

FINISHED FILE

JUNE 9, 2017

1330 CET

ITU

AI FOR GOOD GLOBAL SUMMIT

PLENARY 10: INVESTMENTS, ECONOMIC ASPECTS AND DESIGNING THE  
FUTURE

BREAK THROUGH GROUPS ON INVESTMENT AND ECONOMIC ASPECTS.  
IMPACT INVESTMENT. PLENARY 11, BREAK THROUGH PROPOSALS ON  
INVESTMENTS, ECONOMIC ASPECTS AND DESIGNING THE FUTURE  
PLENARY 12, CLOSING SESSION, APPLYING AI FOR GOOD

Services Provided By:  
Caption First, Inc.  
P.O. Box 3066  
Monument, CO 80132  
800-825-5234  
www.captionfirst.com

\*\*\*

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.

\*\*\*

>> ANDY CHEN: Hello, everybody. Could you take your seat,  
please? We are going to start right away. My name is Andy  
Chen. Some of you might know me but I come from IEEE Society.  
I actually spend most of my life working as an engineer, working  
in the power plant, energy sectors. I retired a few years back,  
as before I became the CTO for large corporations. A little bit  
about myself, and I also recently become partner for venture  
companies, so with different background, organiser thought I  
could do this session and get speakers to talk to you to set you  
up to go to the breakout sessions.

Let me start with, can you flip on the slide, please? The  
presentations? Okay. Ready to go.

For this session, next, please, excellent, for this session  
we talk about investment, economic aspect, and designing the  
future. We have two speakers. One is virtual, and one is in  
person. I will introduce them to you shortly. So can you  
shrink it down? The screen? Is idea of this session is trying  
to get all the different sector work together, government,  
industry, academia, private sectors, and the civil societies, to  
mitigate the risk that we have been talking about for the last  
couple days, that may be posed by AI, and ensure that we will

make the AI for good for human beings, right?

These sessions, this session is the session we are trying to achieve the goal or discuss about it, in the breakout sessions, so that we could bring back the recommendations or the guidelines as we call it, guidelines and recommendation are the same thing. We just tell it differently but when you come back from the breakout sessions we expect you to, the Rapporteur will table the recommendations. Next slide, please.

Maybe I can do it. Okay, good. The screen is a little bit big. It only show half the screen. But I'll tell you what is on the screen. I want to give a little bit idea about IEEE memberships. We have about 400,000 members worldwide. One more smaller. That's great.

We have about 400,000 worldwide. We are the communities for technology leaders. We have all different fields. AI has been a emerging field for us, a lot of focus has been done for us in the area. Recently we published a paper for ethical design alignment for AI. We are looking for comments and input. We had a couple people here, so if you are interested, let me know. We will send you the information to participate in the public consultations for the alignment of designing of AI and AS, autonomous systems.

We are truly global. We have about 550 chapters around the world, and we published standards, standard, that is what we are known for. We have technical community as I said. We publish about 30 percent of the relevant technologies, publications around the world. So that just gives a little idea. One thing I want to point out to you, 42 percent of our members are industries. They are people, professional people graduated and have the careers in professional engineers, electrical engineers, system engineering, computer engineering, as well as electronic engineering. We have about 30 percent of academia, the leading research around the world, and about 11 percent, 10, 11 percent of government agencies that work with us.

That is to give you an idea about IEEE Society. I'd like to move on and introduce you -- can you turn it off? I don't need it anymore. Our first speakers come from Microsoft. Eric Horvitz is the technical fellow and the Director at Microsoft research. His field of interest is quite wide. He got involved in healthcare, transportation, commerce, aerospace. So he is involved in various type of researches sectors. He has a very extensive academic background with many degrees, as well as being the fellows in many professional organisation like ACM, and advanced AI associations. Without further ado, we will bring up Eric Horvitz.

(applause).

(inaudible; I'm sorry, I can't hear or understand.)

>> ERIC HORVITZ: For harnessing AI in the world. I want to point out my experience has been, it takes quite a bit of effort.

(very faint audio).

To engage, AI experts engage deeply and almost become experts and often success takes collaboration with government, industry, academia and civil society. If I can have the next slide, make the comment that there are different models of engagement. I've experienced quite a few over the years. I have trouble generalizing. It has been quite a bit of custom tailoring that goes on with any application. Next slide. I want to point out that for success, I wish I could have been with everybody this week but I had to be here in the U.S., I find that we need to have coalitions around expertise, infrastructure and data. Expertise is very important for even selecting the problems, identifying pain points, understanding possibilities in the open world outside the computer science laboratory.

We often find quite a bit of infrastructure we can harness, human resources, if we engage deeply with domains, the actual specific areas like healthcare, for example. There we find access to data, often hidden, often part of the dark matter of a domain. If you want to go to the next slide, I want to give a couple of directions.

One comment is that I see the word AI used a lot on the programme, regulation of AI, AI safety. But AI is really not one thing. It is not one blue green gas, it's a set of rich sub disciplines and methods, vision, perception, speech and dialogue, decisions and planning, robotics and so on. We have to consider all these different disciplines and methods in seeking true solutions in delivering value to human beings and organisations.

Next slide. One way I think about a beautiful way to, perspective on developing applications is the pipeline, data to product distributions to actions, data to predictions to decisions and back up again, once you set up a pipeline in the domain, to think about how much data is worth, how much more data would be worth, how much better could we do with more information.

Next slide. Next slide. It's often not about automation but about use cases where you have decision support, it is not about autonomy necessarily, human computer richness and collaboration. Next slide.

Let me take an example here of high path areas, this is examples we have been involved with over the years. Take cholera which is a big challenge, 100,000 people estimated dying per year of cholera. I want to point out this model of domain knowledge, expertise and models of engagement. Next slide. It

turns out that with getting appropriate hydration therapy to people in epidemic areas with cholera, you can go from 50 percent mortality rate to a 1 percent mortality rate.

The idea of getting fresh water distributed on time and in the right place at the right time is critical. Another interesting decision problem is there are short acting vaccines for cholera that don't last long. If you have predictions about where cholera might be occurring on the horizon, you can actually direct in a efficacious way where vaccine programmes should happen every couple of years. Five years ago we developed models to predict where cholera would come months in advance, weeks in advance and days in advance, power up planning algorithms that could be used to distribute fresh water and actually as part of design programmes for figuring out where to apply vaccines over time.

Without that kind of predictive modeling which came from multiple heterogeneous data sources, large scale data sources, we would be in the dark. We act in realtime to a epidemic. We don't understand how to optimally set up the plumbing infrastructure to get fresh water to the right place. Next slide.

As an example, in reaction to a epidemic in Haiti, the World Health Organization sent a million doses quickly in response, but it's too late and too expensive to deal with rich predictive modeling, this kind of example of a application area, understanding its value and understanding where methods available today could be applied comes from the engagement of domain experts and AI researchers.

Next slide. It turns out that we have incredible assets if we look for them. For example, this gives you a snapshot of what western Europe in this case North America, but they are all very similar these days, looks like in terms of the number of planes flying over the continent right now. We worked several years ago to engage the FAA, U.S. government regulatory agency for flight, to gain access to ground radar tracks of airplanes. We showed how with a rich machine learning model, you can take all these airplanes as sensors and build rich weather maps for the country and we are showing how it applies to western Europe and Asia.

Next slide. You can build a live service, in this case we call it windflow, live cloud service we joke, that is dynamically updating wind more precisely than the agencies that currently exist based on balloons, all by leveraging existing infrastructure and asking the FAA for data that should be available for human welfare.

Next slide. Here is another example of several years ago we looked at earthquake data and models for how we might help

people deal with disruptions, and understand disruptions and need for surveillance looking at this earthquake that hit Rwanda in 2008.

Next slide. It turns out that there are, there is a cell tower system in Rwanda that could be used as a sensor network, to work with the Rwanda government which provided anonymized cell tower data of three years of cell tower calls. That is 140 cell towers, ten and a half million calls over six days.

Next slide. We could show with AI related technology, planning, decision-making, machine learning, we can use the cell towers, to look at the bulk of the calls that come in around a earthquake to predict the epicenter of the earthquake, within 17-kilometers, outside the country, in the Republic of Congo, from just the cell tower calls, human call level changing, over a region of 140 cell towers. Next slide. We then showed how you can look at disruptions over time in call densities to compute opportunities for systems based on persistent disruptions. Next slide.

Then show with decision theoretic planning model, given the uncertainty in our inferences about where would you most assist, provide help to people potentially in need after a disruptive phenomena like a earthquake, where to look at a triage, where to send surveillance resources, understand more, to ideally collect the uncertainty. Next slide.

We need to start looking at the community, piecing together coalitions of governments, private sector, civil society, to build resources, with the AI development resource of sets of data for use in leveraging for AI for social good. Next slide.

You get a sense of the kind of data that are available right now. These databases can be built and shared widely, and experience infrastructure can be shared widely as well. Next slide.

I want to end by saying, there is so much to do, most of it could be done or quite a bit of it with AI methods available today, the constellation of that method which is a rich fabric of methodologies requires us in leveraging opportunities to work deeply with experts, understand where the pain points are, and where we can get maximal benefit for people. I'll stop for comments there. (pause).

>> ANDY CHEN: Thanks, Eric. Thank you very much. I have a couple questions for you. Can you stay for a little while? We will have the session of questions after Neil speaks and then we can take the question from audience as well. My first question is, where are you and what time is it there? (chuckles).

>> [inaudible]

>> ANDY CHEN: We should give him a hand. 4:00 in the morning. You are a trooper, Eric. We appreciate that. The

second question for you is this is fascinating examples that you give us. What kind of challenges do you face when you try to predict the weather or predict the earthquake? Could you give us an idea of what kind of challenges that you face when you do those kind of projects?

>> ERIC HORVITZ: It turns out (very faint audio).

There is no cookie cutter. We learn so much with every domain [inaudible] there is a long process, working with multiple stakeholders [inaudible] I would be absolutely shocked with [inaudible] outside the laboratory, real world [inaudible] premature or immature for the real world.

>> ANDY CHEN: What you said resonates with this session here. We talked about trying to convince government, convince the industries to be able to adopt similar technologies that you talked about, that is going to be a big challenge for any of the researchers or any events technology. Thank you very much. Please stay on line and we will get you back in ten minutes. Okay? All right. So let me introduce you our second speaker. I met him the start of the meeting, the event two years ago, he is my best friend now.

He's got this unique ability to connect with people. It's my pleasure to introduce to you Neil Sahota. Neil, to my left, is an IBM master inventor, one of the IBM master inventors. He is a global leader in IBM Watson. He has 15 years of working with Fortune 500 clients in many area, government, healthcare, generations, utilities, you name it. He has touched on those areas developing different type of project. He is one of the few IBM corporate leadership programme that actually appear with a NGO to perform community driven projects, places including in China where he did the work and on top of that, he is academia, he is on the board of alumni for University of California in Irvine. If that is not enough, he is also an athlete, he is sitting on the Board of Directors for Orange County marathons, with Orange County kids programme. So let's put a hand together for Neil.

(applause).

>> NEIL SAHOTA: Can you hear me? Good. How is everybody doing? Good. Can we throw my slides on?

While we are getting slides loaded, I am lucky enough that my portfolio is worldwide so I see a lot of great ideas, a lot of use of AI around the world. I can honestly say that in the span of probably two years, it's gone from very small areas like focused in healthcare and education, to actually span around to every major sector and industry today.

I can quote you a lot of great things like the World Economic Forum sees AI as a top six trend, Gartner sees perceptive machine learning as a top three trend. You got organisation

like ACM developing learning programmes. I can tell you that China, their investor groups are committing \$337 billion to AI investments. That's more than 80 percent of the GDP of the countries of the world.

There is a lot of hot things going on. I don't need to hopefully sell you on how important AI is becoming. You have gotten two and a half days of that. But hopefully -- it's not moving forward. Hopefully you are asking yourself a question, this sounds really good, I've heard a lot of great things. I want to use AI, I want to try and learning it. How do I actually do that? That is hopefully the question. I'm going to talk about the future of AI, and how you can actually best make it a part of your solutions.

Is my slide deck ready? I can do this old school if I need to. I don't know where my slides are. I'm going to go old school on this then. I'm going to help you figure out how you can be part of the future of AI by first talking about the past. Let me throw a date out there to you. June 29, 2007. Does that date seem familiar to anybody? Yeah, that actually is the release date of the original iPhone. The iPhone is going to turn ten years old later this month. You think about it, how much have smart phones changed our lives both personally and professionally? I actually read a survey saying that most people ages 20 to 29, if they were stranded on a remote desert island, what is the one item they must have with them, 71 percent said their smart phone. That was interesting, since you have no cell or Internet service and you can't charge it.

But it shows you how important, how a piece of our lives it's become. The real value of the smart phones has become really what is called the app economy. It gave a platform for the development of apps. If you look at companies like Uber or Snapchat, they couldn't have existed, couldn't have had these business models without the smart phone as a platform. I think today we had the same opportunity with artificial intelligence. We can actually change the game by using AI as our new platform to come up with new solutions, new products, new services. You are thinking, okay, I get that, I've heard a lot of that, but how do I actually do that. Right?

Well, to put it very simply, it's two words: Think differently. AI is not like our traditional computing solutions. We are calling it a third generation of computing. It works very differently. It is not about requirements anymore, it is not really about scenarios. It's not even about basic use cases. It is really about interaction, how do we now interact with machines. There is essentially infinite paths to make one thing even happen. We can't code to that. We can't plan scenarios for that. To actually leverage it, we have to

think differently.

Let me share a story with you. I met with a consortium of educators, they are trying to figure out how do we improve the learning experience, how do we achieve number 4 SDG goals about creating a more individualized learning experience. They had a list of 15 use cases, and they are working me through them, and I stopped them after an hour and said, yeah, AI can help you with a lot of these things, but I don't feel like you really are unlocking the real value here.

They looked at me and said, well, we kind of thought the same thing. We didn't really see what is so special about AI. The reason they got stuck was they were thinking as AI is a traditional computing solution. They were thinking how do I automate a manual process or how do I cut out a step or how do I do things faster. To unlock the power of AI, you have to think differently.

I asked them, let's figure this out. What is your major pain point? What keeps you awake at night? They came back and basically said, the challenge is when we teach students, they are in a class. Because it's a very diverse set of students, different backgrounds, different skill levels, different learning ways, we have to teach broadly across to try and reach everybody so we are not fully engaging the students. I got it. We can solve that problem using AI.

Right? We can leverage something like an AI tutor, to supplement what teachers, schools and parents are actually doing. We can do it in a individualized basis. You can use Watson to see what the learning capability of each student is, move at their pace and help them develop confidence and mastery of the subjects. Moreover, if they are interested in a particular subject or based on their personality traits, you can recommend things the student might like and be more interested in learning. Maybe instead of watching some cat videos on YouTube, you can encourage them to watch some videos on astrobiology or astrophysics. Actually try to do something that meets their needs. You can take it a step further and gamify it. Rather than make this I'm sitting at home with Watson, like I'm in the classroom, Watson can adjust the way it interacts with the student based on their age, competency level, vocabulary and emotional state.

They looked at me and they are like, we had no idea you could do that. That is part of the challenge. How do we actually do these things, right? We know that AI helps us create new insights, it can connect millions of dots together, even find dots we don't know about. It can help understand and relate in terms of emotionality for human beings. How can we do some of these things?



It's not just think differently, it's how do we actually do that. The first step is to wipe away our memory. We have to forget about what we know and focus first on the problem we are trying to solve. In a ideal world, you can figure out how can we try and solve that problem. Forget about if it's feasible or not. Let's think about how we can actually do that. The second thing to do is say let's think outside the box.

What can we do that is so radically different, what is that moon shot that we might be able to put in play. Take for example food. You want to teach somebody about how to cook, you are probably going to show them things about ingredients, show them how to cook, have them try different things, taste and experiment. How would you teach a machine how to cook? Machines don't eat. They have no sense of flavor.

What you can do is teach an AI chemistry. That is what we did in a partnership with Bon Appetit, we created chef Watson. Teaching Watson about chemistry, we are able to teach Watson about chemical combinations that produce good flavor but also nutrition. Watson can create original recipes using substitute or locally sourced available ingredients to create nutritious meals. Think about our goal today, have zero hunger, we can now leverage what is ever locally available to create more nutritious meals for people, and combinations we wouldn't have thought of as human beings.

These are the ways you want to think about artificial intelligence. If you are thinking about what can I do for my organisation, how do I get started, it's these two steps. The ideas become more important, the design, implementation has become simple steps. It is what is the right way to go. Throw another quick example, think about self-driving cars. They are based right now on how we actually drive as human beings, like on site. What if you want to think outside the box. What if you can leverage sound, help identify where cars are or children, or even smell, based on exhaust fumes. Could you improve the safety record or the way that self driving cars operate? Machines can do things that we can't. Can we leverage that?

My challenge to you guys is, think differently. Right? Think about what you can do to solve some of the problems you are facing. If you need focal areas, we have 17 great options right here in the SDGs.

I know that you are hopefully, you will take something away from this conference, some of you have gotten great information, you are probably thinking more tactically what can I do. Think differently. Think of some unique ideas, think of those moon shoots and Uber yourself before you get Kodaked. Thank you.

(applause).

>> ANDY CHEN: I have a couple questions for you.

The first one, you mentioned something about investment, this is about investment, right? You mentioned something about China, can you repeat what kind of investment globally is happening right now?

>> NEIL SAHOTA: China invested 337 billion or allocated that much money just for AI projects. I can tell you that is the most of any country in the world. They have got the most publications in AI. They are definitely at the forefront. But this is I think more representative of a way investments are going on generally. If I remember correctly, the IDC has forecasted that AI investment will grow at a compounded rate of 55 percent.

>> ANDY CHEN: The second question for you is, you mentioned about IBM Watson. And it's evolved over the years. We always talk about machine learning, deep learning. Describe a little bit about the journey of Watson, when you started how many years ago, how much it knows and how much it knows now.

>> NEIL SAHOTA: That's an interesting question. Watson actually started in 2006 where all great ideas seem to start in a bar. But the real challenge initially was how do you get a machine to understand natural language, because if you think about how we talk, it's really difficult for a machine to understand, because we use a lot of slang, a lot of idiom jargon. If I said, I'm feeling blue because it's raining cats and dogs, most of you know that Neil feels sad because it's raining really hard. But for a machine, it thinks somehow I physically turned the color blue because animals are falling from the sky. That doesn't compute, right?

We had to find a way for Watson to understand natural language, so rather than look at key words, we had to understanding linguistic patterns and grammatically break down sentences. That led to the forefront, we realized we could do more, we could have a machine mimic the human thinking process. Through observation, through its own experiences, it could actually learn like a human being and become confident in a variety of subjects.

After the Jeopardy challenge, Watson could only play Jeopardy and we were asking ourselves what we can do and realized we had a profound platform that we should open it up for anybody to use.

As a result, we started off in healthcare doing cancer research, partnered with Kettering Sloan, Mayo Clinic, World Health Organization. Since then it has branched out everywhere. Today we are working on project Lucy, where IBM committed a hundred million dollars to enhance the infrastructure of Africa, and we are using AI to basically help in agriculture, so farmers

can improve crop yields with less water, to improve education, to enable more healthcare service access because if I remember correctly, one doctor per 2,000 people in Africa, to help with in terms of energy, in terms of distribution, using renewable energy, even reducing energy consumption through devices.

It's now, Watson is in a space that literally is in every sector, every industry. A lot of what people are trying to do is commercial, and we are trying to push people towards how can we help use Watson and AI in general to meet the SDGs.

>> ANDY CHEN: Thank you very much. Can we bring back Eric? Have a seat. We will take some questions from the audience. When you ask a question, please give us your name and try to keep your questions within 30 seconds. We want to get more people to ask questions and we will try to give you answers. I'll take three questions because we want to get out to the next sessions. Anybody? Go ahead, please.

>> Hi, World Economic Forum. You mentioned Lucy and deploying this technologies to farmers in rural Africa. How do you actually do the education for them to be able to use these technologies, to understand and make sure that they don't make a wrong choice based on wrong reading of the data?

>> NEIL SAHOTA: That is a great question. I hope this is on. That is a great question. What we learned is, it's about the user experience. We leverage design thinking, thinking how would a farmer in a very rural area be able to use this knowledge. So we basically synthesized it so that Watson speaks to them in language or terms that they actually understand. Rather than go into, hey, Watson, analyze the soil content and I see this percentage of this mineral, this percentage of this mineral, Watson tells the farmers, if you shifted your crops over two meters or two millimeters, you improve your yield by X percent. You don't need to water as much over here, based on weather forecast, you should try and plant things facing this way. Again it's all about trying to use language that the farmer or user in general would actually understand.

That is the real power of AI, is that they can actually adapt the way they interact with us human beings to that regard.

>> ANDY CHEN: Two more questions. Anybody else? Go ahead, sir, in back.

>> AUDIENCE MEMBER: Hi, I'm Kenny Chen from Pittsburgh. Do you anticipate Watson or other AI systems reaching a point where they understand humans better than we do? For instance, we have talked a lot about human behavior, in many ways a lot of people see human behavior as the challenge or the problem to solve all problems.

So if we can understand the nuances of things like motivation and behavior change, perhaps we can tap into that kind of

butterfly effect or nudge mechanism of changing systems around the world.

>> NEIL SAHOTA: That is an interesting question. To be honest I'm not sure what the answer is. I think it's partially yes, partially no. There are things that we pick up as human beings and in particular we are not necessarily always rational, that I'm not sure a machine could adequately capture. But the other way around, we have seen things where machines can understand people better than we can. A good example of that is there is a lot of use now of AI tools in behavioral health. We were surprised to learn that people are more open to sharing information to an AI avatar, for example, than a human therapist, psychiatrist or doctor.

We learn that is because of the judgment factor, that even though that person is there to help you, they feel like they are somehow judging me, I don't want to look too bad, don't want to look like I'm scared or weak or whatever it might be. But when they knew they were talking to a machine they were a lot more open in revealing how they felt or what was going on.

In that regard, you can see machines have a better grasp or understanding of a human being than we might actually.

>> ANDY CHEN: Eric, can you hear us? Will you be able to add some more to that?

>> ERIC HORVITZ: I can hear, can you hear me?

>> ANDY CHEN: Yes, we can hear you fine. Did you get the questions?

>> ERIC HORVITZ: Yes, I heard one question, quick answer on the first question.

(audio is very faint).

First glimpse of the application of all the AI [inaudible] work on AI for assisting education, to help people understand the AI method. On the second question, behavioral economics and psychology is very important area for understanding. It's critical in understanding the opportunity for human [inaudible] how to best augment human beings, critical amount of [inaudible] education level. [inaudible] model persuasion, model [inaudible] cognitive psychology of judgment and [inaudible] giving that knowledge [inaudible] people.

>> ANDY CHEN: Thanks very much. I think we are running out of time. But I think the good thing is, you guys are going to break out to your sessions, and you can continue your discussions. So if you know which session you are, please proceed. And we will see you back after your session. Thank you very much.

(applause).

Thanks very much, Eric. We will start the breakouts right away. We don't wait until 3:30, so please go to your rooms and

then we will start right away. Thank you.

(pause).

>> PAUL BUNJE: All right. We are about to get started, the final group -- the final break through sessions of the day. You all have picked wisely, because this, if my colleague Jo Jo and I have done our jobs right this should be the most fun, I hope.

One of the things, for those of you who are not familiar with the inner workings of XPRIZE, you will quickly learn that we like to practice what we preach, so when we talk about crowdsourcing, we mean that in everything we do if you give me the chance to structure and moderate a session, you are all going to be involved. We are going to crowdsource from what we got right here. That is how we are going to end up running today.

Let me give a quick overview before introducing the panel of what we are going to do. So expectations are met.

This is going to be interactive. I'm going to ask all of you to join me in the next 90 minutes of an ideation session. We are going to be brainstorming some of the elements that might go into a roadmap for collaboration. In 90 minutes we are not going to be able to write a roadmap but if we can get some general principles of what should go into a roadmap, and if we can identify some particular milestones, some important things that we think might be on the path to a future state that we want to see, we will have been successful.

To do that, Jo Jo, can you hold up, if you have ever been part of XPRIZE, you know colored paper comes out at some point. Make sure you have a pen. The first thing I'm going to ask everybody to do is a quick brainstorm so we can get your brain moving and your legs moving. I want you guys to do this too but you will have after the brainstorm, we are going to ask our panel to play the role of instigators. Each member of our esteemed panel will offer a few minutes of insight, thoughts, provocation, about what it might look like if we are going to structure out a roadmap for collaboration.

Following that, I want to brainstorm out, I'll give you instructions in a second, two elements, and this is what we are going to report back, of the roadmap, number one, two guidelines, principles, what are general principles that should be included if you were to think about structuring a programme that reaches towards a goal that you saw up there, I'll bring it up again. The second round that we are going to do is identifying specific milestones, things that will allow us to know we are making progress on this broad goal. After each of the rounds, the panel is going to discuss and summarize what you all think might be the highest priorities associated with either principles or milestones in reaching towards the future. It's a

lot to do. I hope you are ready. Let me show you what the role is. Let me introduce these folks. I want to make sure you realize, a lot of them have already spoken and you heard a number of people, great folks. Sitting closest to me, Chaesub Lee here at ITU, one of the instigators and organizers of both this event and many standards and other aspects that make ITU so important. Next is Gary Marcus from NYU, he spoke on the first day, behavioral and cognitive scientist and researcher. Peter Marx leads the advanced projects efforts at G.E., and a big thinker in many ways that you will see very shortly. Another big thinker is Robin Murphy, professor at Texas A&M and the founder of Roboticists Without Borders. Great effort. Last but not least, Robert Opp leads innovation and other change work at the world food programme, critical leader there.

Before we let them speak and like I said, this is your role, I want you to put yourselves in the mind-set of 2030. We heard on the first day about back casting as a methodology for writing roadmaps. We are going to do back casting. The first step in that is for every one of you to take a deep breath, close your eyes, and acknowledge that the year is 2030. We now live in a world where the Sustainable Development Goals have been achieved. And AI played an outsized significant and critical role in achieving those 17 SDGs.

Everybody bought in? What year is it? Is the world a better place? Did AI play a big role? Come on, you guys can at least back me up. (chuckles).

We use this as a device because I want you to be thinking about this. The first step, let's do this, Jo Jo we are going to hand out colored papers here, I want everybody to quickly brainstorm another wedge on here, on this big wheel. I'd like everyone to think of a possible solution, something AI could do. We have got 17 examples here. Clean water and sanitation, improve efficient clean water provision, reduce inequality, build a more inclusive society using disability robotics, etcetera. One example from each of you and here I put up rules, there is important rules, you have two minutes to write down your example. Please write down which SDG it addresses. No more than five words. That is important because each of you will get a chance to in five words state what it is. We only have 90 minutes. Nobody has the privilege of taking extra time. But I want you all to get used to this.

All good? We will get you guys paper as well. Go. Pink and red will do. Let me hand these out. Stefanie, would you mind giving me a hand? Pass these around. Pink and red for this round. Then write it down or just tape this on.

Do you want to write this down? If anyone needs a pen, it's intentional to use color and colored markers. Put one more

wedge on here, what have you thought of that AI could do to help address a global goal? One idea, five words. Doesn't have to be so brilliant. We are crowdsourcing here.

We are not adding a wedge, we are helping one. I apologize. We don't get to write new Sustainable Development Goals. That job has been done, come up with examples of how AI could meet one of these goals.

We should have the Jeopardy music.

(humming).

30 seconds. 30 seconds left.

15 seconds.

All right. Let's do this. I'm going to let you guys kick this off, each of these. Five words so you are not allowed to say anymore, I'll cut you off otherwise, I want everybody to speak. This is your one chance to do so. If those of you who have not yet used these machines, quick instruction, there is a button with a speaking microphone on it. Press it once and only once. The awesome control people will turn one mic on at a time. If it doesn't come on, you won't be able to speak, it will be red. We will go in order. Wait for the red thing to be on. Press it once. It will get activated. Chaesub Lee.

>> CHAESUB LEE: Simple one, good health and well-being, I'm thinking of this individually tailored health.

>> Gary.

>> GARY MARCUS: Healthcare, automated scientific reasoning.

>> Peter.

>> PETER MARX: I'm going to leverage that too. I'm going to say that I have a sister suffering from cancer, seems very applicable, we should cure cancer. For added credit, I'm going to point out on the public safety side, give everyone personal safety.

>> PAUL BUNJE: I like there were two times five words. I feel like you are breaking rules.

>> PETER MARX: Add them up, they are a total of five words.

>> ROBIN MURPHY: Number 11, sustainable cities and communities, use AI to prepare and prevent for disasters.

>> PAUL BUNJE: Excellent.

>> ROBERT OPP: I had to look at SDG2 of course on zero hunger. So it was increase market efficiency for farmers.

>> PAUL BUNJE: I love it. We are going to roll through here, if you don't want to go, you can wave me off. Are you ready? Press your button. I can also run around. That is a good thing. Five words.

>> Knowledge and sciences.

>> I'm going to follow Robert up. Stick with SDG2, five words are optimization of agriculture production systems.

>> I chose the SD G3, healthcare, preventing and predicting

epidemic outbreaks.

>> I chose multi sectoral collaboration, and AI to identify hunger in projects.

>> Solution for 17 goal is everybody helps everybody with the worldwide bank.

>> Goal number 17 changing mindsets to have a systemic perspective on all 17 global goals.

>> Define decision success criteria of 17 SDGs.

>> Segment waste streams to cheaply take out recyclables.

>> Elimination of antimicrobial outbreaks.

>> Perfect.

>> Coordinate under water drones for promoting monitoring.

>> Store medical records on cell phones.

>> Number 8, consolidate global production data, implement.

>> SDG4, revolutionize classrooms with digitalization.

>> I'll say that it will increase the number of working places.

>> I picked SDG17, and improve collaboration and cooperation.

>> I love how many SDG17s have been picked so far.

>> In five words, personalize disease diagnostics via robo chemical pills.

>> There are hyphens in those words, aren't there?

>> Hi, this is for any SDG, it's for understand linkages with other SDGs, very generic.

>> Helps 17 as well, that is for sure.

>> Hi, DG11, design future cities and infrastructures.

>> Future cities.

>> SDG16, improve governance by eliminating corruption.

>> Improve governance. Eliminate corruption. I like it.

>> SDG4, digitize indigenous knowledge and use AI to message usage pattern.

>> Energy harvesting and efficient distribution.

>> Automated mapping of whole planet, would support all SDGs.

>> The mapping stuff, I love this. Geospatial guy. I'm digging.

>> SDG10, replace human discrimination by algorithmic discrimination.

(laughter).

>> 10, assist cultural empiricist reflection on superiority complexes.

>> 1, reduce inequality through literacy for all.

>> Goal number 4, facilitate lifelong personalized education.

>> Probably thought outside the disk, quality of life dictates economy.

>> Nice.

>> Number 2, precision agriculture and sustainable agricultural systems.



>> Very good.  
>> 4, personalized midlife re-skilling.  
>> I want that. That sounds nice.  
>> Thank you. Goal 16, peace justice and provide access to justice for all.  
>> Great. Excellent.  
>> Number 8, on decent work, automate thought concept visualization.  
>> Number 17, intelligent projects management.  
>> I love the rule following. This is great. Peter, you are the only one who has broken the rules so far. This is a good group.  
>> All SDGs, automate [inaudible] process.  
>> Number 16, reduce corruption in government, government and private sector.  
>> Right in front of you.  
>> SDG2, precision agriculture plus AI powered logistics.  
>> SDG10, local people developing local solutions.  
>> SDG10, automated fair redistribution of wealth.  
>> SD G3, better preventative medicine with biosensors in AI.  
>> SD G3, precision medicine, better mental health.  
>> I don't think I've heard SDG5, I'm going to go for SDG5, using AI to eliminate bias and end all forms of discrimination against women and girls.  
>> Brilliant.  
>> All SDGs, satellite data [inaudible]  
>> Worldwide reference base.  
>> Worldwide Web for space. Brilliant. I dig.  
>> SD G3, use AI to bring quality healthcare to remote areas.  
>> This is for all the SDGs, make goal progress data available to all.  
>> SDG5, digital inclusion to reduce inequalities.  
>> Number 2, hunger, optimize critical food supply chains.  
>> SDG7, intelligent energy systems universally deployed.  
>> Number 17, AI created new forms of collaboration.  
>> All SDGs, stress over priority for governments over technology.  
>> SDG4, on education, optimize what to learn and when to learn.  
>> SDG8 meaningful occupation of free time.  
>> PAUL BUNJE: Fantastic. Give everybody a round of applause.  
(applause).  
Did we hit all of them? Did we miss any of those? One? You have it. Press your thing. So I don't have to run down there because now I'm lazy.  
>> I think of what could address all of SDGs, and I think of

financial think tank for good.

>> PAUL BUNJE: Very nice. Financial think tank for good. AI makes the decision.

>> Would be part of it.

>> PAUL BUNJE: We did that for two reasons. I want you to flex your muscles. Now we have a database for how to use AI for Sustainable Development Goals, whether the ideas are all great or not may not be relevant. We heard over the last couple days that examples of how this can work is going to be critical in driving more investment and attention towards the idea of AI for good. That was part of that reasoning.

The other reason is I would like to give our panelists now an opportunity to build off of that, and instigate how you think we can use both ideas but also our plans in developing guidelines for a roadmap, that should say roadmap, for AI for good. Think of anything you want on this. I know you are all big thinkers.

I will let you all go to it.

>> The general issue that one needs to think through is which of these problems can we address with today's AI techniques and which might we need new AI techniques that don't exist and how do we want to balance our investments between commercializing or maybe that is not the right word, existing techniques, and how much towards doing basic research to allow us to do things we can't already do. Curing cancer might be an example of something that we can't do with existing AI techniques, because we don't have AI techniques that can read, they are basically illiterate. To cure cancer we need AI that is clever enough to read the scientific literature and integrate it.

That might be a long term goal, whereas some of the things about optimization we might be able to do right now.

>> PAUL BUNJE: Great point. Maybe one of the things we do next is start to prioritize and draw linkages between what is possible now and in the future. Any other comments?

>> We are sitting here across the street from the original League of Nations, and the U.N., and my reaction, listening to everybody's ideas about how AI can help the sustainable goals, development goals, and almost feeling like we are all looking for AI to give us almost a combination of benevolent partner, if you will, to help humans go make better decisions, there is that element. There is another element to me which is the engineer said which is AI is just a tool for humans to make better decisions. It is like we are aspiring for AI to solve the world's problems.

>> PAUL BUNJE: Great.

>> One of the things that concerns me, I have this major -- I hate roadmaps, I've been on dozens of road mapping exercises, Department of Defense and government agencies, and I think these

are all great ideas without even thinking about the technology, but one question is, what is the incentive? Who is going to pay for it? It is not just the technological driver, and is it something that a government or a meta government agency has to invest in because it won't be taken care of by a smart phone app or something, that would be a normal, you could subsidize it that way. I think these are, there is a category of problems. Orthogonal to the actual category of technological solutions.

>> I love all the references to SDG17. You pointed it out, Paul. But it was quite extraordinary to me, because one of the new elements that was introduced in the SDGs was this issue that it has to come out of partnerships and frankly unorthodox partnerships. And the fact that this crowd has picked up on that and is thinking of that and the vision is for me very positive, and I think, Paul, if your question is about principles of collaboration, my comment is that I think collaboration doesn't, sometimes it happens by accident but not necessarily. I think we need some principles around how the collaboration between the different stakeholders actually works. I feel very strongly that in particular the public/private partnership element is really important to reinforce, and that we do need some general frame of reference that is shared between public and private sector on what those partnerships are going for and as you just mentioned, what it would mean for, who is going to do it and what is the incentive behind it?

>> PAUL BUNJE: Excellent point. It may be some of these themes arise, there are areas where AI is a meta support for a number of different efforts.

>> CHAESUB LEE: Just follow what Robert said, this technical applications is very important, AI is coming from the technical domains, but we have to take into account, look at the 17SDG goals, this is heavily dependent on vertically structured, each vertical has its own ecosystems governed by specific ministries in countries. Even our U.N. system has organisations, have a mandate for each SDG goals. So how we can get rid of this, I believe this is a first event, where one of our goal is how to get together with all stakeholders including public domains to address this, maximize our collaborations, even beforehand to understand correctly what AI means. Many of case our understanding of AI not correct enough. This is a good opportunity.

Rather than principles, we have to indicate what is current existing situation, how we can overcome this given conditions.

>> PETER MARX: To add on to that, I think Robert would agree with me, it's interesting nobody mentioned data during all this. AI is powered by data. One of the things that we have to grapple with when we think about developing guidelines is who is

going to own and be able to access and utilize data to go drive the algorithms that are going to result in the tools and the outcomes that we are looking for. I'm curious to know what everybody thinks about the data side of all this.

>> CHAESUB LEE: Actually, I did. First I highlighted was importance of data. I completely agree, let me make a reference of this AI we can imagine of this engine of the vehicle, each vehicle, industry develop their own engine. If we can think of this AI, AI machine has a engine, this engine should be operated by oil, meant data. AI should be operated with input to the data, for the time being, unfortunately we don't have any standard ways to manage this data. Data is completely different.

>> PETER MARX: I was lucky enough to have had a interesting public job which was I was the chief technology officer for Los Angeles for big city. The mayor made us publish all of our data, open data. It is a fantastic unbelievable resource. We have got a couple thousand data sets, that describe, it is the best collection of urban data ever out there. There is roughly 40 cities in the United States that have done this.

The data comes from those governments, but when you get into companies and proprietary agencies, you get into other governments around the world, everybody's idea of what happens with data is different. I've given the open data, I'll end after this but I've given the open data talk many times across the world. The idea of a city publishing its crime data is very mysterious in parts of Asia, for example, where they expect all the data to be published to the government. Or for that matter in south Asia the comment that I got back was, but you are going to destroy real estate values. Of course in the Middle East it's a whole different take on what you do with data.

We have a lot of interesting stuff, I think that we have to do around data governance and access of data, and motivations for getting data out there, because that is at the end of the day going to power all this AI.

>> ROBIN MURPHY: Not all the AI, I'm going to -- because knowledge, there is data and we are all focused on big data, big data, machine learning, wonderful. All good.

>> Metadata.

>> ROBIN MURPHY: But knowledge has an important role. The decision-making capabilities, things we do is normally components of that are not necessarily with large amounts of data. In fact, you look at data as possibly an input but the challenge for the AI is capturing the knowledge. It is knowledge representation, form that can be reasoned over.

>> PETER MARX: In my defense, I never mentioned the word big. I've always wondered what happened to small and medium data.

(overlapping speakers).

>> ROBIN MURPHY: Information.

>> CHAESUB LEE: The point is, why I'm saying this data management be put on this, today, we have no way to clean up the data. We collect all data without any cleaning. How many percentage of data is really garbage, this should be amazing. Garbage in, garbage out. But unfortunately, we collect all this data, we have no clear ideas how to clean up the data. We utilize this data for our knowledge.

>> GARY MARCUS: Aside from cleaning data and so forth, I agree there is a paradigm for machine learning that is dominant. So ubiquitous it is like water people don't know they are drinking, which is everything is structured in machine learning now as a supervised learning problem, which means you get examples, you have inputs, you know what the outputs are supposed to be for that. But human children learn in different ways. They learn by trying things out, they may only need a hundred trials and not a million trials. The way that the machine learning techniques work, that we have right now is they are knowledge free, they don't reason that much over semantics. It depends on having all these examples. But that paradigm might not stay forever. We might get to a system that can do more causal reasoning as I believe children can do and demands for data might change at that point.

>> PAUL BUNJE: There are some questions in the audience. We will get there in a second. But I want to build off this conversation, because you are starting to touch on themes that have been discussed throughout the last couple of days. That is very important. What I want to see us try to accomplish in the next 60 minutes or so, or 45, is good summary around where we think real priorities might lie. To kick that off, I'd like each of you to take two to three minutes and play the role of instigator. What do you think real priorities are, either in principles for guiding a programme in AI for good, or accomplishments, it could be big ideas, Gary, like the AI you had the other day, these sorts of things. I'll start with you, Robert.

>> ROBERT OPP: A few ideas. The SDGs also are very clear in their underpinning that we should leave no one behind. That is of course also where I'm coming from in organisational background that commitment to leave no one behind in the world. So I actually think that it is really important that as we think through the futures, that we are thinking about people who are the most vulnerable and disenfranchised and somehow have that embedded as a underlying principle.

Somewhat related to that has been, I picked up lots of discussions about the issue of, concerns about privacy and

protection of data which is somewhat related to those who are most vulnerable but I think all of us to a certain extent as well. How are we going to protect the people's privacy as we are essentially generating more data than ever. We talked about that. Is there really any privacy anymore? What are the boundaries around that? As a final thought, I suppose it's come up once in a while that the issue of cultural bias and gender diversity or gender bias in some of these systems, I think that is exactly what you were saying, Gary. The current way that we have AI system learning structured is in that paradigm, that is somewhat free of knowledge and it is comprised of a bunch of examples. But I think we definitely need to have insight built into that somehow that is culturally sensitive to the people around the world. We know the wide variety of cultural differences. We have to have transparency and correctability around that.

>> PAUL BUNJE: Excellent. One thing, these instigators, I would like you to think and reflect on these, what resonates on you so we can capture that. Let's go down the line. Robin, I know you have a slide if we can switch to that.

>> ROBIN MURPHY: That ties into what I'm saying. The big thing for me would be principles of collaboration. Yeah, like really have collaboration, and not just put people from different disciplines in a room. How many of you caught that we went into a nerd out starting with words like semantics which have a very precise meaning in AI that may not quite be the same outside of that little nerd space.

How do you, if you don't know our terminology, how do you know what to ask for? By definition, if you are trying to come up with a innovative solution, how do you predict innovation? Normally you don't. We know from the diffusion of innovation, you get a innovation in and you find the killer app for it about 7 to 20 years later, the really good use of the technology was not the one you expected or the second one. It is a whole iterative process in there. How do we encourage that? In fact, we should insist on it to be meaningful collaboration, with real stakeholders, and to be revised, not just one time we get everybody in a room or we go out and trot around a field or something, and try out a UAB, it is to be revised and include all the stakeholders, a lot of hidden stakeholders.

>> PAUL BUNJE: Can we get Robin Murphy's slide up, switch the two presentations? The point about iteration is a salient one and meaningful one as is the notion of what real collaboration looks like. They may not come up. Let's move to Peter.

>> PETER MARX: Sure. I'm going to take the issue of ethical principles on, and I'll put it in the following way which is we are talking about, look, I would be willing to bet everybody in

this room is affluent enough to have a smart phone or an automobile, have clean water, safe food, all that sort of stuff. The reality is stuff we are talking about we have to take a considered decision about whether we want it to be accessible, equitably distributed, whether the people working on this are going to be diverse and reflect the populations who are going to be affected by this. There is a lot of very humanist considerations here. I'll bring up another loaded word that nobody has mentioned so far, which is labor.

As we continue to automate and rely upon AI, whether it's expressed as robotics or home healthcare or whatever it happens to be, there is going to be labor displacement. One of the ethical considerations that we have to take, have to find a solution for if you will, is what is the responsibility of the technologists and the companies and the entities and the organisations that are going to be promoting all this, to go and address the issues of labor displacement. And make sure that the benefits of all the stuff that we are talking about for 2030 are going to frankly be distributed to everybody.

I say that on, from the point of view of a California Utopian. I'll leave it there.

>> GARY MARCUS: People heard the suggestions I made the other day but I'll repeat them. One is that we might want a way of institutionalizing collaboration and also of sharing AI by having something like Cern as a model for AI. The default right now is a small number of corporations will own most of the IP for AI and that is not necessarily a good model for society.

A lot of problems that we need to solve to do things like solve cancer require work that isn't going to happen in a corporate context where the main goal is to sell advertisements but maybe can't happen in academic context, either, where the typical goal is to find one niche that might be independent of everybody else's niche. I propose to have Cern for AI, and the other thing that I proposed was that those of us who do the research in AI would love something like charity navigator, that would help us know where can we have impact. What are the problems that we might have maximal impact on, where is it likely that if we come up with something good that it's actually going to get put into practice and so forth. Those of us who don't help other people for a living that would like to help people could use help in figuring out where we can best have some impact.

>> CHAESUB LEE: For me, I'm not so much concerned about the technical development of AI, because as a technical engineer myself, working as communication engineer part 30 years, engineers they are devoted to develop this new technologies, AI technology should be readily engaged or not, your concerns is

going up, so they are continuously developed. So this, I don't care about this, just let engineers do by their self.

Issue is after develop those certain technologies, to try to extend this, the benefits to society is sometimes difficult. Especially AI I believe this is one of the nature of this AI has a huge impact, as you can see of these 17SDG goals, impact. We have all our life should be impacted. Most important urgent matter is how we can form this dialogue platforms, autonomous platforms, how we can make it even different, different requirements, but those technical engineers they need some kind of ground to talk with governors, politicians, some business domains, we need something ground of this providing of such a dialogue. If not the case, what about we discuss any subject might be challenging rather individually challenged, AI is a good subject we can have some common ground as much as possible. That from my point of view.

>> PAUL BUNJE: We are definitely starting to hear, there was reference in the audience to the role of AI in underpinning solutions through this type of platform for dialogue or collaborative work in decision-making and the like. That is an interesting theme.

I know there are questions, but in order to move things along, if you don't mind holding those to the end when we can open this up, I'd like all of you to move to the next step and take a quick reflection on what the panel was just describing, and start to brainstorm now for their reaction some of these principles. You have gotten some ideas, principles around collaboration and transparency and the role of intersectoral dialogue. Take two minutes, you will each get another fantastic piece of paper, to write down, and you can write down as many ideas as you want, but only one idea per piece of paper. Make sure that they are only three to four words in this case. Here is how this is going to work. You are going to do that in groups, small groups. Yep. You are going to talk to people. If you can build an AI to do this for you, fantastic. But you only have three minutes. In those groups of choose ten to twelve folks, gather the people that are near you, summarize quickly, I'll tell you when to summarize because I imagine there will be some overlap in what are important ones and choose one person to report out what you think the top five ideas of principles might be. I want you all to, feel free to wander around and such. But I'd like you to listen to those reportbacks, essentially and prioritize. You will have a chance to say, I heard this, this and this. These are the top three. These are the top five.

Take two to three minutes to write down as many ideas, one idea per piece of paper. Only a couple of words. Then I will



tell you when that time is up. I would like you to sort these in a small group. Go. Now we need the music again. If you need more pieces of paper, the lovely Stefanie, lovely Jo Jo, are here handing them out. (pause).

>> Do we get to give points for anybody who uses terms like IoT?

>> PAUL BUNJE: Peter is giving out points to anyone who uses big words.

>> ROBIN MURPHY: Give them a point and I'll take it away.

>> PETER MARX: Negative points for big data?

>> PAUL BUNJE: Is there a prize, beers or cocktails at the end? If you haven't gathered into a group already, take the next four to five minutes as a group .... (pause).

You have about 30 more seconds to sort out your top priority principles.

>> Do we get to give points for, can I use the word exponential?

>> PAUL BUNJE: Do I get to give points --

>> You get ten points if you use machine learning, neural network, spiking neural networks. Deep learning.

>> PAUL BUNJE: Let's come back together. The panel is anxious to hear your wisdom, sort through it, take advice, etcetera. I'd ask everybody to be quiet. The first group to raise your hand and go wins something. Wins XPRIZE swag. I'm sure we have some somewhere. You just started working on it. All right. Will you buy me a beer if you get another minute? Take one more minute.

Is everyone close to at least an idea? I know it's hard. There is a time limit. Yeah, you guys are? We have a hand raised. Be confident. Prizes have a remarkable effect of doing things like driving you to a deadline. Does anybody remember when you had a term paper due the next morning? You would tend to write your best paper that night? That is what we want. I'm cutting you off. Let's do this. This group here, who is going to win some XPRIZE swag, has the floor. You guys did great. You got great answers. I would ask everyone -- we will call you all the group awesome number 1. If I can have everyone respectfully listen to some top principles for a building a roadmap in AI.

>> It should be open source, copyright-free, no patent allowed on it, ethics, it should have ethics, and governance.

>> PAUL BUNJE: Any of the four of you have any questions regarding that?

>> ROBIN MURPHY: Governance like what?

>> PAUL BUNJE: Governance how? How did you mean governance?

>> There should be a oversight, like a U.N. body to make sure it's doing the right thing. But you can't imbue it in law

because laws are meant to be broken. (chuckles).

>> Those are nice principles, but the problems is you need incentives, people have to work on things. As soon as you start to say, open source, there is no licensing fee, when you say no patents, there is no licensing fee, etcetera. I'm curious to know what you think the incentives would be.

>> At the moment, we are money driven. With a full implementation of an autonomous AI system it's going to completely open the wealth system. So we shouldn't really be thinking long term about the money.

>> What if the public sector creates open source AI and the private sector is in better position to --

>> The public sector can learn from what the private sector is doing, a reasonable AI.

>> PAUL BUNJE: Very nice. This may point to a need for practical dialogue, to enable these kinds of conundrums. Who would like to go next? Our big group that insists on more time. These are the good ones, right?

>> Yeah, the heart. We have four major ideas, the first one goes for data and data security and privacy issues, dealing with the transparency, accountability, and the second one with education like clear roles and responsibilities formulating rules and regulations and the third one is like having a decentralized research work, decentralized governments, and --

>> Blockchain.

>> Yeah, and the fourth one is kind concern for AI and U.N. leading the research, transparency, empowered U.N. ...

(silence).

>> PETER MARX: You are attacking the power issues that exist around AI. Who, if you will, the power structure, that may very well exist around AI and decentralizing that power across individuals and even world organisations, if you will. I guess my reaction to this, high principles, absolutely for the best, what was the term that you used, unlocking the wealth economy, was the last group? But I'd be curious to know how we transition to such a thing, how we get there.

>> We are using an autonomous machines running for eight years as a model, which is the blockchain. Decentralized model here, we mean by the machine to be decentralized and not a human intervention.

>> PAUL BUNJE: One interesting thing I'm picking up is the interaction of some of the different technologies and applications of technologies and how that may underpin this work.

>> If I can be skeptical, it is not clear that we can get to something as complicated as machine reading with a completely decentralized research approach. It's one thing to decentralize

something like blockchain and another possibly to decentralize research planning. That is what we do now, we have been working on AI for 60 years and we haven't made any progress on either machine reading or common sense reasoning and so forth. There may be some pieces that we can't do in a purely decentralized way.

>> PAUL BUNJE: Another group that would like to go. Here in the back.

>> We came up with five principles, gender equity and inclusiveness, common and shared values, pragmatism, continuous impact ...

(silence).

>> Resonates with me. The challenge what you are describing there is almost like a melding, if you will, of the private corporate world with the public world with the academic world, which is, these are the high principles, when they come together that is where they are. Part of the challenge that we have is, in fact we have a lot of structures which segregate if you will incentives, things like patent law and the rest that make it more difficult to get to that common set of goods that you describe there. I'd be curious to know what you think, Paul, about how we get there.

>> PAUL BUNJE: I'd like to ask others how we get to interesting places. One thing I'm picking up on and it's relevant, despite the age, Gary, of research into AI this is a nominal conversation about applications. There is a significant amount of Utopian perspective on how we can develop this out. I'm not hearing as much pushback as I might otherwise think, which is meaningful, to me anyway, because I read that one of the things we can do is set forth in advance principles for future work, if we can accomplish that or not, may become another question. That may be the practical nature. You were referring to this as well, of how we develop it. We turn the dystopian mics off.

>> That is what usually happens to me. There is a cartoon of someone offering hard truths and other people offering happiness, something like that, the line for happiness is much longer than the line for hard truths.

There are real issues to the fact that places like Google have all the money right now. Other things being equal, if you are good at machine learning ...

(silence).

You have to be at least somewhat mindful of the financial pressures, IP pressures that we are alluding to and so forth. The set of values about egalitarianists are obviously values that we want but we have to be realistic about, what is the mechanism that we are going to use to get there. How are we

going to move from AI being a tool for advertising and selling people stuff they don't really need, to a tool for solving SDGs. Talent that we need, how do we protect IP if we need to do that, etcetera.

>> PETER MARX: There is an implicit tension here, which is that we may end up in a world where a very few companies have all the data and all the algorithms and all the power, which may be measured in money if you will, and then there is another world which is empowered, everything has been made available and accessible to everybody, if you will.

In fact, actually the folks who are probably the most in need of it, folks who are looking for clean food and clean water, there is five billion people who have no ability to call a first responder. There is no 911 for five billion people out there, for example. Are those people going to be empowered? Or are we going to parcel it out to a few large powerful entities.

>> PAUL BUNJE: The points you are making is the reason I asked everyone to accept the idea we are living in a Utopian world in 2030. The question becomes how do we accomplish it considering the barriers and lack of opportunities. In other words, if we are going to embark on an AI for good programme, how do we structure practically --

>> I'm imagining now a review paper written in an academic journal in 2030, and it begins by saying the astounding thing about the last 15 years is not just the major advances that were made in AI, figuring out how to get machines to read, but the way in which the whole structure, the fabric of society was changed from a capitalist versus society to one in which people focused on Sustainable Development Goals and restructured everything from the IP framework to how people were paid to do the work.

>> PAUL BUNJE: I did ask them to be provocative. I like that. Let's go to another group, an opportunity to summarize. Robin? You had something.

>> ROBIN MURPHY: One of the things I'm confused on.

(off microphone).

We seem to be focused on the word principle as abstract social principles whereas a nerd engineer I went for principles like theoretical, physical, physics type principles, if you are going to do collaboration, the first principle is that stakeholders have to put out a meaningful project, something that they are going to have and that the researchers would have access to it, it would be open access. And everybody commits to do it for a minimum of ten years. None of this swoop down, we have a solution, oops, never mind, we are moving on, something else kind of thing.

Getting those types of actionable concrete principles for

collaborate, because you can't collaborate if you are not all on the same page, in the same room, working on the same problems, having feedback, no matter how good everybody, what their good values are, you still have to have mechanisms. That gets back to you as well.

>> PAUL BUNJE: Great. Let's let the last group -- we have three or four more groups, if we can summarize and then I'll turn back to you all to finalize it out.

>> Hi, our group came up with two principles that seem to tie with whatever the panel was saying just now. The first one is no one left behind, and the second one was to take a holistic approach.

>> PAUL BUNJE: We will get the rest, then allow the panel to pull us all together in the end. Then there will be a coffee break following this. I'm not going to tell you when the coffee break is because I'm taking moderator's privilege. Other groups? In the back here, very back.

>> There was I think two ideas here, one was having a mechanism or platform for having transparency for the AI that has been developed in applications. Another idea was having a certain measure for good, right, so you have some certain maybe forms that you have to fill in, and there is this platform, like the charity navigator that one of the panelists mentioned, that based on the good score of your application or your whatever you are developing, you can get maybe more resources, or you can be rewarded in a way, so that you have incentive to go to this platform and prove the measure of good of your direction basically.

>> PAUL BUNJE: One more comment.

>> Yeah, so the point about transparency was sort of adding, the points of the transparency is the intentionality. Being transparent about the intentions of the collaboration, and for the collaborators to be transparent about the intentions that they each have within the collaboration, what they are expecting for their AI systems, system or systems they develop to actually do at the end of the day, so people aren't caught off guard and things aren't used for things that people did not want them to be used for.

>> PAUL BUNJE: Intentionality, tie practically to what you were suggesting around stating specific goals and programmes.

>> ROBIN MURPHY: There is intentionality. Transparency is such a weird thing to me, because being old in AI we learned in the '70s if you build a AI system that anybody is going to use, you have to explain what it did. It has to explain what it is. It's good software engineering. The fact that we are rediscovering it says a lot about the failure of people like me as a professor of computer science in teaching decent computer

engineering skills that we are making sloppy systems. We are not even making it transparent unless you have something like a deep net in which you are in a different league all together.

>> This is the issue, the field is dominated by people who don't know much about computer engineering.

>> Or AI.

>> I showed the XKCD slide, pouring data into linear algebra, stir and hope for the best. People who are postdoc level or graduate student level don't know much about computer engineering, about how to build large scale systems and how to do verification, all this stuff. They are working off of a paradigm that is naive. Transparency is like sudden idea from above that they hadn't even considered because they don't have the background to think about it.

>> There are few people working on basic technologies. Most people are just preparing meals using ingredients off the shelf.

(silence).

Frying pan, exactly. But there is a fundamental thing here, which is that a lot of these principles come down to the fact that the technologies and the tools and ingredients and even the meals are outpacing, if you will, the human's ability to go and work with them. A lot of stuff we have been talking about is very high above the plane of the ground if you will. The people who are actually working on water systems and farmers who are working in the fields, the engineers, the field engineers working on jet engines, things like that, they are not thinking about AI today. One of the challenges that we have is we have to think city planners, for example, everybody thinks that cities have these masses of analysts, who are sitting there collecting data, saying how we need to make a better city. I have yet to find that army of analysts.

>> The opposite problem can happen too. There is a army of analysts in New York, NYU is ready to go.

>> Thinking of cusp.

>> Yes, it's not clear we have the right data for the right problems yet. Something will come out of it but it's not clear yet that there is a killer app for something like that maybe because we don't have the right data. Maybe people haven't asked the right questions. But even when there is a army --

>> PAUL BUNJE: Is that a indication of a lack of community between the data gathering side and the problem definition?

>> PETER MARX: I think what it is, human structures believe it or not are more slowly moving than technological innovations we have been coming up with. 2030 is closer than we think. I'll put it that way.

>> PAUL BUNJE: Let's gather last few comments. How many groups have not gone? Raise your hand if the group hasn't gone.

1, 2. I'm going to let these groups go and then we will open it to questions. Did you have a summary as well? You are the third. Here and to the back and to you all in the front to finish off.

>> The first one has been mentioned, transparency. We had inclusiveness, what was already mentioned about leaving no one behind. Then but separately about Omni science so the knowledge of the AI should be made available to lots of different people and not just be clustered within one group that is knowledgeable at the time. Trying to foster bottom up type of collaborations. The last one is about complementarity in these efforts when you want to collaborate so that, with there being competitive process as well but when there can be collaboration, they are complementary to foster that.

>> We have five rough principles that we collected. The first one is democratization of access, usability, to enable people to reach this artisan period of AI and tinker with it. The second one is multistakeholder approach to involve a variety of different groups in society. Then joint ownership of ideas and solutions and collaboration. Another one a little softer, spirit of opportunity, while at the same time recognition of the risks. And then perhaps this last one, the facilitation and managing the collaboration and partnership.

>> PAUL BUNJE: Excellent. Some common things are starting to emerge which is helpful. This group here.

>> My point is by slight abuse of terminology to use transfer knowledge across domains and data. These 17 things are never going to succeed if they are done in silos as other people point them out but from the same token the data and knowledge you learn one is going to have to be applied to the others for pure efficiency reasons. Some of it can be technical in terms of formal knowledge transfer but also a diversity of people who are trying to solve the problems.

Then a comment regarding the last panel discussion, my being a 50-year-old something, I would caution other 50-year-old something not to complain about computer science students who don't know how to use soldering irons and assembler language and punch cards, because gosh, you know.

>> Wasn't quite the complaint.

>> ROBIN MURPHY: When I teach AI classes, becomes a reflection on the computer science curriculum.

>> How many people here have programmed, anybody, that is it, we are all barred from further conversation.

>> One principle to add, probably to base that collaboration programme on a formal general problem solving framework, for example, the mathematical equation model, and formalizing our goals based on shared language understandable to all or most.

Yeah.

>> PAUL BUNJE: Question, comment or addition?

>> I wanted to clarify the panel's interpretation of what we meant when we said transparency. I think the conversation got derailed a bit in terms of how it was interpreted. When we are talking about transparency, I think there was some mention of sloppiness of software programming and not being able to explain, and in the context of AI the techniques that are being used right now, there are two aspects in transparency.

One, in actually declaring what those cost functions are on which the AI techniques are being optimized, so in the case of targeted advertising, we are saying what we propose from a transparency perspective is that you declare that this is going to be used to, let's say, figure out what your mood is. Then if based on that they are pushing alcohol sales on you, then that is ethically not right.

(silence).

About certain neural network techniques currently, they behave like black boxes because we can't interpret how they arrived at specific decisions and there are efforts being made, as an example, professor's lab to add interpretability so I wouldn't attribute that to sloppy software programming which was I think brought up. I wanted to clarify.

>> ROBIN MURPHY: Transparency, and I apologize if I misspoke, transparency isn't necessarily a sign of sloppy. But if you are writing an AI system, if you are using let's go old school, an expert system, you want to have transparency. If you did not put transparency in that system, where you could see, then you were by definition a sloppy programmer, because you should have known better. This is like going back to the idea of being over 50, you know, we don't really need to know about World War I or the Civil War or French revolution or any of that other stuff because it doesn't matter because we are all in a whole new bright shiny age.

As far as computer science, the history of AI, we are still using ASCII as a underlying code, as a representation. That goes back, it is nice to know why we do certain things. I'm not going to let us off the hook for that. But I thought your social transparency on what is going on with the systems that are mining us is incredibly important and worthwhile.

>> Technical point quick. The problem is not that the cost functions are hard to describe. Those are simple but you wind up with ten million parameters. I'm sorry your child was killed by the driver of the car and the reason is parameters 313 and 314 were negative numbers and should have been positive numbers, that's not satisfactory.

>> PAUL BUNJE: Let's wrap that up now. You all have heard,



you get the last word, which is useful for you all. We have heard a lot of ideas around what some of the principles are. As the panel, you have the prerogative to define where you think real priorities are. I would ask you again to put yourselves in 2030, to your point, this is not very far off. What were the most important things that we as a community, trying to build some sort of programme under this banner of AI for good, that what were the most important principles that you believe got us to this date in 2030 where the SDGs have been significantly helped? I'll go down the line, if you want to offer two or three and we will capture these together.

>> CHAESUB LEE: Definitely this collaboration is most important part, we are talking crowdsourcing, challenge of collaborations, that would be one of the way that is the most important part, if you really have SDG17 goals, basically inclusiveness has hit me because we had a lot of effort for shutting the digital divide and now AI is coming to us. We are wondering about we may already have a intelligence divide, it will be heavily impact of this widen of this digital divide. Those are concerns in my mind. So crowd collaboration will be one of the important part.

>> PAUL BUNJE: Excellent.

>> GARY MARCUS: Looking back, until 2020 there wasn't much progress in AG I, artificial general intelligence, it has only been now around engineering. It was that time AI opened up and people collaborated on big science projects for AI, why are we stuck doing this narrow thing. This collaboration was important, a model of collaboration where people work together not to build better advertisements but to build the next generation of AI. Even though it was too big for one person to do by themselves.

>> Peter.

>> PETER MARX: I guess I think about the way that our managers, leaders, elected officials, agencies, departments, all those entities, if you will, will use AI and these tools to go and accomplish the 17 Sustainable Development Goals. It comes from the basic, one of the things we can do is we can tell everybody in a position of influence or authority to go and say, look, you are not going to accomplish those goals unless you are using the emerging technologies of which AI is obviously an emerging technology, to go and accomplish this.

I hate to say it this way, we are in a room where we are all bought in, if you will. But much of the world doesn't know what we are talking about yet. We have to get them to a place where we go and say, look, you have to consider the use of modern technologies to accomplish the SDGs, and I'll put it there.

>> PAUL BUNJE: Another flavor of inclusion, to build off of

what you were saying.

>> Absolutely.

>> Robin.

>> ROBIN MURPHY: In terms of what I've heard, the thing that is the most, the concept that is the most actionable and would have the most impact is the idea that better risk projection, what are we using for metrics, how are we -- where is even the research into better metrics on this? How do we project risk when we are using AI in very complex systems, in a multi dimensional systems we get the harder it is for us to mentally simulate, and then when we try to computer simulate something, it's hard to simulate it correctly because we forget something because you are simulating what you pulled in. There is a lot of work to be done there and it would be very valuable work.

>> PAUL BUNJE: Rob, bring us home.

>> ROBERT OPP: Our vision of 2030 that we put on the cards or in the notes will simply not be a reality unless we are inclusive and diverse in the approach. I hear that coming out loud and clear.

I personally believe that. It is simply impossible to drag people kicking and screaming into that vision, if they haven't been involved, if there are only certain views being represented, and it's only being programmed by certain mindsets. I think that is the biggest one for me. The other piece of it is that that probably is enabled by some sort of the notion of democratized access that is coming through. The ability to have more accessibility from people that does not necessarily depend on purely income level or geography, but the ability for people really to access and use systems as quickly as possible in very low cost and simple, culturally adapted model.

And that probably comes as a result of collaboration that we have been talking about.

>> PAUL BUNJE: There are fantastic themes there. I'm reading in some really simple high priority efforts that I think that the entire community can take forward particularly around inclusion, collaboration, access, and much of what I heard from the group reflects that and underpins it with more substance. I thank all of you for your insights and thoughts. Round of applause for our phenomenal panel here.

(applause).

(silence).

For the final couple of plenaries beginning at 4:00. Thank you all. Thanks, you guys, that was great.

(break).

(switch of captioners at 1600)

\*\*\*

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.

\*\*\*