

FINISHED FILE  
AI FOR GOOD GLOBAL SUMMIT  
JUNE 7, 2017  
2:00 P.M. CET  
PLENARY 2: TRANSFORMATIONS ON THE HORIZON

Services Provided By:

Caption First, Inc.  
P.O Box 3066  
Monument, CO 80132  
1-877-825-5234  
+001-719-481-9835  
Www.Captionfirst.com

\*\*\*

This 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.

\*\*\*

>> URS GASSER: Ladies and Gentlemen, please take your seats. We would like to start. Thank you so much. Being myself Swiss I'm compulsory about keeping time. So please take your seats. Let's get us started.

So good afternoon and welcome to this Plenary Session No. 2. My name is Urs Gasser. I am the executive director of Berkman Klein Center for Internet & Society at Harvard University. And we launched an AI governance with the MIT lab. Thank you to all the speakers from the morning. We have a fantastic lineup of speakers. We'll roughly follow the mode of this morning's session; that is, we will have a series of short presentations followed by one question that I will ask. And if time permits we will then also have a bit of a discussion among the panelists towards the end.

So this session is actually a sort of a deep dive on some of the themes that were introduced earlier today, in particular in the session before lunch. And the question before us now is what's at the horizon. What are kind of the next generation of AI based technologies that are developed in the labs but that are also increasingly adopted and people are using. And what do we expect from these emerging technologies. How will the world

look like in maybe 20, 30 years from now realistically speaking.

I know from the presentations that I have seen that we are very much following the tone of the second panel to take a realistic look at what is happening right now and probably a bit less buzz words and type. So we'll take kind of a fresh perspective on that topic. We'll not only look at what's ahead, what's the trajectory of the technologies, we will also talk about the impact of those technologies. What are some of the main challenges and again some of the core themes that were introduced earlier today. What are some of the opportunities and how can we harness these benefits of AI for the social good. We'll do so, we will address these questions and issues from multiple perspectives. Again also as a follow-up on this morning's conversation. Acknowledging the need that we have to bring the technologies together with the social scientists, with the ethicists, with people who think deeply about Human Rights, about ethics, about the law. And so we will have a very multi-disciplinary session then.

Our first speakers, I will briefly introduce the speakers and you have links to the bios, is Professor Francesca Rossi who is a distinguished research scientist at IBM Watson research center and Professor of Computer Science at the University of Padova. She is not only developing AI technologies but thinking about the ethical questions in the development and related to the behavior of AI systems. She has many affiliations, among them is that she is a fellow at the European Association for AI. And she also sits on the board of the partnership on AI which was introduced already this morning.

She will talk about the opportunities on the horizon and the promise of a symbiotic partnership between the machine and humans. Francesca.

>> FRANCESCA ROSSI: I thank the organizers for inviting me here to give a talk. I am happy to be with all of you. It is a very multi-disciplinary environment. Can you hear me? Thanks.

So I think that's the only way to take -- talk about these issues in a very multi-disciplinary environment with moving forward one initiative, one event at a time and being concrete about the next steps.

So yes, so I will talk briefly about the opportunities I see on the horizon, but also the challenges, especially the ones related to ethics. So the main point is that I think that the future but even a little bit present but more in the future will be about some symbiotic partnership between humans and machines. And the reason we think that is that more and more in academia, which is my whole career has been academia but now in a corporate environment, I see this is the best way to achieve the best results in terms of using AI in our everyday, personal and

professional life. So we have seen a lot of transformations since the industrial revolution. You -- everyone knows about them. We have seen electricity coming and automobiles, mass production and then information, telecommunication. And now we have seen this great revolution which I believe is possibly greater than all the previous ones, which really we are -- we have the opportunity to greatly augment human intelligence. With machines that based on the AI techniques can really learn over time how to improve their behavior, can reason in ways that we humans have less capability to reason. And also hopefully and necessarily they will interact with us in a very natural way because without that natural interaction then the symbiotic relationship cannot take place.

Why I do think the future is human plus machine together and not just machines alone? Because I see more and more in many deployed AI systems that this is really achieving great results, in health care, in various kinds of industries, in education, in space exploration and in many other application areas. And the main point, the main reason is that I really believe humans and machines are very complimentary. So we are very good at asking ourselves some interesting questions, machines are less coded. We are very good at common sense reasoning, which is understanding how the world functions and reasoning on top of that understanding. We are also very good at value judgment based on our also emotional situations.

But on the other hand we are less good at other things that machines are much better at than us. So machines are much better at patent discoveries, statistical reasoning, large scale mathematical reasoning and so on. So they can really help us in those aspects that we are not very good at. So that's why this complementarity between human form of intelligence and the machine form of intelligence is what has to be exploited in order to get the best result and in order to make better decisions.

Our life is full of opportunities and needs to make decisions in everything that we do. And so it is really important to be able to make better decisions. And what do I mean by better? Better depending on the context and the environment can be more efficient. More confident because based on more knowledge coming from more data, more grounded, more insightful, making decisions that can consume less resources. If you think, for example, at our planet or even in very less scale environment and even more ethical decision and less biased.

So better may mean many different things in different contexts, but I think that these synergies between the human and machine form of intelligence will really bring all of that. I think it has the potential to bring in that.

So let me give you a few examples that are existing already. So here, for example, is an example where Watson and IBM is helping doctors to find the best diagnosis out of some analysis of -- in this case it is breast cancer kind of diagnosis. So, for example, here the doctor can conduct the diagnosis and can decide whether it is an lipoma or necrosis or whatever are the options. And then help by the machine, the doctor, double check whether this is the most plausible thing. And if the machine thinks it is not then can ask for an explanation and give feedback on why -- why some other diagnosis is actually the best explanation.

So here the doctor is really helped by the machine in trying to understand what's the best diagnosis in this particular case. But also learn more a general concept that can help him in the future as well.

The second example is about ALS. And so in this case in collaboration with doctors, starting this dangerous disease, the doctors have examined some genes to understand whether they are responsible and correlated to the ALS. But they -- they can't examine all the genes. So they asked the machine to rank the genes according to what the machine thought and they look at the top ten ranked genes. And eight out of the top ten are found to be correlated with ALS and five of the eight are not known to be correlated with ALS. Here the machine is helping doctors discovering new correlations between genes and the disease. And so it is helping them to focus their reasoning because otherwise the doctor would have to, you know, go through all the genes or randomly look at them without having some initial ranking that the machine based on all the data and information can give for them.

The third example is about creativity. Machines cannot only help in checking data, to help doctors but also checking, you know, the creative landscape to help artists to develop and to create. And so in this case, this was in Barcelona where the machine looked at all the works of Gaudi in Barcelona. Look at the style of Gaudi. And then suggested according to this tile, some shapes and forms and materials to use to create something. And this is what the artists came out with based on this suggestion and interaction with the machine.

So another example which is more general is that IBM, we are also focusing on social good specifically which is mostly related to this event here. So we have a program which is very interesting where students and post-graduates from various Universities come to IBM are mentored by inside researchers in order to tackle a project to develop a project related to some social good initiative. So in collaboration with some Institute, some foundation, organization that can support with the data

they tackle one problem and I'll give you one example. For example, here about the Zika virus. This was a project that was done in 2015 where a student using machine learning and data science together with IBM researchers they focused on identifying primate species that can host the Zika virus. And this is helpful to understand the strain of the virus and how to cure it better.

So with all these potential, of course, we know that whenever we have such a powerful technology there are also challenges. And as we mentioned in the morning already these challenges can be of many forms. Some of them are technological challenges. So one of them as we mentioned algorithmic accountability or capability for these algorithms based on AI techniques to explain why they are suggesting a certain diagnosis, for example and not another one because without that capability of explaining to the doctor why they are doing or not doing something or suggesting or not suggesting something then the doctor will not be able to trust the machine correctly.

Of course, we don't want the doctor to trust no matter what the machine says. We want the doctor to achieve the right level of trust in this machine. So the machine should be able to explain why it is suggesting some diagnosis and not another one. So that the doctor can know the potential of the machine and also the limitations. So not necessarily everything that comes out of the machine is good because as Yoshua said we know especially techniques based on machine learning they are very good in various tasks. And they have a very small percentage of error but they always have some error. And the error that they have are not what you would expect from a human being. It is a very strange kind of error. Like the two pictures we saw this morning of things that look similar to us but they were interpreted differently by a machine. So definitely algorithmic accountability and strainability is very important. And we still are not, you know, are not there yet to understand how to make them in the best way. And then also to make sure that when we deploy an AI we have to make sure that they work according to values that are lined to human ones.

If you want to give a doctor a decision support system that will support the doctor in whatever decision is to make, diagnosis, therapies and so on and you expect the doctor to follow some ethical guidelines, you also expect that machine, that decisions super system to be aware of those ethical guidelines and to follow them while making those suggestions. You need to understand how to embed these ethical values in to a machine. The data bias problem has been mentioned because data is everywhere. But -- and the more data, the more the machine, you know, as behaves with less error but it doesn't mean that it

is not biased. It doesn't mean that it is diverse enough. That it has all the features that you need to capture really that space of the decision making.

And so this ability to be aware of these biases, possibly mitigate them and alert the final human decision maker that there could be some bias introduced by the data. It is crucial in making whatever decision capabilities provided by the machine then it is taken with the right level of trust and not more or not less. And there are also some global responsibilities that have to do with the ethical development and the deployment of the AI system and impact on the workforce. There are various things that we can do as companies or even noncorporate environment. We can define principles and guidelines and recommendations on how AI should be developed, what the purpose should be.

For example, for IBM the purpose is that we don't want to replace human intelligence but augment it and trustworthy of how we deal with the data of humans used to train the machine learning systems. And we want to support, you know, this evolution in the workforce. But in general I think that devises many principles. And in the last two years here is a complete picture showing all -- many initiatives that have started in the last only two years or even less that have all very overlapping goals. So understand the impact of AI on the society, on people, on organizations, on Government, on policy making and so on. And so understand what it means to ethically develop, deploy AI and how to embed ethical principles in the AI system. Value alignment, data issues, data privacy, policies and so on. And all of them are very related to each other.

And in fact, one of my goals in the short term is to try to understand the relationship between all of them because we don't want, of course, this initiative to do redundant work but want to be complimentary and join forces and -- but this is just to tell you really the whole world from a different perspective is tackling these issues that we are discussing today.

I will speak just more in depth about the partnership on AI that has been mentioned already. So I represent one of the founding members there, IBM. But you may know that this has been funded by six companies which are the six that you see in the top row here, Amazon, Apple, Google, IBM and Microsoft. There are many other partners. Some of them are companies but also many are NGOs, Civil Societies, academic associations and so on. They want to study the impact of AI and educate people, policymakers, government, associations about the real potential of AI, the current state of the art, the limitations, what policies can or should or should not do about regulating AI research, development and deployment. And in general, and this

has to be done in a very multi-disciplinary environment.

However the role of the companies here is to provide the voice of the customer. So the voice of the real life deployed AI systems that once you deploy them in to the real life you really can see the impact and you really can see the challenges and the issues that are raised. So that's where the companies can play a role here. And then they can discuss together with everyone else.

So to summarize, for me the future is really about, you know, human plus machines. I have seen many, many examples where this is the way to go if you want to have better decisions. Whatever better means to you. And it is not the human alone. It is not the machine alone, but the human and the machine together. And I think this has the best potential to really solve the societal challenges that we face everywhere in the world to close or mitigate the digital divide and also to be very impactful in Developing Countries. Because whenever I mention this health care domain application that I mentioned in my talk I always think and the examples that I have are from our society, like first world society. And you can see some improvement there. Helping doctors to make better decisions in this kind of health care system. But, of course, the potential, the real potential is to help doctors in the Developing Countries where they don't have the support. They don't have the kind of information that our doctors have. So that's where the real impact will be, I think. Okay. I'll stop here for now. Thank you.

(Applause.)

>> URS GASSER: Thank you very much, Francesca. If you don't mind to grab a microphone, I have a question for you. I liked very much your framing to think about humans and the machine as some sort of a partnership. And partnerships, of course, involve some sort of bidirectional conversations and interactions. Now today naturally as you focus on AI much focus is on the technology side of this partnership, right? It is about the machine. And we learn where we stand, what the future brings and the like. But I was wondering from your perspective, particularly being a computer scientist, what is somehow the human dimension? Where do you see things going in the future? What do humans, what do we as individuals have to bring to the machine to make it a productive relationship? What does it mean for our mindset for how we think about this interaction?

>> FRANCESCA ROSSI: Yeah. So when I presented this idea of the partnership between human and machine, I described some examples of things that we can provide. We can do much better than machines and vice versa. So, for example, asking the right questions, this is something that we do very well. Or even, you know, value judgment or common sense reasoning. These are still

things that we don't know how to embed in machines that I don't see a very short-term solution or how to imbed these things in to machines. While -- so these are the things that we really can provide. And without them the machines will not be able to make good decisions by themselves. And so really we are needed and we are needed to the machines with these things that we can provide in order for them human plus machine overall system to make a better decision. And also this will, you know, in some sense free us from, you know, trying to do the other things that machines can do naturally much better than us, like, you know, very complex, you know, repetitive reasoning or statistical reasoning. Our brain is not that good at doing that. And so we can focus on the other things that we are very good at that we can provide.

>> URS GASSER: A risk that somehow we start actually to Delegate our questions and our judgment to the machine. Because we ask Siri if we want an answer. So yes, we ask questions but somehow the judgment is increasingly delegated to autonomous systems. Do we need some sort of new or enhanced skills of critical thinking as we interact in this partnership? Would you say there is also a kind of a mental process required that we need to readjust looking at some of the examples that you provided?

>> FRANCESCA ROSSI: Of course, there are many things that we are delegating to machines. Using a calculator to do complex division that before people were doing by hand and now we don't know how to do it. It requires some more time. So we are already delegating things to machines. But I don't think that we are very close to delegating this really very human specific, you know, capabilities.

And overall I think that this are -- there is a need to evaluate the overall hybrid, you know, human plus machine system rather than evaluating the error of or the behavior of decision making capabilities of the machine alone. We should never think of the machine alone. But the two together which brings in to place also these capability of interacting of building trust between the human and the machine over repeated interaction over time.

>> URS GASSER: So Francesca's presentation also highlighted some of the ethical challenges and questions. And I think that's a terrific segway in to the next presentation. And I am really delighted to invite Salil Shetty to join the podium. He will talk about the idea of new ethics for AI that is rooted in Human Rights and Human Rights principles. So there is a really strong link between these two presentations.

So Salil is the Secretary-General of the other AI, the Amnesty International movement where he leads the work to end Human



Rights violations in many parts of the world. And that's another connection point. He has done tremendous work to bring Human Rights agendas and issues at the forefront also in the global south. Before joining AI he was a director of the UN millennium campaign and was deeply involved in the promotion of the Millennium Development Goals. Salil, without further ado.

(Applause.)

>> SALIL SHETTY: Thank you. Thank you all for giving me this opportunity. I must be one of the few sort of interloopers in this audience of technologists and policy experts on technology. I was going to mention, the post-lunch session is always the most difficult one. So thank you for bringing us in on this one. But Francesca's talk has certainly woken me up. That's the worst place to be sitting because the sofas you start reclining. Do keep an eye on these people to see if they are awake.

One of the biggest problems we in Amnesty International have with AI we have the same name. I am often confused when people talk about AI inside of Amnesty as to which AI they are referring to. We have two AIs, Amnesty International and Artificial Intelligence. You have heard and you will hear a lot more in the course of these days about the amazing things that Artificial Intelligence could do in the future.

I want to talk to you more about Human Rights today and tomorrow. So no surprise on that front. And I want to talk to you about this in a kind of setting which makes it clear that as far as Amnesty is concerned we are very clear that there is huge possibilities and benefits from Artificial Intelligence.

So the AI for good is an amazing idea of having this conference. The question is who is it good for, and we are going to talk more about it in the coming period. I think in the eyes of the general public and certainly within Amnesty International as well it is difficult for people to connect what we do as kind of bread and butter work. Our bread and butter work as you can imagine is freedom of expression stopping torture, that's the sort of work that Amnesty typically does. But we are looking more and more equally at the issue of Human Rights in the workplace, access to health care and education, economic and social cultural rights as well.

And in that context I want to just refer to one anecdote and I want to come back to that at the end of my speech and this is I think two years ago I was in a rural part of the Czech Republic. And I was taken to a school to look at the situation of Roma girls' education. And this is a Roma family. And I spent some time talking to this girl. She must have been 10 or 11 years old and I think her name was Dinka. You could see from the eyes how bright and how enthusiastic she was. But at the same time the

story that she had to go through A to get in to that school as a Roma girl to get in to a school in the Czech Republic you can imagine this part of the European Union is not the easiest thing to do. And as soon as a group of Roma children had managed to get enrolled in to the school, a lot of the non-Roma children left because they don't want to be in the same school as the Roma children. So let's keep Dinka there and we will come back to her later.

I do believe that the discussion on Artificial Intelligence is at a fork in the road. And I think we have clear choices in front of us. And that we want to be discussing these choices. And one of the things which Mark McGanda had talked about from our country was the concept until there. And the until there concept whenever you want to make a decision the way in which you could make it is by thinking about that last person as to what does that last person -- how does it affect that very last person. Who is at the end of the row. And in some ways Dinka, the Roma girl or a girl in Afghanistan that's a good kind of check in to see as to what does it mean to this girl.

So humanity may live on one planet at least for the time being. But we inhabit different worlds. At one end is prosperity and wealth, technology that's continuously evolving to fulfill any need we can imagine and the future prospects of incredibly long and healthy lifespan, space travel and much more. At the other end is the reality of poverty and injustice that has been scarcely improved despite the arrival of the Internet and mobile and Big Data and now Artificial Intelligence and in between them is one of the biggest threats of health of our society and that's inequality.

If we look at global development indicators there is a very positive picture. Extreme poverty, undernourishment, access to education, clean water is increasing. The world has achieved significant progress in the past 25 years. But at the same time economic inequality has been rising. And I always also talk -- when I talk about economic inequality I always try to refer also to voice inequality, that those who are left behind have no voice. Even in the OECD countries including the wealthiest countries in the world, income inequality is at the highest level in ten years. Nine times of that poorest 10%. Up to seven times from 25 years ago. These numbers have been thrown at many of you.

The one which came out earlier this year, that eight individuals, the wealth of these eight individuals is more than half the population of the world which with 3.6 billion people have as much wealth as eight individuals. Around the world from Asia to America and from Europe to Middle East we are seeing what happens when inequality grows and it is ignored. Anger

simmers and the poetics of blame thrive, discrimination against women, Muslims, black people, states, dividing the world. Inequality -- when we need to work together to solve the global problems we face but growing inequality is neither an accident nor is it without consequence.

Inequality and political instability generates direct consequences of human choices, Government policy and corporate practice. The way companies and Governments handle technological innovation is a key part of this story. To take one example, one of the most important phenomenon over the past few years, 1.3 people working in the big economy and they have no guaranteed working hours or sick pay. They lack many of the rights that workers in regular jobs enjoy. And at the same time technology companies have made billions by providing platforms for those so-called flexible working arrangements. Of course, businesses can and should innovate but are they paying their fair share? If you look at growing corporate profits together with wide ranging cuts to health care, education and other public services even in the west, we know that something isn't working. Tax evasion tactics have allowed hugely profitable companies from paying very little tax while whole job categories disappear and are replaced by low paid insecure jobs. This is a rule in life. If you take and take and don't give back people will get pissed off. Some will argue that for business what matters is growth and profits. And this has to be their priority. But then they can't complain of the growing disillusionment and reducing trust levels in business and governments.

Societies must plan for and mitigate against the risk. And this is critical in the case of Artificial Intelligence which has the potential to create huge economic disruption. As economic inequality grows there is another kind of technology, fuelled inequality that's growing at an exponential rate that relates to data. There is already a huge asymmetry in power between companies and Governments and individuals and Civil Society on the other. Control a handful of companies and Governments exercise over unimaginable amounts of personal data. Whether it is to sell as ads for electronic surveillance programs data gives formidable power to those who control it. And we heard how Big Data analysis and micro targeting of waters was used in U.S. elections and the Brexit referendum.

While the use of data in politics is not new the difference is how powerful these techniques have become in a short space of time and how Artificial Intelligence could super charge all of this.

So what will happen tomorrow? Let's fast forward 20 years. What world do we want in 2037? Technology is chief among them, Artificial Intelligence will shape tomorrow's world. That's a

certainty. And I don't want to present a binary picture but let's for the purpose of this conversation visualize two scenarios. If we continue as we are, hundreds of millions of jobs can be lost to automation and largely replaced with insecure jobs with little protection of workers' rights. We may have some social protection schemes but they will allow people to survive with little hope for the future. Artificial Intelligence will be used across the board in health care, education and across public services. Robo cops will patrol streets. And wars will be fought by killer robots reducing the cost of war for the wealthy and powerful but not for the poor. We already know how data driven systems from financial to predictive policing applications can discriminate against individuals. AI systems will become the gatekeepers deciding who can access health care and not. Those are power and access to the fruits of the data economy a handful of countries and companies will be those that continue to gain while the vast majority of people are left behind. This brave new world may entrench global inequality on a scale never imagined. And this will come with massive upheaval and disruption. There is a possibility of another world and this conference and initiatives which follow from here offer as an opportunity to take a critical step in identifying principles for the ethical development and use of Artificial Intelligence.

Companies could take notice and Governments could respond. In the future we could have Artificial Intelligence systems that detect and correct bias and data, rather than doubling down on the human bias. We could have automation that takes people out of dangerous and degrading jobs but also educational and economic policies that create opportunities for dignified and fulfilling jobs. And Governments could ban fully automated weapons systems so that killer robots never come in to existence.

This is a future where enormous power and potential of Artificial Intelligence is harnessed for the good of humanity, promoting equality, freedom and justice. It is a future where open source AI allows innovators across the world to harness the power of technology where explainable AI is developed and used allowing for AI decisions to be interrogated and challenged and with clear legal accountability systems to ensure that the rights and responsibilities of users and developers are clear. In short it is a future. AI is a technology where Human Rights principles is the same as a core design and use principle.

But this can only happen if we start out now with an objective to protect human dignity, equality and justice. AI is built by humans and it will be shaped by human values. If we build AI systems that are a mirror to our current societies they will be

riddled with historic biases and inequalities of our societies. But we can do better. When the states, they were not reflecting the world that we live in. But an aspirational world. We must today challenge ourselves to be aspirational again as we prepare for a future world where AI and technology are integrated in to every aspect of people's lives.

Fortunately we already have the Sustainable Development Goals, the 2030 goals which almost every country in the world is signed up to the United Nations just two years ago and a core principle of the Sustainable Development Goals is the principle of leaving no one behind.

So this is a principle which would guide us as we think about the discussion here today. And it is heartening to see the interest and dedication to ensure that the development of AI is ethical. We believe that such ethics must be based on Human Rights principles. These are universal principles that have developed and matured over nearly 70 years and have been applied in the national context across the world. Governments have binding Human Rights obligations and corporations have a responsibility to respect Human Rights. We strongly believe that enshrining AI ethics and Human Rights is the best way to make AI a positive force in our collective future.

Amnesty International exists to bring about a world that Human Rights and freedom are enjoyed by everyone everywhere. As the world's largest people's movement of Human Rights, today our work is as much about campaigning for the release of prisoners of conscience and the protection of the world's 21 million refugees as it is calling for a future to promote freedom, dignity and equality and justice.

And so I take this opportunity to announce Amnesty International's Artificial Intelligence and Human Rights initiative, I invite you to collaborate over the coming months and years to enshrine the protection of Human Rights in the development and use of AI.

So going back to Dinka who I referred to at the beginning of my speech, let's set our mind to think about what happens when Dinka grows up. What happens when our children and grandchildren grow up and how will they judge us on what we did and decided today. Thank you.

(Applause.)

>> URS GASSER: Thank you so much, Salil, for a fantastic talk. You made many important points. And I hope we get a chance to discuss them with you also after this Plenary. One point you made is to clarify that also Human Rights and the principles have a dynamic nature. It is not a static concept. It is an evolutionary concept itself. And so looking at it from the other way around, do you see some sort of an opportunity

there that AI may actually help us to re-envision Human Rights whether it is in the implementation which is more obvious by looking at the principles themselves and that against the backdrop, against discussions that we had over the past decade in the light of the digital revolution do we need to come up with new types of Human Rights. Do you see similar dynamics applied in the future for AI?

>> SALIL SHETTY: Sure. There is nothing static about human rights. It is constantly being interpreted. The UN system and other bodies are constantly adapting and interpreting what it means in a very practical sense. So absolutely. But the basic principles, you know, of nondiscrimination of equality of justice, of fair play, of humanity, those I think are core values which you in a sense those are the guiding values and those are the ones which you really should think twice before you compromise on it.

For example, we take what's happening in Europe around the elections more recently. You have extreme sort of populous parties pushing for anti-refugee rhetoric and many of the mainstream parties then try to start sounding like them in order to norm, to -- that becomes more of the mainstream. And I think absolutely you have to adapt to the reality in your context but you don't compromise on the basic principles.

>> URS GASSER: The theme of inclusion is also actually a perfect segway in to the next talk by Professor Vicki Hanson, who serves as the President of the Association for Computing Machinery or ACM in short. She has many positions. I'll just highlight two. She is a distinguished professor and codirector of the Center for Accessibility and Inclusion Research at the Rochester Institute of Technology and holds a chair of inclusive technology at Dundee in Scotland and led the foundation at the IBM Watson research center. Vicki, I hope you can expand on this key concern, how do we include or create an AI ecosystem that is inclusive and doesn't leave out the last person. Yeah. Thank you.

(Applause.)

>> VICKI HANSON: Thank you. That was a hard act to follow. There is no doubt about it. So yeah, you already gave way to my opening line here. I was going to say that I was asked to say a few words about ACM, which is the Association for Computing Machinery. I had a discussion at lunch. This is not the Academy of Country Music. Those of you looking for that I do apologize in advance.

So ACM is the oldest and largest society for computing researchers and practitioners. And even though it is headquartered in New York you can see this is a very global organization. This is a map of where our members are from around

the world. And ACM is a volunteer driven society. So the activities that we do are all based on the interests of the researchers, practitioners and educators who are our members. Our main strength is our technical leadership. And just going to show you a little bit -- oops about this here. The clicker is a little -- got a bit of a delay built in here. So we sponsor about 200 conferences a year around the globe and are in collaboration with many more. The proceedings from those go in to our digital library. And I wanted to say that ACM has been working to take a lead on a new conference on AI and ethics which has become an important topic today. Among those special interest groups we were told that ethics is the most critical challenge they are facing these days. That was the motivation behind getting together a new conference.

And as another example of our leadership one of the things that we have is a major awards program and our largest award is the ACM Turing Award. This year it is being given to Sir Tim Berners-Lee for his work in inventing the World Wide Web.

This is a very exciting year for us. This is the 50th anniversary of us having given the Turing award. And as a big anniversary we will be celebrating the award next month in a day and a half event that has technical discussions from Turing laureates, from our own prize winners and from dignitaries throughout the field. And they are going to be talking about topics relevant to AI. There will be a panel on advances in deep neural networks and one on challenges and ethics in computing.

And I mention this because that event will both be live streamed on June 23rd and 24th, but it is also going to be videotaped and made later. So maybe some of you could go to the ACM website and take advantage of these discussions that we are going to have by a lot of luminaries in the field.

I want to thank you for inviting me to this Global Summit. Thank you Stephen Ibaraki and the other organizers for having me here. This is clearly a critical time for AI. And I will say that Artificial Intelligence is not my area of expertise. So I was excited to be invited here to learn and I have already learned a tremendous amount. Looking forward to the rest of the event.

The reason that I have found this whole topic of AI for good so inspirational is, for example, this sentence that I found in the report on the 100 year study of AI, saying the measure of success for AI applications is the value they create for human lives. And this is something that I believe deeply in. My own background is in human computer interaction, not AI. So I'm particularly interested in how humans and machines interact with each other now and in the future. I am particularly interested in how machines will augment the abilities of humans. My only

background is particularly in accessibility research. So I'm interested in people with disabilities. And as you can see in this slide disabilities come in various forms. So this shows the fact that sometimes a disability is a visual impairment. It could be a mobility impairment, inability to move a hand or someone who is unable to walk. It could be a hearing loss. Or it could be a cognitive impairment. Some problem of thinking or processing information. And very typically these days what happens with aging there is a combination of these impairments that people may have. So someone may have a visual problem. They may have problems with their hands and all of this makes it very difficult to interact with the world as soon as you get a complex of problems.

At a very high level these disabilities create problems with communication, with mobility, with life skills, and with technology.

I'm just going to say a couple of words about each of these. So in terms of communication as we all know the ability to communicate is really a major factor of being human. We need to talk to each other in terms of our professional lives. It is important. It is important for entertainment. People who have problems with a hearing loss or aphasia, cerebral palsy, they affect a person's ability to communicate and interact with others. Mobility, we all want to move around the world and we want to move independently. So visual impairments, in fact, make it very difficult for someone to navigate in the world independently, particularly in a new environment. People who have a cognitive disability have trouble navigating the world because of getting lost. You think of wheelchairs. If they are in a place that's not wheelchair friendly it is difficult. Sometimes more surprising is people who are deaf or hard-of-hearing also have trouble going to new places because of the fact that it is difficult to hear environmental sounds, emergency sounds like an ambulance, for example, or it is difficult to be in a hotel because the hotel may not have a fire alarm system that would alert them if there were a fire in the hotel.

And life skills, this is just the fact that there are certain tasks that we have to be able to do for ourselves day-to-day to stay out of the care facility. And one of the things that people are starting to talk about these days is the ability to have some kind of digital literacy and interact online with all the Government services that are coming out and that may be a daily activity. The fear that people can't do the basic skills of daily living, people may have to go in to care facilities. That's the normal accessibility thing that people talk more on. I am not going to go in to it for a lot of time.



That's the kind of thing people with motor disabilities, for example, use a mouse or keyboard. But one of the things I want to say is I have been using the words impairment and disability. The World Health Organization defines impairment as a health or body problem. So, for example, someone might have a hearing loss. However disability is a complex set of problems and it is actually created by a society. So someone has a disability when they have a hearing loss and somehow can't interact in their environment. They can't hear the danger sounds. They can't talk to someone. They can't hear something on the computer because it is all in sounds. People with a hearing loss are disabled in environments in which information is presented only as sound. And so if we are developing applications with AI or any other way, that are using sound only, then we are disabling people who have a hearing loss.

Okay. Why is this important? I have been talking about disability. It doesn't really seem perhaps that maybe there are that many people with disabilities in the world, right? You probably don't see that many in your day-to-day life. Actually there is one billion people in the world who have a disability of some sort. Why don't you see them? In part it has to do with aging because as we grow older more disabilities do happen but it also has a lot to do with the fact that a lot of the people in the world who have a disability seem more less hidden. They don't go out in society. Often they are at home.

One of the things that I was going to say is that this issue of disability and inclusion relates to the UN Sustainable Development Goal in reducing inequalities, No. 10 over there on the chart. It is a little bit different than the kind of reducing inequalities we heard this morning, but there are huge inequalities for people with disabilities right now.

So, for example, people with disabilities experience poorer health outcomes in the general population, lower educational achievements, significantly lower employment. The employment rates are only about half for the -- for people with disabilities as opposed to the general population and much higher poverty levels. So there is a huge disparity for people with and without disabilities.

So my goal here is just to do some awareness raising of accessibility among those of you here who are working on AI so that this kind of disparity doesn't get exacerbated in applications that are built using AI.

I just want to give you two examples about how technology isn't always a good thing for people with disabilities. And then talk a bit more positively about things that can be done. So just to give a first example, I currently live in Rochester, New York, which is the home of Kodak photography. I used to think

that Kodak had to do with film. It is mainly the paper that photographs were made on. People sent their photographs in and they developed them. Rochester was full of buildings that had dark rooms in them. Who worked in the dark rooms? It was largely populated by people who were blind. When digital photography came out which was invented by Kodak and they couldn't figure out what to do with it. When digital photography came out all those dark rooms were put out of business. And all those workers who were blind and visually impaired lost their jobs.

And another story that's sort of similar, you probably all know that Alexander Graham Bell invented the telephone, right? And he did it because his wife was deaf. And he actually wanted a way for her to be able to feel sound so that she could communicate better with other people. It had an unintended consequence. The telephone became the greatest disabling technology for deaf people. They couldn't have a high paying job because they can't use a telephone and engage in calls. There is a lot of work that goes on in developing technology that doesn't think of the consequences of disabled people along the way.

Okay. So with this background in mind there are some examples where AI has been helpful. Here is a current example. This was announced a couple of weeks ago. Facebook, I realize the print is so small. The one on the left we finally made it. If you are a blind person, you have no idea what that is about. On the right it says a Saturday night splurge. Someone who can see can tell what the -- what it means to finally make it. And on the right you can see what the Saturday night splurge is. Someone on Facebook got the idea to use image recognition. And now what's happened with Facebook here is if you are a blind person you can get the translation, you can get the image recognition and find out what these images are about.

Another thing that's big in disability has always been the idea of independence. And I remember that virtually every blind person that I met in my life said they wanted to drive a car. I always said this wasn't going to be possible. It turns out there is some tests going on and it is a little hard to see it, but the gentleman in the middle of the screen who has a white cane and he is blind. And they are starting tests so that people who are blind can do the autonomous driving. And this is a dream come true for many, many people.

I am not going to talk about robots because I'm running short of time. But this is something that's very popular these days. They are being talked about is the future of AI for elderly people living on their own who experience a lot of social isolation.

And in closing I'd just like to say -- just like to say, if I

get to the right page, that I want people to just think about the idea of developing for a wide range of abilities when thinking about new AI applications. AI discussions often focus on the technology itself looking at things that are faster or more integrated. But these things also need to be accessible and useable to a wide range of people. We have all heard stories about -- this has been talked about before about the bias and datasets. And one of the things that I really hope gets included in some of these datasets in the training is that more people with disabilities have input so that the training will include a wider range of ability.

Last sentence, so other people here I know are going to talk about these data driven and algorithmic biases a little bit more. But consider data from people with diversity in AI development so as not to disadvantage the billion people worldwide whose needs differ from those of healthy young adults. Thanks.

(Applause.)

>> URS GASSER: Thank you, Vicki. A very quick question only since we are running a little bit short in time. So given the nuance picture you painted and also the helpful reminder that there are massive inequalities and access barriers, this morning where we were energized with this vision of we will make the world a better place and everything will be great and democratized, do we need to engage in expectation management? Do we really have to soften our message or tailor our message?

>> VICKI HANSON: Okay. Tailor our message, yes, I think it can be a very positive message. This particular population gets left out a lot in thinking. It has been left out a lot in past development and I don't want that to happen again as we start all new technologies.

>> URS GASSER: Thank you. Our last speaker who promised he will keep us right on time, which is amazing, thank you. It is now. It is Professor Gary Marcus who is a Professor of Psychology and Neural Science at NYU and published on many different topics ranging from human behavior, neural science and AI. And he was previously the CEO and founder of a machine learning startup Geometric Intelligence. And he is a frequent contributor to the New Yorker and New York Times and author of several best selling books. And it would be delightful to have your final remarks of this panel at least.

(Applause.)

>> GARY MARCUS: I have no slides here. I see myself but that's not what I want to show you. I don't know what's going on with the clicker here. I see. While we wait I just want to say it is an honor to be sharing the stage with some of the people who are speaking so incisively about human dignity.

And here we go. So I'm going to speak today from the perspective of someone who has worked in the academy as a cognitive scientist, professor for a couple of decades and also in industry as a founder of a company that was just purchased by Uber. Wendell tells me I was the first person to write about the driverless car. And you can take seriously what I have to say. I don't know.

My aim today is to provoke. I am going to start with some contrarian statements on AI and I am going to conclude with suggestions of what might be necessary to move forward if we are going to achieve AI. Where is the field right now? We are not nearly as close to strong Artificial Intelligence as many believe. I think there has been a trend throughout the day where we started with a lot of enthusiasm. And there has been a little bit of skepticism filtering through. And I will keep the trend in that direction.

So here is some optimism from Andrew Yang who was the Harvard business review. If a task takes -- Andrew is half correct, there are some things that people can do in less than a second that AI is good at. So, for example, you can probably distinguish between Tiger Woods and a golf ball and so can your favorite deep learning net. AI is pretty good at recognizing objects. Typically using convolutional networks that were developed in part by Yoshua and so forth.

The basic idea is you have Big Data in and you have statistical approximations out and it works pretty well. I have used one of these apps that does deep learning. And this was in my hotel room in Hong Kong over the day. Water bottle and water bottle and third one on the right it came up with pen. This app that does recognition. If you probe more deeply there is something not quite right. Here you can describe this picture and you might come up with something like a group of young people who are playing a game of Frisbee. And you might look at the next one is a person riding a motorcycle on a dirt road, but if you saw this one you would not say that it was a refrigerator filled with lots of food and drinks. If you were a patient in an Oliver Sacks book but you would not if you were a human being. Yann LeCun and I have these debates. My line is the hallucinations and they are part of where we are right now.

Deep learning has got at certain aspects of perception but perception is more than categorization. And cognition is more than just perception. So to achieve its destiny, AI is going to need to go further than we have already. Whoops. So a good deep learning system might be able to identify the dog at the bottom or might have trouble because we don't see dogs are ears in this orientation. And if the model doesn't -- the real point you can tell that it is a dog holding up a barbell. And if you are a

human being you have thought how did the dog get so ripped. If you were my four-year-old and say who is the elephant standing on a tight rope. The way I like to think about it a cognitive scientist is there are many things that go in to intelligence, perception, common sense, analogy, language, reasoning and so forth.

And what we have made real progress on we meaning Yoshua and not me, what we have made real progress on is perception but the rest we haven't made that much progress in the field. In terms of understanding what problems can we actually solve in the humanitarian domain. If a typical person can do a mental task we'll probably automate it using AI. What you are doing in your humanitarian organization if you can get it done fast a machine will do it for you. So we have labeled examples. So if you don't have the labeled examples then the techniques are not going to work. You have five or ten. The techniques are not going to work. If you can do less than one second of thought and we can gather an enormous amount of directly relevant supervised data, we have a fighting chance.

So Go is exactly opposite. If your problem is like Go we can solve it. You can gather as much data as you want and keep iterating over and over again. But if your problem is not like Go, like it is politics in the real world we may not have the tools for you yet.

The stuff I have here in green we do pretty well. Speech recognition especially if it is in quiet rooms with native speakers maybe not so well. We can image recognition as long as the world is pretty bounded. We can do natural language understanding in narrowly bounded domains. Siri can tell you what time a movie is playing but not answer general questions. How much do transistors cost and this is eluding the exponential when the Mac plus was released. Wikipedia can tell you when the Mac plus was released, but no search engine that will come with those two pieces of information together which we would hope for in a general conversational AI system. And AI is great at advertising and targeting, but I think that Yoshua and I would like to see AI do more than that.

Automated medical diagnosis, we are making progress there on the visual side but less so when we have unstructured text from doctor's notes and so forth. We see a little progress. I would love to see how well the Facebook stuff does. My guess is it is not that hard to fool. Domestic robots for elder care but the eye is nowhere near strong enough yet. We would like to see driverless cars but safe, reliable driverless cars are farther away.

Here are some impediments I think to reaching strong AI. First is engineering. Machine learning is hard and difficult to

devise. You made modules and bigger modules. And you put them together to make larger modules but we can't really do the same thing in machine learning. And the talk at Mtech a couple of years ago, Peter Novec will be here by Skype, the way that things work now your data works and you test the model on Tuesday. And it seems to work on Thursday. And then suddenly it is Christmas and all the assumptions break and the model doesn't work anymore. We have no procedures for reliably building complex cognitive systems yet.

This is by John Scully's nephew. The idea you can build the systems and they work in certain circumstances but you don't have those guarantees when you change the rest of the system downstream it is going to work out. XC, this is your machine learning system says one person. You pour the data in and collect the answers on the other side. What if the answers are wrong? You are just supposed to stir the pile until the answers look right. And here from the ACM is I recently had a piece that was on the cover with Ernie Davis. The artist drew a robot sitting on a tree branch, and the point I want to make statistics is not the same as knowledge. You have a lot of robots cutting a lot of tree limbs. You don't want a single robot to cut the wrong side, the tree limb fall down and hurt itself and possibly other people with its chain saw. You can't solve this problem with 50,000 labeled examples.

The third problem there is bias in the field to assume that everything is learned and I don't fully know why this bias persists but it is pervasive. Yann LeCun, he wants to make -- pour in a lot of video and everything will emerge from unsupervised learning. And I wish him luck. I doubt it will work.

I was trained in developmental psychology and I have been improving and pariting. And the best arguments are arguments that say look, the brain of the human being starts with something in it, an innate language device and sets of places and things like that in. Nobody ever believes me when -- I started making them about baby Ibexes. Here is a baby Ibex and it is climbing along the hill. If it falls once it is out of the gene line. What about building something in? But it is not learning trial by trial in the way that our contemporary machines. We need more innateness if we are going to build intelligent agents. We show you our robots which don't have a lot of innate structure and don't do as well.

So we lost the sound but that's okay. There is a nice little piece of ragtime here, but you get the idea even without the sound.

So we heard a lot about exponentials this morning. I want to advise a little caution. So in some fields of AI, in narrow AI

where you have that tightly constrained problem the rules never change and so forth. There has genuinely been exponential progress. And you can plot it by looking at scores on chess. There is no data on strong AI. And here we have Eliza in 1965 and there hasn't been progress in general Artificial Intelligence. If we are going to fulfill the destiny of AI we want to get the strong AI. There is a lot we can do now. There is a lot of fruit to be gained in the next ten years and also a lot of things where we want genuinely intelligent machines.

I give you a proposal with three premises. Almost certainly going to -- you go back to my pie chart. There are many different components to intelligence. Many different components to Artificial Intelligence. We need people working together and psychologists and people working in classical AI and working in neural networks and security and program verification all working together. It is not going to come out of one lab. The human brain is too complex to be understood which one individual and true AI probably is as well.

So we are going to need I think a really strong interdisciplinary effort. And the second thing, in the ideal world AI would be a public good, not something that's owned by one corporation or eight individuals or something like that. But we are headed on a path where that is what's going to happen. And there are companies like Google that are busy patenting very basic ideas in AI, some that I think are indefensible ideas. If you are a startup company, is Google going to sue me for using this idea that's twice as old as Google. Google don't use their patents defensively. They are never going to sue or many of us know that Google has started some fierce intellectual property disputes recently.

The world has changed. And it could change more. It could change where these eight people who own AI or whatever it might hypothetically be aggressively pursue intellectual property claims and no existing approach to AI approach can get us to the next general -- corporate AI tends to focus on what can we do with the techniques that came out from Yoshua's lab. How can we commercialize the stuff that's most recently been discovered and academic labs seem to work too independently of each other. A bunch of small academics may not be enough to get where we need.

I am going to make my proposal. The proposal is let's look down the street from where we stand right now to CERN. It is a global collaboration from thousands of researchers to build technology and science that could never be built in individual labs. Maybe we need to have a model like that for global AI. Lots of people doing AI for the common good.

Two more sides. Wendell asked us what is it that -- what's one detailed problem we can work all on together. I want to give

a meta answer to that. I have been thinking about how to do AI myself and I don't know which of the many problems that are out there I should address myself to. And what I have been wishing for and maybe someone can help us with is something like charity navigator and not for the individual, how should I spend my thousand dollars or five thousands and who can move the levers. People like me who do AI research don't know about. We need some help from other people in the room. Where is there data available. Maybe there is some way to sort of metaize to coin a terrible word Wendell's question. I will send you with this wonderful African Proverb. "If you want to go fast go alone. If you want to go further go together." Thank you very much.

(Applause.)

>> URS GASSER: So thank you so much for a fantastic closing talk. We have reached the time for a coffee break. We would like to ask the panelists for a Tweet link statement where you are optimistic or pessimistic about where we are heading to in our field. Who would like to start? Salil.

>> SALIL SHETTY: I am in the business of optimism. If Amnesty International starts getting pessimistic the world will not be in a better place. It is AI for good but good for everyone and not for some.

>> FRANCESCA ROSSI: I am optimistic because I think that AI will be shaped by humans. And I am generally optimistic about what human causes are.

>> VICKI HANSON: I would say I'm optimistic now. So much excitement and different things can be done and all kinds of creative ideas are going to come out of this new work.

>> GARY MARCUS: I am optimistic about the opportunities and inertia won't get us there.

>> URS GASSER: A generally optimistic note with a caveat. Let's thank our speakers. Let's have coffee.

\*\*\*

This is being provided in 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.

\*\*\*