FINISHED COPY

ITU AI FOR GOOD GLOBAL SUMMIT
POPOV ROOM
PLENARY 5
"BREAKTHROUGH" PROPOSALS ON PRIVACY, SECURITY, ETHICS AND SOCIETAL CHALLENGES

GENEVA, SWITZERLAND
8 JUNE 2017
1300 CET

Services provided by:

Caption First, Inc.
P.O. Box 3066
Monument, CO 80132
1-877-825-5234
팕719-481-9835
www.captionfirst.com

***



***

>> REINHARD SCHOLL: Thanks and good afternoon, ladies and gentlemen. One, two, three.

Can I invite colleagues who are going to report on the breakthrough groups on stage? Okay, Joe, I think you are ...

You're all set on stage? Andy? We will have slides, yeah.

Andy and Joe. Are we ready?

Okay. So good afternoon, welcome back. So this is the session where the breakthrough groups that have been meeting before lunch are going to report back to Plenary. And we have the rapporteurs or leads summarizing their proposals. So you will see for each of the breakthrough groups you will see one slide which explains the two or three proposals that the groups have come up with. So each of you has about three-ish minutes in time to do that. Then maybe Marcus and I ask a follow-up question and then we go to the next one.

We would -- getting to that. Next slide, please. So we would like to get your feedback on these proposals. So in order to get the feedback, please do the following. Go to the application, to the app. Visit the AI for good event app. Click on schedule.

>> MARCUS SHINGLES: We will share where the AI for good app is.

>> REINHARD SCHOLL: Find the respective breakthrough session. Then scroll down to the bottom to find survey. So scroll down to the bottom to find survey and then answer two quick questions. I'll say this one more time. You go to the app. You click on schedule. You find a breakthrough session. You scroll down to the bottom where it says survey and then you answer the quick questions.

Yes, Richard?

>> (Speaker away from microphone) -- require a Google account or something to use which I refuse to do. Can you get the app without some kind of login or do you have to have a Google account?

>> MARCUS SHINGLES: In the app for the feedback, we are not calling it voting because obviously we are not going to be voting on something here that we spent 90 minutes curating. So the objective of this exercise is with these groups, it's all the breakout groups that we do over the next two days we have 90 minutes. There's only so much you can accomplish in 90 minutes. We realized going into this, in those 90 minutes, having a group, taking this group and having you go into the different breakouts, that some of you would actually have a lot of difficulty getting alignment on the guidelines. That there's a parallel universe happening in the room because there's different perspectives, people from different disciplines. That is part of what we wanted to learn.

So the process itself is what we are trying to do to learn from it. We didn't expect we would come up with specific guidelines out of the exercise. The objective was let's get initial reaction across these major topics that we think are important. Let's get some of the feedback. Let's see if that feedback is all headed in the same direction or headed in different directions or even coming and clashing with each other. That's okay.

When the teams share out, they are going to talk about the things that as we pass this summit, as we build a community to build on the concepts we are building on today, how complicated is it going to be, what type of messages, what types of perspectives are being included to get a sense of the order of magnitude of trying to come up with guidelines around these topics. So it is meant to be a little bit of a frustrating

process because everyone is coming in with did different perspectives. The teams will share with you what they formulated and if they had trouble formulating. That's part of the process.

>> REINHARD SCHOLL: So the feedback we close and tomorrow at the afternoon Plenary, we give a high level feedback.

>> MARCUS SHINGLES: You can almost think of the feedback as hitting the like button. You get feedback, it's a Pulse, that's all that is meant to do.

>> REINHARD SCHOLL: So next slide, please? So this was on the breakthrough on.

>> MARCUS SHINGLES: Enhancing privacy and security. Okay. Do you want to give an overview of what generally you felt the group perceived as what privacy and security even meant? Contextualize a little bit and share what the guidelines were?

>> SEAN McGREGOR: Sure. I think there was broad agreement that privacy and security is important, but how we put that into action is where we were coming into considerable disagreement and uncertainty in how to proceed. So we had an idea that the artificial intelligence really benefits from data and we require broad access to the data to build as powerful model as possible but we don't know how to trade that off with the negative side effects of disclosure, of all the private information that enables it.

So the first item. I don't know if you are able to see the screen in front of me, but the first item is the necessity of continuing the conversation going into the future and to bring together all the stakeholders as we started here today to drive more towards that consensus.

There's a lot of disparate organisations out there that have different viewpoints on privacy, both in terms of how people are put at risk and in terms of how privacy concerns stimulate or prevent the development of important new AI technologies. Bringing stakeholders together to drive those compromises and figure out the good middle ground is something that we need to do.

So the proposal 1 is labeled as continuing the conversation, to assign, identify or convene a world governance bold or group to lead or coordinate on security and privacy issues and develop these concepts at the international level.

>> REINHARD SCHOLL: Sean, was there, just to get a sense of the agreement, was there sort of agreement on this or different voices whether this is needed, world governance body?

>> SEAN McGREGOR: We didn't have much of an opportunity in the session to actually voice disagreement with the proposals, since these were all summarized at the end of the session. So I

think that we'll discover any potential issues vis-a-vis application that everyone has access to.

>> REINHARD SCHOLL: All right.

>> SEAN McGREGOR: So moving on to the second proposal, which would be what this body could potentially do once convened or potentially do this independently is to create model laws, regulations of some sort on security and privacy that is representative of this consensus and then encourage countries to adopt those independently. It doesn't necessarily need to be at the international level since there's limited enforcement capacity there. It could be adopted by individual countries.

Also the corollary to this is the potential for a broad international agreement that could be developed. But that would take a lot more work than this, which could be more immediately achievable at the convening of some body.

>> MARCUS SHINGLES: Was your feeling that these were U.N. roles? When you say governance body, by the U.N.? Or by industry? Or by a consortium? What was the thought on that?

>> SEAN McGREGOR: I hesitate to go too much into that at this because of wanting to be representative of the group. And we didn't get to that level of granularity. I imagine that it would be largely consistent or my read of the group would be that they would be encouraging the U.N., at least playing a role in it.

>> REINHARD SCHOLL: Okay. All right. These are the two proposals.

>> SEAN McGREGOR: There is a third one.

>> REINHARD SCHOLL: Okay.

>> SEAN McGREGOR: There we go. So this is one I labeled as creating public goods. This is where we identified the promise, but also the insufficiency of a lot of privacy enhancing technologies or approaches in terms of differential privacy, multiparty computation, homomorphic encryption and others and that there's a need for strategic investment in creating these public goods and that can be done at the international/national, or the national level. And it helping everyone in the course of developing these and being able to both produce AI systems that are useful and also still have privacy and security.

>> REINHARD SCHOLL: It is a bit complicated at the end here.

>> SEAN McGREGOR: Yup.

>> REINHARD SCHOLL: Okay. These are the three proposals. Let's try to get your feedback. So again to repeat, it's just to get a flavor for what you think, whether this is something that could be carried forward.

>> AUDIENCE: What is the point of the app?

>> REINHARD SCHOLL: Yes, sir?

>> (Speaker away from microphone.)

>> REINHARD SCHOLL: You have to wait until the red light goes on. You click the button. Press it once and just wait. Wait a few seconds. It will get on.

>> AUDIENCE: One and two seem to me horrendously general. They could apply to every group that has met today.

Item three I think is a bit more specific. But one and two look very, very general indeed.

Are they specific to privacy and security?

>> REINHARD SCHOLL: I think we are going to pretty much in realtime, the groups have been meeting for 90 minutes and scrambling to get some text. That's what we have. And let's just get your feedback. You say yes or no and that's what we have.

We need to rush so that the other groups also have time to present their proposals. Let's go to the app.

I'm not taking any interventions from the floor. We don't have enough time here. So let's just go to the app and just to see, get a flavor or show of hands who is going actually to give us feedback, who is going to do this on the app.

Okay, all right. Good. So we will get some statistics.

>> Again, just to be polite here too because for this to be presented and not take a lot of feedback and questions on the spot here is not because the feedback is not important. The notion here, we have to have some initial starting point. It could even be a starting point that we come later and we disqualify, but there is not a lot of progress we can make in 90 minutes or even in a day. The goal here is to put some notion, posit some notion and then post-summit we will convene smaller groups to think through that. It might be acceptable when you think through it, of course, to even iterate them to where they are disqualified, some of these are replaced with other concepts. It's a starting point. That's all we are trying to do here.

>> REINHARD SCHOLL: It is supposed to be a fun exercise.

>> It's supposed to be fun. Oh, yeah, it's supposed to be fun, too.

>> REINHARD SCHOLL: Let's move to the next group, ethical development of AI. Andy, please.

>> ANDY CHEN: Our moderator is Robert Kirkpatrick. So we had a very active discussions. I think they were a lot of common themes coming out. And I think we get at least 20, 20-some recommendations. We filter out the top two. Those resonated with all the people in the room. The first has to do with transparency. That is raised by -- a lot of people are thinking you need transparency on the design of the AI. Should there be a consequence of errors, how do you trace it back and

what degree of transparency you need to have. In the engineering world that's what we do. Transparency on your design and monitoring. How do you monitor and collect all the data that you need so you can do the investigations for anything that will follow. Transparency is one of them.

Then there's degrees of transparency. At a minimum, if it is not critical but it could be meticulous transparency. That is Wendell put it, that would be a critical system. That's the first proposal.

>> REINHARD SCHOLL: Okay.

>> That's actually a recent topic. There is an ex-employee from one of the major search engines coming out talking about how sophisticated the AI is on the search and how it could lead to some unethical understanding of using your search.

>> ANDY CHEN: Right.

>> Second one?

>> ANDY CHEN: I do the second one. Proposal 2 is coming up.

>> It's on the screen. Below there.

>> ANDY CHEN: Okay, I didn't read the bottom one. The second one comes out loud and clear is this, when we come up with all these ethical guidelines, we got to be mindful about it is just not benefiting the particular sentiment of the populations. It has to be all inclusive. It has to be consulting with the Developing Countries or even the impact on the countries where they lose the jobs, the impact that they have with this AI being out, I think that's a very important issue for us to include in our consultations.

>> REINHARD SCHOLL: Have you had discussion on how to engage them proactively?

>> ANDY CHEN: We haven't gone that far. There was really a lot of ideas. I think this is a good thing that we want to go into further, for sure.

>> Is that the third one? Just the two?

>> ANDY CHEN: Just the two.

>> REINHARD SCHOLL: Okay. Let's go to the feedback.

So we'll give you ten seconds to give us feedback on that.

And then we go to the next one. To the next group. All right, if we move to the next slide. Future of Work. Alexandre is going to present.

>> ALEXANDRE CADAIN: Hello, everybody, even if this is the first formulation of our proposals, let's say a starting point for further development, I wanted to share quickly how we actually got there. The idea was to imagine actually how the Future of Work will be and we wanted to try and work beyond the usual fears and dreams, actually, regarding obviously the Future of Work and we tried to set a realistic view on what is actually

happening today and how we could actually prepare ourselves for the future.

So in order to get to those points, we asked our great panel five different questions. I am going to share with you. First it was what is the best case scenario for the future? Then how do we make sure humans will be prepared for the uncertain futures of work and position to play a meaningful role?

What are the opportunities and risks we foresee with managing human talents? How do we mitigate the risk of completely replacing human contribution with machines in Developing Countries? And last, how do we mitigate the risk of engendering further inequality.

Each level led to several guidelines. We wanted to narrow them to two, which are for the first one, encourage social dialogue to continue and consider what data should be open source and what skills are required for the preferred future of work.

>> REINHARD SCHOLL: Could you explain what is meant by open source? That term is not always clear.

>> ALEXANDRE CADAIN: The term? Open source? All right. The idea here was to actually understand that we need great transparency towards data. A point was made by several members of the panel that actually data today belongs to few actors. We wanted to make sure that to avoid further inequality in the development of AI we would make sure that more people would have access to useful data in order to promote AI development.

>> REINHARD SCHOLL: Okay. Good.

Then the second one?

>> ALEXANDRE CADAIN: The second one is about promoting the maintenance of infrastructure in which we include governance, to democratize data, knowledge, and skills.

>> REINHARD SCHOLL: So infrastructure. You said that includes government? What else does it include? Maintenance of infrastructure?

>> ALEXANDRE CADAIN: The idea again was to make sure that both in developed and Developing Countries that we actually can develop an infrastructure that is able to give access to data, knowledge, and skills to the majority of people and again, the idea here is to promote its birth on how to design it in several countries and to assure its maintenance.

>> REINHARD SCHOLL: I guess I'm biased because I work in ICT. When I hear infrastructure, I think of network infrastructure, but you think of more generally the infrastructure. Okay, good.

All right. So we would like to have your feedback on that. So we give you a few seconds to say what you think of these proposals. First one to encourage social dialogue. To

determine continuous, what data should be open source and what skills are required. And the next one, promote maintenance of general infrastructure to democratize data, knowledge, and skills.

With that we go to the last but not least breakthrough groups on Humans and Machines.

>> JOSEPH KONSTAN: It doesn't take AI to figure out I'm next. We had a wonderful session. I see many of you out there, that's great. The rest of you missed out. We looked at broadly issues around human centricity and the diversity of humans. We heard about where there needs to apply AI technology for literacy in rural India. We heard about great potential in medical diagnosis. We heard about saving people in radiation cleanup. And from that we first got this to six proposals, did our own round of voting and I did some painful merging to see what you have here.

You will see overlapping themes. One is human centric accessible transparent design. The idea that we need the systems that encapsulate AI to deal with cognitive literacy, language, to be cultural sensitive to where they are deployed. If they don't understand power structures, they can be ineffective, insulting or worse. And to provide that transparency and access to internal models and learnings that you've already heard about. We are wrapping that all into one sort of mega-proposal around a set of principles.

Number two, encoding ethical, legal and human rights criteria into the AI systems themselves, fail safes that are wrapped around those systems to prevent harm, and the evaluation criteria that are being used. While we recognize that the U.N. and others have articulated some universal principles around human rights and other topics, these can't be limited to universal principles and have to be anchored to local values in the places where these different systems are deployed. That is going to require an interdisciplinary engagement of philosophers and ethicists and social scientists and representatives of the disenfranchised communities that these systems are designed to serve.

Slide next, number 3 there for governments, NGOs, industries, you name it to get out and prioritize will these. We have a longer list of these. At the top of the list were a few of the things you might imagine that include applications that have the chance to save human lives or enhance health, safety, and education. Applications that have the chance to deliver education, wellbeing to the most underserved populations around the world. Some of these are things the market will take care of. Many of these are things the market will not take care

of unless the governments, the foundations and others around the world make them a priority.

Those are our three proposals.

>> REINHARD SCHOLL: Okay, thank you. Can we go back to the first slide so we can take a look at it? The first is human centric, accessible, transparent design.

The second, ethical, legal and human rights criteria in AI systems. And then the third one, go back to the third one, articulate and fund high priority areas.

You said, Joe, the market may take care of some but not all of the others. Can you give a couple of examples where the market may not take care?

>> JOSEPH KONSTAN: I think in the developing world we had a number of examples that came up. The one that is top of mind is we heard about how roughly 50 percent of women in rural India lack the literacy or even basic grade school level education. And the availability of teachers and the delivery of education there could be immensely enhanced by AI technology. It is not clear that certainly these women are not going to be the ones who can pay for it. It is going to take concerted governmental, foundation, or other efforts to say: Wow, this is the kind of challenge.

In fact we were just talking about the fact that this is one of the great things that the XPRIZE Foundation and some of the other foundations that are out there trying to advance applications of AI for positive good are seriously looking at and taking steps in.

>> REINHARD SCHOLL: Okay, thanks very much. We would like to get your feedback on these three proposals. And I think in general we can applaud the breakthrough groups. They were the first to do that. Thanks a lot.

(Applause.)

>> REINHARD SCHOLL: The next group will have it is even a bit more difficult. There's a coffee break and then the breakthrough groups are going to meet and you come back immediately and we are going to present the results. Thanks a lot for this one.

Before we move to the next Plenary, Marcus is going to tell you about some of the results of the XPRIZE competition.

>> MARCUS SHINGLES: There's a few groups here we want to recognize. We'll do two later. You saw creative art in the hall. That was competition done with the IT of the MIT Media Lab. We'll talk about that in a bit.

I do want to recognize somebody here today. This is an example of how AI is being used for good, where we are crowdsourcing innovation. Let me spend two minutes on this and I'll introduce base I will. XPRIZE a few months ago announced a

prize with Qualcomm, $10 million Tricorder -- Does everyone know what Tricorders are? From Star Trek, the medical device? It was based on that notion.

Rod Rodenberry, Gene Roddenberry's son, the creator of Star Trek, was part of our team. This was inspired by the Star Trek medical device.

This was the criteria: Develop a medical device that weighs less than 5 pounds that can diagnose you against 13 disease conditions better than a board of certified physicians. We had the competition run for five years. We worked with UCSD, University of California/San Diego. We had the FDA, Federal Drug Administration involved to oversee it. A rigorous process. We had independent judges and 300 teams signed up five years ago.

Those 300 teams got down to six earlier this year, late last year. We got down to two this year. We just awarded the $10 million competition about a month and a half ago, of which two and a half million dollars went to the winning team, which was Dr. Basil Harris. If you want to stand up, Doctor, for a moment.

(Applause.)

>> MARCUS SHINGLES: If you want to come up here real quick, Dr. Harris. So the final solution that took five years to create, this is an interesting story. He's actually an emergency room doctor in Pennsylvania. And him -- is it two or three brothers? Him and his two brothers are all ER doctors.

>> DR. BASIL HARRIS: No.

>> MARCUS SHINGLES: Get my facts straight here.

>> DR. BASIL HARRIS: One is a programmer. One is an electrical engineer and urologist.

>> MARCUS SHINGLES: One is a programmer, electrical engineer and an ER doctor. And the need of the device was for his own personal use. When he has patients come to the ER, he wants them to come with data. I think probably, we talked about this. That was probably the motivation why he did so well. And they use artificial intelligence. Why don't you describe -- he 3D printed the parts.

>> DR. BASIL HARRIS: Here is part of the Tricorder right here. Thank you, Marcus, and thanks to the XPRIZE Foundation and the Qualcomm Foundation. This was the inspiration from building the Tricorder. Our device and all the finalist teams did amazing jobs building a Tricorders. And my dream in the future as an ER doc is that a patient will be able to come in to the department and come to me with a diagnosis from their home device with reliable information, real vital signs, real data that I can trust, and then move to the next step. That's where we are going to start seeing real efficiencies in the future.

Thank you for that.

>> MARCUS SHINGLES: Name a few out of the 13 disease conditions so people have an understanding.

>> DR. BASIL HARRIS: This was a cool demonstration project that spanned two illnesses, pneumonia, urinary tract infection and chronic diseases like COPD and diabetes. They span multiple body systems, acute diseases, chronic conditions, really to show what a system like this can do.

And we are already moving forward with getting full FDA clearance on the components of our device, as are the other teams that are still moving forward. This technology is coming. I think it is going to empower a lot of people around the globe.

>> MARCUS SHINGLES: What was would be example of where artificial intelligence was used in the device?

>> DR. BASIL HARRIS: The Tricorder as envisioned by the XPRIZE, building a simple Tricorder wasn't enough. You had to build the AI of the physician into the device. It is not on the show where Dr. McCoy examines you and looks at you and makes the decision. The device is making a diagnosis, seeing what is happening and what is important to make the diagnosis. The artificial intelligence of that decision making is controlling the whole device and plus all the artificial intelligence, machine learning of the signals coming in to make small decisions on the way. Very cool.

>> MARCUS SHINGLES: So that was the winner of the Tricorder, Qualcomm XPRIZE. Thank you.

>> DR. BASIL HARRIS: Thank you, Marcus!

>> MARCUS SHINGLES: I know you brought your daughter here. She's how old?

>> DR. BASIL HARRIS: I would like to thank Mira my daughter, helping me to demonstrate Tricorders out in the lobby.

>> MARCUS SHINGLES: Come up and visit them. What time are we coming back now?

You're not leaving? Transition?

>> REINHARD SCHOLL: So we move immediately to the next Plenary Session. Thank you very much, everyone on the panel here.

(Applause.)

***



***