idea of this meeting. AI for good is just a fantastic topic. Although I guess when I think of it, I would focus on the goods part, not the AI part. I remember my colleague Pat Winston said that AI is like the raisins in raisin bread. By that he meant that raisin bread is mostly just bread. That's the core stuff and then the raisins are sprinkled in to make it more exciting.

I think we look at problem solving for good. AI can be a component to that, as can any of the appropriate tools. Just think about what tools are appropriate and sometimes that is going to be AI. That's very important and we've heard some great ideas here. But sometimes AI is not the biggest part of the problem. Sometimes it's a truck or a shovel or a fleet of satellites. At Google we have done some work with trying to cover the planet. For the first time take a picture of every spot on the planet every day. And we were able to achieve that. We sold off our assets to a company called planet ink. So we can now get a picture of everything. This is very important for things like monitoring how crops are growing, for disaster relief, and so on.

And mostly it is a communication satellite issue, but there's a little bit of AI there because there is low resolution and high resolution images. You have to decide exactly where are we going to spend these resources, for what high resolution pictures are important to take next. There's an AI planning problem to do that. And when we can do that, we can help the farmers. We can help prevent famine and so on. AI plays a small role in everything that is done there.

Another example, my come for example Eric brewer did a project on treating blindness in Africa. What it was was going

out to the villages with nurses that could triage various diseases that people had. Some could be treated right there. And some had to be sent to the hospital to work. And there was a technical component, figuring out remote imaging and communications to be able to make a link between the remote villages and the hospital in the city. There's some very technical work about how that was done. This was done again today, perhaps you could throw in machine learning processing to look at the eye disease.

The important part was making the trek out to the village and the high-tech communication was less important.

And another big part of it was understanding the culture. Saying yes, I can come here and people want to be cured, but you have to forge a partnership with them. And in particular, one of the issues was for some people you say well, if you go into the city and go to the hospital we can do an operation and your sight will be restored. But there might be a three, four, five-day walk to the hospital and people were not going. They said I don't believe you, that that's going to happen. I'm not going to take that trip.

So they came up with a very clever low tech solution which was they handed out a little certificate that said if you show up at the hospital you will get an operation. And if you hand that certificate in, you get the operation. If you don't have the certificate, you still get the operation just as well, but it was building trust, was the key point here. That little certificate was part of the trust. And the relation between the people was part of it as well.

And the high-tech part was a small part of the solution. So once you understand the problem, once you built this trust, once you have that level of understanding, then I think one of the keys is communication and dissemination of information. AI certainly plays an important role there. Stuart touched on it in terms of education and being able to do one-on-one tutoring and do that well. That's certainly a goal that we've all been working for. I've done some work on that.

And I hope we can achieve that and make progress on it. But there's other things that are easier as well. Just having the bare minimum of knowledge and being able to spread that and communicate it. You don't have to say here is a complete diagnosis for how to treat disease. Sometimes it's important to say when you are doing this, it would be good to wash your hands and make sure that everybody knows that. So the basics are often more important than the high-tech machine learning.

And in terms of the communication and dissemination, there are teams to help out with that. We've done some work with Google for nonprofit, making all that available so that

information can get out there. And that's true from the organisations to the customers they work with and it's also true within the organisation, that if we solve the problem, if we have the best practices, then we shouldn't have to solve that over and over again. Certainly we need to be able to apply it to specific situations. Every situation is going to be different. But we should have some repository of resources and best practices and some information retrieval and machine learning capabilities to be able to access those resources.

Important role can be played in building that repository, in funding and supporting the creation of the repository, and in funding and supporting the individual organisations and the individual people that go out there to be able to solve these problems. Stuart recommended World Economic Forum, and that might be a great forum for it. There could be others as well.

So there's lots of problems out there. AI can be one of many tools, but I don't want to place AI above the shovel, the truck, and all the other important tools that we have at our disposal. So let's use it appropriately where we can and I think with that, we can start to do good. Thank you.

(Applause.)

>> DOREEN BOGDEN MARTIN: Thank you very much, Peter. And we are very glad that you have not abandoned your fixed line phone as more and more people get rid of their fixed line. We are very happy today that you didn't do that.

Also I just wanted to pick up on your point about the repository. We have heard that from the session that just reported back, that that is something that has come up throughout the week about the importance of sharing experiences, best practices, and not reinventing the wheel. And that repository suggestion has been made by others.

As we do have you on the line and we can hear you, we were going to open up for questions afterwards, but perhaps I could pause here and ask colleagues, maybe we can take two questions? And then we will let Peter go. So Peter, if you would agree to stay with us for a minute I could perhaps take two questions from the room.

Okay, please, go ahead. If you can state your name and then ask your question. Go ahead.

>> AUDIENCE: Thank you very much. My name is --

>> DOREEN BOGDEN MARTIN: You have to push once on the microphone and wait. There you go. Please, go ahead.

>> AUDIENCE: Thank you very much. For this excellent presentations. My name is Barbara Bulse, Global Development Impact. I am very excited about the potential of AI. However, I haven't heard much about how AI can be used to foster collaborations between various stakeholder groups to truly solve

the problems. We know we have the solution. Financing is there. Technology is there. The key challenge is that we do not really collaborate the way we could. Thank you.

>> DOREEN BOGDEN MARTIN: Okay, thank you, Barbara. Perhaps I can take one more question and then we'll turn back to Peter. Anyone else have a question?

Gentleman in the back, please go ahead. State your name, please.

>> AUDIENCE: My name is Francois. I come from Paris. My question would be, you said we need other tools to solve the problems. I say we also need other intelligence. So you know, we talk about artificial intelligence, but how can we optimize the combination of human and artificial intelligence? Of individuals and collectives?

>> DOREEN BOGDEN MARTIN: Okay. Peter, can we turn back to you for those two questions?

>> PETER NORVIG: Great. So the first question about collaboration and what role can AI play there. Again, I think AI plays a minor role in terms of collaboration and I think the most important part is, one, you need some level of tools, of being able to create materials, whether it's written or spoken or video or whatever, and have those available.

Then you've got to get them out to the right people at the right time and you have to make the right connections of one-on-one to have people talk to each other and work with each other. And once those connections are made it seems like AI is out of the loop. So the role for AI may be in finding the connections, of saying here is someone, here is what their interests are. Here is another person and here is their interests and we have a lot of experience in AI in doing that kind of collaborative filtering, recommendations and so on. We can apply that to the problem of collaboration saying it would be good for you two guys to collaborate.

Once we've done that, now I think it's a personal one-on-one problem and not an AI problem. It would be best for AI to get out of the way once people are connected, and let them handle the collaboration themselves.

The second question was around the combination of human and artificial intelligence and working together. I think that's a great area. We have seen a lot of work on that in terms of being clever about figuring out what data do you want to collect, what human usage do you want to make of that data. So we have done a lot of work in term of things like language understanding, speech understanding, and machine translation. And when we go for a new language, there is, most of the work is in figuring out how we want to take advantage of the humans that interact with the artificial intelligence.

So we go into a new country or a new area, find native speakers of the language, have them say some words, have them have a dialogue with them and get sort of the key pieces of information that we need from them. It's a lot of work behind the scenes for the AI to figure out what are the important questions to ask and how can we best make use of this resource we have, the human resource. That's an area that has been studied a lot and I think being able to take advantage of that partnership of human collected data and human labeled data and then raw data and use active learning to intelligently figure out what the right mix of those is, that is going to play an important role.

>> DOREEN BOGDEN MARTIN: Thank you. Thank you very much again for joining us. You are welcome to stay on the line through the rest of our discussion, but let's give him a round of applause in case he leaves us.

(Applause.)

>> DOREEN BOGDEN MARTIN: Thank you. Now I am going to turn over to another remote presenter. We have Eric Horvitz on the line. Eric was with us at the beginning of the afternoon. You may recall, I think it was 4:00 a.m. for him and maybe he went back to sleep. Otherwise we say good morning to you, Eric. Thank you for spending a long evening following us remotely. We are delighted to have you here.

Eric is a technical Fellow and Managing Director at Microsoft Research. Eric, over to you.

>> ERIC HORVITZ: Can you hear me okay?

>> DOREEN BOGDEN MARTIN: We can hear you. Are you on a mobile or fixed line?

Go ahead.

>> ERIC HORVITZ: Okay. So --

(Bad feedback.)

>> ERIC HORVITZ: I can't see the projector, so I don't know where I am with the slides. Let's just start with the head of technology. I think it's clear to me that we are barely scratching the surface of what is possible even with today's technologies. And by trying hard we learn about the bottlenecks and where we need to go next and what kinds of technical leadership we need and advances that we need.

There is so much to do in healthcare education, transportation, and there are a number of solutions, even laboratories that have not been translated into the open world. It turns out that translation is not so straightforward. It involves the larger understandings of domains. But when we have meetings like we had this week to prepare for advances, as well as rough edges ahead. These tease are large opportunities, I

think, but also causes for concerns. Stuart pointed out some of those earlier in the week.

Next slide. I think it is critical to invest hearts and minds and dollars on the aspirational challenges, and as well as on the rough edges and dangers. It is implicit we need to work too as we go. The communities working with communities interested in applications. Also need to continue to find technical bottlenecks and next directions.

The graphic I put up in Plenary 10, this captures the idea that AI experts have to engage deeply and vice versa. There needs to be a joint sharing, going all in on understanding, with the nuances of a domain end-to-end, bringing in all these relevant stakeholders and working through different models of engagement. I mentioned a couple examples earlier. Cholera, activities around cholera, diagnosis, understanding how to grapple with traffic situations, for example, various kind of automation.

Next slide. So there's quite a bit, quite a few things we could be doing. Just this week ITU, this group here got together. I loved sitting through the last sessions quietly. I wasn't sleeping. I was doing some of the read outs on the breakout groups on shared goals, guidelines and best practices. That's fabulous, but also doing spot studies on concerns as they arise and opportunities. In February I led a meeting, shared -- we had blue teaming around problems spanning a fairly large ontology of possibilities. I found that eye opening for more of the attendees. We have to do more of that kind of thing.

We have to do processes for studying and monitoring. I have been involved in a 100 year study on AI which we will continue to do fresh studies over time and it is endowed to go for as long as Stanford University exists. There are research centers that have quite a bit of work going on in academia and research.

There is work going on with multiparty stakeholders and initiatives I have been involved in. We talked about the partnerships, but I thought I would bring it up and provide background quickly. It dovetails beautifully with the work going on this week. I hope there will be incredible synergies, for example, the output of this week to be injected into the partnership on AI. Let's go to the next slide.

So the partnership of AI started with a discussion among researchers at Amazon, Google, IBM, Microsoft and Facebook, later involved Apple and we brought it to the board of directors and other sectors, AAA scientists the nonprofit foundation world. Basic goals of this organisation which will be a 501(c)(3) nonprofit would be largely to develop and share best practices on research, development, testing and fielding of AI technologies. I started out with a focus on industry, but of

course that could be broadened. Other aspects include advancing public understanding, providing an open platform for discussion and engagement, and identifying nurturing, aspirational efforts in AI for social purposes.

Next slide. One of the early things we did as a group of companies or as researchers at companies was to think about tenets, eight of them that we all could agree to. I put two of them up here. I found it a fabulous challenge, to agree on eight very important basic assertions such as we are committed to open research and dialogue on the ethical, social, economic and legal implications of AI and that we seek to develop AI research and technology that is robust, reliable, trustworthy and operates within secure constraints.

Next slide. Today the board of directors is a multiparty board that has within it from the founding companies as well as the AI research area, ethicists, privacy experts, and representatives of AI scientists and the foundation world.

Next slide. The board got together and identified seven pillars of activity that really align nicely with some of the discussions going on this week for the programme, per my listening in. I couldn't be there.

First area is safety, critical AI, second, fair transparent and accountable AI, collaborations and inclusiveness and bias in data sets and resulting systems and recognizers. Four, collaboration between people, another one on AI, labor and the economy. A fifth one is on deeper social and societal influences of AI. We explored this in the disruptive AI outcomes meeting. Some of them are interesting, deep, and implicit and are concerns.

Then AI and social good, looking at aspirational projects and technologies that could advance the technology of AI, and special initiatives. One that is interesting is of interest to all members. How do we address the data divide between large companies and NGOs and academic world who don't have access to such data sets?

Next slide. So recently we announced a growth of the partnership. We announced the first initiatives.

These are the kind of groups that are joining now as full partners, Human Rights Watch, the Electronic Freedom Foundation, the Center for Democracy and Technology, UNICEF, and I'm happy to say XPRIZE.

We also had other companies joining. We this is growing in that space but we are now going international with Intel, Sony, SAP, SalesForce, and other companies joining up. On the next slide, efforts discussed by the board that are now underway. The ITU and XPRIZE folks are discussing Working Groups. The

partnership on AI is standing up topic specific as well as sector specific Working Groups.

Another project that was a passionate area for us is funding people outside of the area of computer science. People in the nonprofit world doing governance, privacy, civil liberties to fund these people to become educated, to travel to conferences where they can learn to engage, to help with the divide between the technical folks and understanding that people outside the technology will have.

Other areas, looking at prizes and recognition for pieces of research and AI and people in society. And we also would like to do a an AI grand challenges programme. The model that we looked at was less so XPRIZE style and more so national academy of engineering in the U.S. which listed 12 grand challenges in the sciences several years back. We would like to stand up with the committee what are the key grand challenges in AI when it comes to AI, people, and society? We hope to see great collaboration there between the work going on here this week, the XPRIZE folks, ITU and so on.

Moving to the next slide. So what we can do, I would love to see the full transcript an notes from this meeting and slides. I think the board of directors of the partnership would love to engage. Let's think about how to work together in some new ways. I made an earlier comment, well, also mention Microsoft. I want to see Microsoft in the final talk. My next slide will mention very, very briefly because we hope that some things going on at Microsoft will be best practices of sorts in terms of stuff that is in process and looking at other companies as well.

Go to the next slide. Microsoft set up a panel called ether. AI ethics and engineering and research. This group reports to the senior leadership team and with monthly reviews, challenges, issues around ethics, AI, people and society and the kind of topics we talked about that the partnership on AI has highlighted. We would love to see the kind of committee at other corporations. What AI is doing is providing a conduit where best practices including process across industry and the nonprofit world can be shared.

So next slide. Let me just say, fabulous meeting. I wish I could have been there for all of it. I loved the few hours. I love the idea of the coalition and collaborating broadly. Thank you very much.

(Applause.)

>> DOREEN BOGDEN MARTIN: Thank you very much, Eric. So I think we have time perhaps for two quick questions for either Eric or Stuart. Please.

>> AUDIENCE: I wanted to clarify both the facts and the reasoning about the CERN proposal, first because the numbers Stuart cited is already spreading through the Twitter verse it's off by an order of magnitude. It was seven and a half billion dollars and not 65. There were earlier pieces of equipment that were significantly cheaper than that, I suppose. So it wasn't quite about that.

But also in the Twitter verse someone pointed out the point is not about the equipment, it's about the collaboration. It's unfortunate that Stuart wasn't at my talk. I'll try to reprise the argument as quickly as I can. One of them was about public good versus private good. So other things being equal, it is likely a very small number of people will wind up owning most, small number of corporations. It was very much about the inequality issues.

Part of the motivation is that the public might have some ownership of AI that otherwise might not. And so the price of the equipment is not really the relevant measure there.

The other point that I would like to make is, and it was implicit in the talk, it is not clear that we figure these things out. Maybe we don't need a 65 billion or even six and a half billion dollar piece of equipment. We have been working on strong AI for many years and made no progress. We made progress on things like categorization and probabilistic logic, but on the whole, common sense, natural language, and some people would call AI complete. It might be for those problems that we need a large collaborative effort and it doesn't happen through large academic or corporate labs to fix, give the topic of CERN AI more thought. Thank you.

>> DOREEN BOGDEN MARTIN: I am going to give Stuart and Eric an opportunity to make perhaps a brief comment, if you can, and then we are going to hear from our two closing speakers.

Please go ahead, Stuart.

>> STUART RUSSELL: Yes. So I don't disagree with the argument that it is important that the ownership of major progress in AI remain in the public domain. It remains to be seen whether creating a new institution with a very large number of researchers is the right way to achieve those goals. I also don't agree that we made no progress towards strong AI. By which I think you mean general purpose AI. I think everyone understands that there isn't going to be a single magic bullet that produces general purpose AI and we haven't found the magic bullet yet. I think there are a lot of conceptual components. Some of them have actually been in place for a long time. I mean, I think I would say an understanding of formal logical reasoning is a component of general purpose AI and it has been in place for two and a half thousand years.

So I am not one who believes that all the progress that has happened over those millennia is a complete waste and hasn't contributed to the goal of achieving general purpose AI.

>> DOREEN BOGDEN MARTIN: I'm, sorry, going to stop you there because some colleagues actually need to head to the airport.

Eric, if you want to say something for 30 seconds? Then we have to wrap up. If you are still with us?

>> ERIC HORVITZ: Yes, I'll use that 30 seconds to say there's uncertainty about how to go forward.

(Feedback.)

>> ERIC HORVITZ: We have addressed the concerns and ... we should be inclusive. We should have groups lick the AI partnership, which -- there was an effort last year ... others to really ... we need collaboration.

>> DOREEN BOGDEN MARTIN: Thank you, Eric. Of course if we had all of the answers we wouldn't be here. I know there's lots of colleagues that would like to say other things. This is the beginning of a discussion and we look forward to discussing with you further and collaborating if you further.

With that I am going to turn over to Marcus who at this point doesn't need any introduction, I don't think. Except to say that he is the CEO of XPRIZE Foundation and the co-organizer of this event. So Marcus, over to you for your closing remarks.

>> MARCUS SHINGLES: Yes, I'm going to stand up because I'm going to be grabbing some, I want to recognize a few people and say thank you. First of all, Secretary General, thank you for hosting us today. We appreciate all your support. Thank you very much for this.

(Applause.)

>> MARCUS SHINGLES: And I wish I could give you something fancier, but I have a backpack called AI For Good and that's yours now. That's what you get.

(Laughter.)

>> MARCUS SHINGLES: Some of the people here, so Reinhard Scholl, where is he? Everybody give him a hand. He was instrumental to the event here.

Stephen Ibaraki, over here.

(Applause.)

>> MARCUS SHINGLES: And many of you know Amir Banifatemi, who is -- yeah.

Amir is currently running our global IBM AI XPRIZE. You know him because he was critical in pulling this together. Thank you to the three of you. We are creating a really productive environment over the last few days.

I also would like to recognize Yvonne Cooper from the XPRIZE team. She is -- thank you, Yvonne.

A lot of the experience you had here, all the hallways and the atriums and exhibits, that was produced by Yvonne and the XPRIZE team. I hope you experienced it as more than just a seminar.

One of the last things I'll mention, the Secretary General and the team and I have been discussing the momentum from this session. I'm very cognizant, because I go to a lot of conferences myself, of the excitement that is generated at a summit or discussion like this. And I think it is -- two things I would say, to set expectations, it is not easy to move forward after a summit like this even though expectations are high and it is exciting, but I'll tell you I think there's strong commitment on things we can do together. XPRIZE is committed as a nonprofit to work with the UTI and the team here. Many of you who are here that we have been sitting with. We are committed, we have money we can invest in this. One of the things that we are working on currently which is relevant, we are doing a future of housing roadmap or impact map. We are trying to define what the future of housing needs to look like in ten, 20 years to be healthy, safe and affordable.

Dr. Paul Bunje is leading that for XPRIZE and the function with Xenia Tata who was supposed to be here. They are going to be in India next week with a lot of experts on the future of housing. And Lowe's Home Improvement has invested as underwriter two and a half million dollars for that initiative. What I've suggested to the team here is that one way we could take this forward is to tie it to a current initiative and bring this community together in January, February, March, to actually go through a process where we look at the future of housing which hits about four of these SDGs and look at how tying that to AI would be a way to keep momentum here in a practical way. Ideally the future of housing will lead to grand global competitions, XPRIZES that we launch so we see the fruits of our labor, seeing it manifest as something tangible in the future.

Marcus, we had breakfast this morning. There was a suggestion that we look at a global competition XPRIZE around machine reading. That would be a benefit to this committee if we could leverage the global crowd to do something around machine reading to progress the industry. I know that the Secretary General and I talked about if we do something moving forward, maybe we should add a couple other key disruptive technologies like block chain for good. Crisper for good, AI for good. It's more like a trifecta. There are some things presenting real risks and benefits.

Can you believe it, we are five minutes early after three days. Thank you very much for coming. Any final words, Secretary General?

>> HOULIN ZHAO: Yes, I have.

>> MARCUS SHINGLES: The five minutes is yours. Thank you, everybody.

(Applause.)

>> HOULIN ZHAO: Yes, fantastic. This is really good. My colleagues said it is marvelous that this late hour of Friday you are still here. I add a few observations over the last three days I invited a few people to have breakfast, to have lunch. I found that Marcus sometimes comes late, or earlier because they are interested in the panel rather than staying in the restaurant to have food. That's great.

So I remember seeing Rupert at the ITU telecom in '79 and I started the event telecom since 1971. I myself was in '79 I saw Rupert, but they are not at all like Sofia who captured our imagination yesterday. Nowadays we are talking about how AI can transform our future to the better life. But in fact we are also talking about how we can transform AI's future for the greater good. More than 470 speakers and participants attended this conference, close to 5,000 people live streamed the event online.

Through social media and the international media attention generated by the event, I should believe that figure. More than 100 million people have seen, heard, or read about our force to utilize AI for the benefit of humanity. It would be too hard to recap everything that has been said, heard, or Twittered over the past three days. Well talked about innovation. Education. Disruption. Standardisation. Even building. Privacy, safety and ethical issues were an important part of the discussions. So was the need to democratize the development of AI. Especially when it comes to debt. For some speakers AI holds the promise of exciting new possibilities. Others were more cautious, mindful of the challenges that lie ahead.

As Director General of the world organisation said AI is a new frontier. We don't have the answers to many questions yet. And we are not even sure we know all the questions that needed to be asked. But all participants covered, converged on one point. We need a platform, bringing together governments, industry, and other stakeholders to discuss those complex questions and generate ideas that we will allow us to use AI tools for good.

Over the past three days many of you urged me to keep the momentum going, including our outstanding speakers at this last meeting.

I will update the Secretary General of the United Nations, Mr. Guterres, about the discussions that took place during this conference. We will share the outcomes from the other events

such as WSIS Forum that starts right here next week. You are all inviolated if you plan to stay in Geneva.

Ladies and gentlemen, we are at an inflection point, all of us together need to create a good environment for AI innovation to flourish. AI gives innovators the tools to allow them the capacity to innovate. Let's give them the opportunity to use AI to help solve some of our greatest challenges, starting with the SDGs. That is why it is so important that post developed and developing economies, multinationals and SMEs, universities and startups, all join forces. I'm pleased to advise you that to further promote AI for good, ITU has decided to put a theme for the World Telecommunication Information Society 2018 with Enabling the Positive Use of Artificial Intelligence for All. Of course, we invite you to join us next year to promote this event.

The first special issue of the new ITU journal, ICT Discoveries, is dedicated to AI, which we will issue in 2017. The deadline to receive Articles is not closed yet. I cordially invite you to submit your Articles to us.

Let me thank the XPRIIZE Foundation, our sponsors, the Kay Family Foundation, Wiki Omni, Word4App and Price Waterhouse Coopers. I want to thank all the U.N. agencies that give us their support, the committees that helped to organise the event. In particular Amir Banifatemi from XPRIIZE and Stephen Ibaraki. My special thanks to the speakers and moderators, the participants of you. And ITU staff who have been engaged in supporting this event.

Let's give them a big round of applause.

(Applause.)

>> HOULIN ZHAO: This is the end of this event, but it is the beginning of ITU's cooperation with you. At ITU we hope that you will continue to play an important part in this conversation and dialogue and that you will join us to discuss how we can continue together on this journey.

Whether by participating in some groups as we just heard or by correspondence like email to share our views. We all need to come together to encourage and support the development of artificial intelligence. Last word, have a safe journey home. I hope each of you will return home full of new ideas and a renewed sense of purpose and possibility.

Thank you very much. Safe journey back.

(Applause.)

>> Can I just formally announce that this is the end of this wonderful event. Thank you very much.

(The event concluded at 1800 CET.)

This text is being provided in a rough draft format.
Communication Access Realtime Translation (CART) is provided in
order to facilitate communication accessibility and may not be a
totally verbatim record of the proceedings.
