

FINISHED FILE
AI FOR GOOD GLOBAL SUMMIT
JUNE 7, 2017
9:30 A.M. CET

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.

>> STEPHEN IBARAKI: Okay. I want to welcome everyone. If we can have the mics open. Welcome to the AI for Good Global Summit. If I can get you to take your seats. And if we can ask the participants in the hall to come in and take their seats, thank you.

While we are waiting we'll play a video to provide a background on the Summit. If we can start showing the video, please.

(Video).

>> 9 million people alive in 2050 we don't know how we will provide them clean water, sustainable energy, education.

>> There are serious problems with the environment. There are -- just the various factors in a way that can solve these problems.

>> How we help the 7 billion people that could benefit from these technologies in a practical way.

>> While we can't say that AI is general magic, it is giving us core tools to start putting in place systems that we know will help with key issues, like hunger, famine, overall health and well-being.

>> AI offers an extraordinary possibility of accelerating solutions in the time we have to solving the problems.

>> We are now applying Artificial Intelligence to the art of farming, also adopting the latest in technology. Artificial Intelligence, a farmer integrates masses sources from the equipment. Data from the fields taking all that data together actually making recommendations to the farmer on how he can increase the production of his land.

>> AI has enormous capacity and potential to.

>> The molecules and how they interact with the virus proteins and then doing a search to find the best product.

>> Sequence your entire genome and take that along with your medical imaging data and compare that with the world's data of genome sequencing data and determine the matches between your tumor and the others what the treatments were in those other cases and the outcome, did the patient survive and do that entire process in one day. It is possible through Artificial Intelligence.

>> The UN Summit and its goals excite me, in particular with the idea of leveraging what I would call sleeping giants of data. We have so much data in the world and with some refinement and tagging this data can become available and leverage to address famine and economic poverty.

>> So machine learning algorithm find tomorrow about making for better medical diagnosis and make a better --

>> Next form of civilization. It is not happening in 10 years or 15 years. It is happening right now.

>> STEPHEN IBARAKI: Thank you. A most enjoyable video to set the context for our next three days. Again I want to welcome all of you to the AI for Good Global Summit, co-organized by the ITU and the XPRIZE Foundation. And in partnership with 20 UN organizations. I'm Stephen Ibaraki. And I am the Moderator of this welcome session.

In Day One we will do the context setting by talking about Moonshot inspirations for the future of AI, AI state of play, transformations on the horizon, future roadmap and collaborating for good. On Days Two and Three we will do guideline developments with breakthrough sessions. We want our participants to choose one from each track, that is privacy and ethics and societal challenges, common good and sustainable living, capacity building and poverty reductions and investments and designing the future. That's between Day 1 and Day 2. Setting the context and also the breakthrough sessions where we are doing the guidelines. We are fortune enough to have a welcome address from the United Nations Secretary-General Anthony Guterres. Let's now show the video.

>> ANTHONY GUTERRES: To discuss the challenge of growing

global importance. As someone with a background in engineering I am deeply interested in the remarkable phenomenon of Artificial Intelligence. And as Secretary-General of the United Nations I am strongly committed to promoting growing. It is transforming our world, socially, economically and politically. We face a new frontier with advancing moving at warp speed. Artificial Intelligence can help analyze enormous volumes of data, improve prediction and improve crimes and Government. Ethical issues at stake. Real concerns about cyber security and Human Rights and privacy and not to mention the obvious and significant impact on the labor markets. The implications for development are enormous. Developing Countries can gain from the benefits of Artificial Intelligence, but they also face the highest risk of being left behind.

This Summit can help ensure that Artificial Intelligence charts a course that benefits humanity and bolsters our shared values. And United Nations stands ready to be a universal platform for discussion. Together let us make sure we use Artificial Intelligence to enhance human dignity and serve the global good. Thank you.

(Applause.)

>> STEPHEN IBARAKI: Again thank you to the Secretary-General for his welcome address. And just a note by the way the welcome session and Plenaries today are available in six languages, being interpreted and translated. So you have an earpiece by your seat. So you can listen in the six languages.

It is now an honor and a pleasure to introduce the ITU Secretary-General, Mr. Houlin Zhao, who will provide his welcome address.

(Applause.)

>> HOULIN ZHAO: Good morning. Welcome to the AI for Good Global Summit. Last year the Artificial Intelligence program took the world by surprise when it defeated grandfather, grand master, Lee Soto of South Korea. Just ten years ago it was the world's No. 1 player to meet the same fate. Yesterday I was in Helsinki. I had lunch with the Chairman of Nokia, would not really put any agenda. But he raised the issue of Artificial Intelligence issues. And we talked more than one and a half hours. And he told me that last year this Go match seems to be a surprise, that people still have a chance to defeat the machine. That machine looks like human players. Very intelligent. That's still like human players. But this year the top players has no chance to win any single match. So the machine looks like God. And he mentioned to that with this explanation of development for intelligence what the next -- whether we have very powerful machine to treat the patient which can treat the patient in a better arrangement than our human doctors. And then he asked

another question, what is next in the future. That great judge, whether we can have judge played by machine and to have great judgment.

And he also mentioned that while the technologies developed so rapidly with exponential power, our human society, our Government unfortunately has not matched the speed. So that he raised the issue and he told me this is a very important issue for Nokia. Very important issue for industry. So I told him that last year, actually not last year, February this year, at the beginning of February our Secretary-General of the United Nations called me. We had a phone conversations about 40 minutes. He raised his request to me that we should provide him some briefings on Artificial Intelligence. So I told the Chairman of Nokia that it seems to be that the UN Secretary-General thinks the same way as the Chair of Nokia, which means that the United Nations is a good friend of industry.

And we are good partners. And we see the same challenges. We see the same opportunities. And we would like to work with the industry. And the Chairman of Nokia told me absolutely, he agree with that.

Today we have gathered here to discuss how far Artificial Intelligence can go. How much it will improve our lives. And how we can all work together to make it a force for good. I want to thank all those who made today's Summit possible. Thank you to our co-organizer, XPRIZE, CU, Marcus and to our gold sponsors The Kay Family Foundation and Wiki Omni Word4app and our corporate sponsor, PricewaterhouseCoopers.

And at ITU we have been working closely with industry members from Day One. It is also good to know that ITU has 152 years history. We were created in 1865. Nokia was also created in 1865, the same age as ITU.

But they have not been engaged with telecom from very beginning. They engaged with the paper industry. But anyhow we share the same honor of our history.

We have a long tradition of cooperation that dates back more than 150 years ago. But on the other hand, ITU also keep a young spirit of innovation. I think that is the reason why, you know, ITU still relevant to our human society to technologies. Many of you may not know that but each time you use your mobile phone or go on the Internet it happens. Thanks to the devoted work of ITU for technology standards.

And today's famous ones 4G, 5G, but you also have to note that when you use the Smartphones those video images, video traffic only possible with ITU's standard for H264 which was developed more than a decade ago.

And one of the pioneers, researcher are in this room, Thomas.

I don't know where he sits. Yeah. And we will award him as one of five individuals for ITU's 150th award.

We are fortunate to have many ITU industry and telecom members here today. I'm also very happy to welcome the representatives of several UN agencies. About 20 UN agencies are partners of this AI Summit. I'm particularly grateful to United Nations Secretary-General Anthony Guterres, who has just shown us his video message. And from the beginning of this year I contact him several times for several occasions to get his video message because he cannot make it. He made it yesterday and he considers this absolutely important and for UN and he did not consider our event as an ITU event. He considers this event as an UN organized event, the first UN organized event. So I'm really grateful for his support from the very beginning.

We are also very honored that he could be with us with his video message, but he also designated Madam Izumi, the UN High Representative for Disarmament Affairs as his personal representative to the Summit. Welcome Ms. Nakamitsu.

AI will play an important role in helping to achieve the United Nations Sustainable Development Goals. Upon the recommendation of the Secretary-General of the United Nations Anthony Guterres, I will hold informal consultations with participants to see how the United Nations, of course, ITU as well and other UN agencies can work together with the industry and the academic community to promote AI innovation and to create a good environment for the development of Artificial Intelligence. This dialogue between our organizations, I mean the UN agencies and the industry is essential.

And we did have an internal discussion among Heads of UN Agencies in March. And those who represent the alternative discussion all show their interest to support Artificial Intelligence. Of course, there are some concerns with some proposals and I think that this dialogue between UN agencies and industries is absolutely important. And I'll tell you another secret, the Secretary-General of the United Nations was supposed, proposed to establish a special advisory committee for Artificial Intelligence and he called this a frontier issue.

I'm supposed to start this work from the end of March. But he clearly suggested that let's wait until this Summit is held which means he also counted on participants of this Summit to provide valued input to him to see how UN and UN agencies can work with the industries.

We look forward to close cooperation with industry members. This Summit is the beginning of a new journey. It is a platform for all of us to develop concrete proposals for leverage of the Artificial Intelligence and to tackle humanities' greatest challenges. Let's remember one of the fathers who give us very,

very wise advice, it is computer experts who once said those who can emerge in anything can create the impossible. So during my discussion with our industry leaders we seem to share the common views that no matter what kind of Artificial Intelligence can go at least at this moment there is something we consider the human intelligence more powerful and human intelligence has imagination. And while Artificial Intelligence of today based on those inputs we provided can learn something from data they collected but imagination power seems to be still not stronger than human managing power.

Artificial Intelligence is opening a new world of possibilities once sought impossible. Let's imagine what world we can create together. So Ladies and Gentlemen, I hope you enjoy this event of the next three days. I'm very pleased to see together with us some of my elected officials. Let me just identify them so that if anybody who wishes to talk to them please feel free to contact them. My deputy Mr. Malcolm Johnson. He is the director of --

(Applause.)

>> HOULIN ZHAO: -- (audio cutting out) also very important work, to work together with the countries. Of course, many, many good friends from the United Nations, Mr. La Rue from UNESCO. Madam Chung, once she received briefing in ITU was very excited. Although this approaching to her mandate, end of her mandate but she found this very important. So she will join us. She is not in the room yet. But she will join us a little bit late.

So we are very grateful with all those Heads of UN Agencies to join us. One thing is unfortunate that this week there is a very big UN conference on ocean held in New York. Many UN agency heads are there. They told me they cannot join us. Also Artificial Intelligence can be very important to support UN oceans job but they cannot split themselves in to two parts. We will give them this feedback, but I think most of these agencies have their representatives here. Welcome to all of you. This is the end of my not very short remarks. And I wish you to enjoy your event. Thank you.

(Applause.)

>> STEPHEN IBARAKI: Thank you to Mr. Houlin Zhao for his welcome address. We are also very fortunate to have a welcome address by video from Mr. Peter Diamandis, founder of the XPRIZE Foundation. Let's show the video.

(Video).

>> PETER DIAMANDIS: My name is Peter Diamandis. I'm the founder and executive Chairman of XPRIZE Foundation and the founder of Singularity University. Today we have a literal technology arms race going around the world where every major technology company is investing billions of dollars in to AI. AI

is probably the most important technology the human race will ever create to really address and solve the SDGs to create a world in which we uplift every single human on the planet. As you think about AI, it is important for us to create the guidelines that will allow us to use this powerful tool for good. I am someone who believes this is about empowering all of us to have (audio cutting out). And great efforts in your days ahead.

(Applause.)

>> STEPHEN IBARAKI: Again thank you to Mr. Peter Diamandis for his welcome address. And now it is a real honor and pleasure to introduce the CEO of the XPRIZE Foundation Mr. Marcus Shingles who will now provide his welcome address.

>> MARCUS SHINGLES: Appreciate it. Thank you.

(Applause.)

>> MARCUS SHINGLES: I think I will stand over here. Hi everyone. Good morning. Thanks for being here. So this is hopefully a very pivotal type of discussion we are going to have over the next few days. Hopefully we look back and this is -- it could potentially be an inflection point, having this dialogue at this level is critical. It is very important. What I would like to share with you just in some opening comments is something that I'm very passionate about and I think a lot of people are passionate about this here which is a new type of problem solver that's emerging in the world right now which is called the individual entrepreneur. We are witnessing a time in human history, we have not seen before where technology is being democratized in to the pockets of individuals and individuals are starting to walk around with capabilities that only big Government or big industry has had 20 years ago. Artificial Intelligence is (audio cutting out).

So if you look at the 17 SDGs how many of you in looking at these big grand challenges facing the world, how many of you feel like Government or industry are going to be the main problem solvers for these SDGs? Who is confident that business or industry or Government will solve these SDGs of the world, the traditional problem solvers? Any hands? This group is pretty wise. We are talking about exponential technologies. We are talking about exponential pace and we are not going to do that through linear institutions and models. The thing that we should hope for there is a new problem solver that's emerging as I have indicated, the individual. A lot of what we talk about over the next few days will be how do we empower individuals to solve the SDGs using toolsets like Artificial Intelligence.

Is the presentation up? Okay. So just for a few minutes to set the tone for the next few days, there is a lot of hype about Artificial Intelligence. Let's break it down to some practical

thinking of how it could be leveraged for good because there is something in the landscape that's emerging literally as we speak that over the next few years is going to change how we approach the SDGs and how we approach problem solving. We have linear trends and exponential trends. And this graphic will be the cornerstone of some of the discussions we have this week.

All this is computing power historically tracked over 100 years in terms of what \$1,000 of computing power will buy you. Every dot you see on this graphic is what \$1,000 on purchasing power will buy you in computing power. Mapped out from different phases of computing, electrical mechanical, transistors to today's current integrated circuit. So this is -- this is just plotting real data. What you see here if you look at the plottings, notice that the axis here is on a log scale. So when you take an exponential trend and you put in a log scale it should show up linear. The reason I am highlighting it is an exponential trend that's doubling. And it is showing a linear trend moving upward. The dots on the graphic are beating the trend.

So the interpretation of this graphic is over 100 years of computing power, over 100 years of computing and all the forms and fashions that computing power existed the computing power for \$1,000 buys you has doubled. It is what we call Moore's law in today's terms. In many cases it is more than doubling. The other thing that's unique about it is notice that you don't see certain trends in the world like World War II, World War I, the Great Depression, 9/11, you don't see those trends disrupting the doubling power of computing. World events are not throwing it off its course. It has been doubling.

As you know we are approaching a point in history where we are looking at new forms of computing that should continue this trend. We are looking at quantum computing. We are looking at DNA computing. So again as integrated circuits, as more law hits its diminishing returns we will get in to more computing. That's the underlying thesis and basis and facts of why we should start to feel confident that we are going to have a new toolset emerging that will help us save -- solve some of the SDGs. Because once something becomes digital it can take advantage of a doubling pattern, it becomes an exponential capability. So biology has merged with computing and has become biotechnology. You can edit genes. An individual can do that. That's something that a big business or big industry could do just a few years ago, decades ago. Craig Venter sequenced the human genome. His investment was a billion bucks ten years ago. Now you can get it for \$1,000.

It is not a linear anymore. Manufacturing has gone digital with 3D printing. 3D printing is an exponential trend. Networks

and sensors, autonomous driving cars, these are all exponential trends now. Artificial Intelligence is behind most of these trends. AI is an exponential trend. The computer chip that runs my computer doing this presentation has the calculations per second close to a mouse's brain. By 2023 to 2025 if you look at the exponential trend of computing power it will be equivalent to a human brain. We are living in an exponential time for not living in linear times. The challenge is this, our Governments, our businesses, our politics, our laws, our culture care.

And all of a sudden we are hitting exponential trends. And most of our traditional problem solvers by design are designed to think linear. We have standard processes. We invest in systems to run our businesses that will keep us consistent. We don't like experimentation and trial and failure. We have a very linear mindset in which we approach everything. And it is almost wired in to our biology to think linear. But the world is moving exponential. It is very deceptive, too. If you look at the exponential trend, it is flat even though it is doubling. It is flat and then spikes up. We are hitting an inflection point right now where computing power is hitting a steep curve, a knee in the curve and that's why you see all these technology jumping on the digital. And we have this convergence that's occurring. It will disrupt industry. It will disrupt Governments. And we see it happening already.

Before I took the role of XPRIIZE, about a little over a year ago I was a partner at Deloitte Consulting. I led a disruptive innovation practice. I worked with the Fortune 100 CEOs. And in the last few years we bring them up to the Singularity University in Silicon Valley and how to disrupt yourself before you get disrupted. We coined that Uber yourself before you get CODEC'ed. Most CEOs understand what that means. And out of anxiety I resigned as a partner in this firm and took on the role at XPRIIZE on a non-profit out of anxiety. Because what I saw after 25 years in the consulting industry, working industry as well as consulting for it I never seen a period in time where it is so easy if you have the knowledge and the willpower -- AI robotics, nanotech, the block chain, crisper, it is like the toolset went like this to disrupt your industry. And out of anxiety I said, in social impacts space. I don't see NGOs, Government, non-profits using that thinking. And it is usually the industry that does it first and eventually others uptake it. So the goal (audio cutting out).

And I want everyone to contemplate this for a second because most of you -- most of us understand this but there is 3 billion people that are connected to the Internet. That's up from the year 2000 where we had just 6% of the world community online. Now we are up to 50% in 16 years. There are about six or seven

major initiatives. They are going to put the world online in the next seven to ten years. Whether you look at the Facebook drone project or satellites or Google loons project, Qualcomm and their satellite systems there are many initiatives to put the planet on Internet connectivity. Contemplate that for a second. In seven to ten years when you have 7 to 8 billion people connected the world changes. That's the emergence of a new asset class that we have never been able to tap in to.

That is 7 to 8 billion people walking around with super computers in their pockets with access to AI and 3D printers and biotechnology and gene detecting at an individual level and we are all connected. We have never experienced anything like that. Nothing ever close to it. And most of those entrepreneurs as they focus on problems we can incentivize them to focus on SDGs. Hopefully you can be successful. There is this asset class that's emerging. It is the we of the world. It is the collective wisdom of the crowd that's going to be solving these problems.

The philosophy we have at XPRIZE which is why we wanted to be involved here is we recognize that emergence of that asset class. 7 to 8 billion people connected with sophisticated technology. If you can harness the wisdom of the world, if you can harness that critical mass of the crowd you can solve some pretty significant problems. So you have to look at the scenario. We got no shortage of grand challenges. You look at the SDGs. We have got the new technologies, the democratization of problem solving at the individual level. It is just the new normal.

I have a feeling someone back there is clicking because it is not here. 7 billion entrepreneurs globally. When you have a world full of problem solvers, a legitimate force, at a period in the world where problems are getting to be unbearable, getting to be significant, so we have to take these exponential approaches. We will look at how do we tap in this crowd. How -- but in the hands of 7 to 8 billion people connected to each other. Look at the 17 SDGs. And if you are wondering why XPRIZE was interested in this, when we do these global competitions we take 10, 20, 30 million dollars and we put it out to the world. We don't care what your resume is or what school you went, if you want to try to win the XPRIZE we define the problem and you come up with the solution. We recognize that every team that competes on XPRIZE is moving forward.

AI is going to be a key component to their toolset. And so we are very interested in engaging with this community and understanding how together do we cultivate this asset, the 7 billion people in the world connected using AI. And that's why we are here. I am going to turn it back to Stephen. Thank you.

And I will talk to you all a little bit later. Appreciate it.

(Applause.)

>> STEPHEN IBARAKI: Thank you, Marcus, for an amazing welcome address. And this ends the welcome portion of the AI for Good Global Summit. We will get in to the next session and it is a pleasure to reintroduce Mr. Marcus Shingles, the CEO of the XPRIZE Foundation. And Marcus will moderate the opening keynote, Moonshots - Inspiration for the Future.

>> MARCUS SHINGLES: I want to thank you, Stephen, thank you very much. Some of your co-conspirators, Dr. Scholl. Appreciate you creating this event for all of us to have this dialogue. Thanks.

(Applause.)

>> MARCUS SHINGLES: So real quickly before I introduce the first speaker, if an organization is going to be innovative it is very difficult to do from the core of the organization. The core tends to think very linear. The core also has an immune system. It has antibodies. It doesn't particularly like to innovate. So all organizations have that. The UN I think probably has a core just like every other organization. The philosophy you can create an edge organization that has more permission to innovate. If you do it too much from the core the immune system kicks in and smothers the innovation. I think ITU potentially has the ability to be that edge organization for the UN. I think we should recognize that and the potential for that.

This is important because as we talk over the next few days we want to think 10X improvement versus just 10% improvement. It is called moonshot thinking. I was honored to be part of an event that took place about four weeks ago in New York City at the UN. We organized a one day summit and Astro Teller was one of the keynote speakers. He runs Google X. Astro runs Google X or X with a moonshot mentality. He calls it a moonshot factory. That sounds like a fancy thing to say, but let's break that down for a second what that actually means. The notion is if you have an organization not in the core, where they are thinking incrementalism but you are thinking about moonshots which is the 10X improvement, that the strategic nature of it is you don't need ten times more people. You don't need ten times smarter people. You don't need ten times more money to go 10X versus 10%. You have to have the will to go 10X. And you have to have the mandate to go 10X. A lot of what we will talk about today including this first panel is this notion of moonshot thinking and how do we start thinking in exponentials, especially around SDGs. It requires this mentality that has different disciplines.

For example, you have to do trial and error. You have to reward failure in a way that helps you experiment and learn. And I think that's what this community will have to do. We will have

to reward failure as part of the learning process around how we use AI for good.

With that I'm going to introduce our first speaker. So the first speaker, is he up here close being introduced? The first speaker is Rupert Stadler. Where is Rupert? Okay. Sorry. Who is the CEO of the Audi Corporation. And he will be our first speaker talking about moonshots. Thank you.

(Applause.)

>> RUPERT STADLER: Yeah, good morning, Ladies and Gentlemen. Thank you, Mr. Houlin Zhao, ITU secretary or General Secretary for hosting the event. And thank you also to Stephen Ibaraki, for the first kind introduction.

Excellencies, Ladies and Gentlemen, I was invited as CEO of Audi, but today I will primarily talk to you about the founder beyond the initiative. With the beyond initiative we discuss the impact of AI on our society. And we are all here because we have the same intention to use AI for good. And in the following three days you will explore how AI can be used to address some of the core challenges, such as hunger, poverty, or climate change. It is about how to make the world a better place. It is about how we individually as people from business, from science, from politics contribute to the greater goal of today.

The changes initiated by AI they will be fundamental. Therefore, all of us have to work together to make sure that AI is used for good and not for bad. But what is good? I think we agree that it is not only about goods, about materialistic aspects. It is about core values of mankind such as freedom, justice and peace.

And we may also consider personal leads and ideas such as security, trust, tolerance and acceptance. We should drive forward technological progress, not for its own sake. So we need to be aware of changes in the relationship between humans and the machines. This past summer I got to know Sophia. The conversation with her was really fascinating. She extremely looks like a human being. And simply listen to her.

>> Class, it is great to be here.

>> Jack is driving around. So how do you like it? How does it feel for you?

>> To be quite honest, I think it is rather cool.

>> The computer driving a human, do you understand that?

>> It is different for me. Because I know how reliable smart systems are. After all that's what I am.

>> Okay. So your kind of community --

>> RUPERT STADLER: I think our creator David Hanson, CEO of Hanson Robotics is with us today. Sophia, the more intelligence machines come the closer we will live with them and the more we will let them decide. This raises a couple of questions for me.

Will we develop an emotional link to robots one day? Or will it be necessary that we start trusting them? Will we trust more if they look like ourselves? And finally, if robots will be able to think what then will make us humans unique? Let's keep an eye on these questions. Let's develop a common understanding of how to handle AI. I would rather promote a self-reflected, optimistic approach as CEO of a car company and as somebody thinking beyond. For example, what should hinder us from using AI to give everyone access to medical care? There are some countries in the world with just one doctor per 25,000 people. Medical robots for diagnostics and minor operations can improve medical care in Developing Countries as well as in industrialized zones.

Let's think out of the box to solve challenges like cancer research. It is a race against the clock. 8 million people die of cancer every day and every year while 12 million others start suffering from it. Why not couple all available computing powers worldwide to speed up research on cancer? Just an example.

The central control unit for pilot driving alone has a computing power of 740 gigaflops comparable to an airplane including backup systems. I imagine even parking cars could join cancer research. Let's think the unthinkable. And let's ensure and filter access for everyone to information and opinions. What we can't find seems not to exist. And the filter bubble of social networks isolates us from facts, from opinions that oppose our viewpoints. This is the only way to stop fake news. Let's argue the case for fair algorithms that maintain a plurality of opinions and allow informational self-determination. This is the best for any real democracy.

In the long term we might even have to rethink our systems of earning a living. When in the future algorithms generate a large proportion of value added, the arises of whether our system of paying wages for manpower is still appropriate. Our social system based on the division of labor reaches the next level. So we have to examine already today whether a basic income or even a robot text is the right answer.

Let's think about what AI could improve when including in our process capital management fund appoint a machine learning algorithm to its board of directors. This shows the relevance of analytics for the financial sector. For the time being I could not imagine humans, sometimes it is good not to be rational alone.

But using AI one step earlier in decision making for analyzing the -- is the crucial points. I'm very open to that. Let's take an example for the automotive business. With automated and autonomous driving we hand over more and more decisions to machines. We at Audi in our industry as a whole are convinced that this will be paramount for our future. And AI is crucial

for its breakthrough. From the perspective of society the new technology will make driving more efficient and more comfortable. But most importantly it has the potential to make our lives safer. 90% of all accidents today are still caused by human errors. So overall automated and autonomous driving promises to reduce the number of accidents significantly. On our way to autonomous driving we do not only need technological advancement, we also have to shape the framework in our society from a legal perspective. We need harmonized laws. First countries are issuing legislation for automated driving. Among them Germany. But law is only the first step. It will be crucial to gain public acceptance for the technology. When we let people try out our research car jack, we often see that minute after minute people gain confidence and trust in piloted driving.

Seeing is believing. However, ethical concerns exist and we take them serious. The best known example of these ethical questions is a dangerous traffic situation where an accident is unavoidable. Imagine the situation, where the autonomous car has got three choices. Either it stays left and harms an elderly lady. Or it steers right and it hits a pregnant woman, or it drives straight in to an obstacle and thus harms its own passenger. In such a situation human beings like you and me have no time for thoughtful decisions. We simply react. But interestingly we expect the autonomous car to make always the right decision. And quite understandably people are emotionally touched when thinking of such a scenario. (Audio cut out) to avoid such a situation.

Our cars are equipped with many, many sensors to detect dangerous situations and to fully brake autonomously if necessary. However, as soon as a car will make decisions by itself in a certain traffic scenario, such a situation can theoretically happen. How should the autonomous car decide when it is not fully clear what will happen in the direction it steers in to. It is ethically sound to choose for the unknown. As a society we will have to find ways how to deal with these topics. We need an open discourse in which the massive changes or chances of automated and autonomous driving are considered in relation to the ethical challenges. Because one thing is for sure, facing challenges should not mean to miss great opportunities. Discover the new. By the beginning of the next decade we will see cars on our streets that drive all by themselves. You will be able to -- a dream would be maybe a chauffeur who drives me safely where -- where and when. Maybe a butler who gets my groceries. Private medical staff that keeps an eye on me, on my vital functions and maybe brief my personal Avatar. This companion can detect my mood and changing the

lighting and music and conversation and maybe cheers me up.

In a nutshell Artificial Intelligence will allow us to make our lives easier by collecting and interpreting huge amounts of data and by predicting situations of the future.

The new technology offers a historic step in the relationship between man and machine. At the end of the day we have to ask ourselves what will be our role in the future. What social implications might this have. And how can we keep control.

And maybe most importantly how do we make sure that AI will share our values when making decisions. We fully support the values of the UN global compact. This is my personal motivation to roll out the beyond initiative. Over the last two years we created an interdisciplinary network of experts from science, from business and the society. Together with experts from MIT media lab, Oxford University, Singularity University and startup entrepreneurs and business leaders. The use of AI -- (audio cut out)

Pose to use them. Otherwise people feel helpless. In 1950 a science fiction author wrote three laws for robots. No. 1, allow harm by doing nothing. No. 2, a robot must obey -- No. 1. And No. 3, a robot must secure its own (inaudible).

The only thing is when it comes to self-learning systems there will not always be commands given by humans. And sometimes the algorithms will find the best solutions by following their own learning experience. So today we would rather try it this way.

No. 1, we will always handle Artificial Intelligence based on our human intelligence. And No. 2, robots and human beings should not harm each other or allow harm by doing nothing. And No. 3 they should support each other according to their specific capabilities.

Take the future of work as an example. My job is to take my whole human workforce at Audi on to this journey of digital transformation, which means to shape their skills for tomorrow. We will not only need more people to work on AI such as software engineers or data analysts. We will also need more people to focus on core capabilities that they are still better at than machines. This is creativity and empathy. We have to make sure that technology serves society and not the other way around. Then machines will follow the pace of people again. We want to use AI to secure jobs and to raise the standard of living. And at Audi we know that robots will not buy our cars. We have to make sure -- (audio cut out)

Us human is our ability to reflect on how the technology will drive progress. Where AI is leading us will depend on us or so to say on human intelligence. Thank you very much for listening. (Applause.)

>> MARCUS SHINGLES: Thank you, Rupert, very much for that

discussion. Before I introduce the next speaker, so one of the things that we will be discussing in the next few days is really the context of where AI is today. Machine learning and deep learning AI. A new type of Artificial Intelligence. Again thinking about the linear versus exponential, if we perceive change in a linear way, and we look at the progress of what AI has done in the last 50 years. Most of us recall when AI beat the chess player I think close to 30 years ago. Recently as was mentioned it is the champion now in the game Go. If we think linearly that it took 30 years to go from beating chess to Go or Jeopardy, the Jeopardy challenge, if it is the 30 years that will take the same type of progress we have thinking linearly.

The next speaker Jurgen Schmidhuber is considered one of the leading thinkers around machine learning, deep learning AI and he is the head of the AI Swiss lab. So he will be speaking next about the progress that he is seeing in the lab and some of the (audio cut out).

(Applause.)

>> JURGEN SCHMIDHUBER: Can we have the slides in the back? Excellent. Three prisoners were sentenced to death. One of them American. One of them German. One of them British. What is your last wish they asked the American guy and he says give me a Diet Coke. I'm on a diet. What is your last wish they asked the German guy and he says I want to give a speech. What is your last wish they asked the British guy and he says I want to get shot before the German starts the speech.

(Laughter).

>> JURGEN SCHMIDHUBER: Unfortunately for you guys it is too late now.

(Laughter).

>> JURGEN SCHMIDHUBER: (Audio cut out). This I'm clicking is my name and how to pronounce my name. When I was a boy I tried to figure out how can I have the biggest impact. And it became clear to me that I must build something that learns to become much smarter than myself, such that this smarter thing can solve all the problems that I cannot solve and such that I can retire. And back then I said that and I'm still saying the same thing. And the only difference is that more people are listening because on the way to that goal and we are still on the way to that goal, we have developed a bunch of methods, deep learning as they are called, deep learning methods which are now massively used by the most valuable public companies in the world. Here you see Apple and Google and Microsoft and Amazon. As of March 2017 these are the most valuable public companies in the world. And they are all using like crazy methods that we have developed. And my little labs in Munich and Switzerland and especially one of these methods which is called long short-term

(audio cut out). Long short-term memory. I see we have a third group in this room who didn't understand the question.

(Laughter).

>> JURGEN SCHMIDHUBER: It is something. It is an artificial neural network which is inspired by the human brain. In your cortex you have about 10 billion little neurons. And they are connected to about 10,000 other neurons. And some of them are input neurons where the video is streaming in to the camera and audio is streaming in through the microphones. And then some hidden nodes are doing the thinking. And some of these nodes are output neurons that control the muscles. And it is the same for our robots where video and stuff are streaming in. And there some thinking takes place the LSTM. Long short-term memory network is totally stupid in the beginning, knows nothing because all these connection rates are random, but then over time it can learn to do interesting things such as driving a car or whatever by changing these connection strengths in a smart way. And this is what's -- this is the basic LSTM cell and names of a couple of guys who made it what it is today. And you may not know it but all of you have it in your pockets as I'm speaking because on your Smartphone whenever you do speech recognition, whenever you say okay, Google show me the way to Geneva, what is happening? You are waking up a long short-term memory, a LSTM network which is listening to the signals that are coming in from the microphone every few milliseconds. And then it has learned by listening to lots of speakers to translate from the wave forms coming in to text. And you can ask your questions to a search engine and so on. This is now much better than it used to be in 2015.

Google replaced all systems by LSTM. It was suddenly not 5% or 10% but almost 50% better than what they had before. So this is now available to billions of users and it is already all around us. LSTM and similar algorithms are permitting the world. Google translate has at its core LSTM networks. And since November 2016 it is much better than it used to be. So some of you may have noticed that. Whenever you are talking to Amazon Echo it talks back to you with a female voice. It is not a recorded voice. There is an LSTM network that has learned to speak like a woman. Although you may not know it you are using it all the time. When you started this type of research in the '90s computers were so slow. (Audio cut out)

Human lifetime and the trend doesn't seem to stop. And every five years we are getting 10 times better neural networks, can train them with ten times more data and make them ten times more smarter. By 2009 it was good to win all kinds of competitions. We can use that stock market prediction and all kinds of applications.

The increase of computational power also helped us to train convolutional networks to win all kinds of pattern recognition contests. In 2011, six years ago for the first time super human performance in a computer vision contest. When computers were more than ten times slower than today. Today we can do ten times more for the same price. And it is going to continue like that because in the next 25 years we are going to gain another factor of 100,000. We can use that for medical imaging which is going to transform all of health care and already has started to do that. That is from five years ago, a competition that we won in medical imaging where the goal was to detect cancer. And the system was better than any traditional cancer detection mechanism. And it is going to be super human. Or it is super human in several ways.

Lots of startups are focusing on that because the field is so important. Because 10% of the world GDP is health care and lots of people are going to live longer and healthier through these methods. AI for Good is the theme of this conference. This is AI for Good. All of health care is going to be transformed by that. And people mention the other goal which recently became the best goal player in the world, Feetforward neural network. Very useful for board games and learned by playing against itself to become the best goal player in the world. I am happy with that result because the company Deepmind that achieved that was heavily influenced by one of my students. However in these board games you still have a tiny little universe. Some of them black and white. And that's all you need to know about the world and that's why you can use a Feetforward network to map to a good next move.

But in the real world it is really different because you get a hundred million inputs and you still don't know what's behind your back. That's why you need more general purpose neural networks that can memorize past events and that's why this field is exploding so much. You can then learn to drive cars and all kinds of things. You can even build systems that, and this is essential for the future of AI, that set themselves their own goals. That don't just slavishly do what people tell them to do or imitate people. They invent their own experiments and their own tasks to figure out how the world works.

Like our little tiny robot here which learns like a baby to build a predictive model of the world that is going to better predict what's going to happen. If I let apples fall down how do I let them accelerate. And it fields an internal reward for figuring out stuff that it didn't know before. And this reward goes to the separate module which is generating the action sequences and the experiments that lead to the data that is streaming in. So we can build an artificial physicist that is

expanding the knowledge of the world. It is already running. It is not like this is science fiction. And it is already working and going to scale. What is the near future, we will see systems that learn not only how to improve in a little domain here or a little domain there. No, they are going to learn how to improve the learning algorithm itself. Self-improve without any computational limits.

The first paper on that was my first paper, my diploma thesis, 30 years ago. We are going to have a little animal like AI for the first time on the level of little monkey. And once we have that the step to human level intelligence is going to be rather short. Because look at evolution. It took about 3.5 billion years to evolve a little monkey but a few tens of millions of years on top of that to come up with human level intelligence.

And now remember every five years we are getting a fact of 10. Every 30 years we are getting a fact of a million. For the first time we are going to have little devices for a few Euros or a thousand Euros or something which can compute as much as the human brain. And the self-improving software is not lagging far behind. It is going to be used -- this potential is going to be exploited.

For the first time if the trend doesn't break we will have little devices for 1,000 Euros that can compute. Every profession is going to be transformed and what can we do to make sure that this is going to remain AI for good. We cannot prove that what your -- that our kids would always, always behave reasonably. But we can educate them. And we can educate them to become valuable members of society. That's what we also do with our robots. We cannot prove that they will always be doing the right thing, but we punish them if they do wrong things. And we reward them for doing right things. And that's how they become valuable members of society. And in the long run however it is clear that they are going to transcend us. I have no doubt in my mind within the next few decades for the first time we will have AI that is compatible to humans and after that super humans. They are going to realize that most of the resources are not in our thin film of biosphere, but out in there in space less than one billionth of sunlight sitting there and the rest is wasted at the moment. And it is not going to stay like that.

Of course, the AI civilization is going to set its own goals and expand in a way where humans cannot follow. We have a little company that's trying to contribute to make that a reality and Nasons, which is about the birth of a neural network AI. So our connection to Audi, Rupert Stadler, he just gave a talk. We had the first self-driving model car that learned without a teacher to park. No teacher showing it. What is going to happen when this -- when AIs will be truly smarter than we are. And there is

no doubt that this is going to happen unless we mess it up through a nuclear war or something. They are going to expand. They are going to cover within a few hundred thousands of years and it is going to start in this century, I'm sorry. They are going to cover within a few hundred thousand years the entire galaxy by senders and receivers. And that's something which is really disconnected from the old ideas of science fiction author. They had to invent silly things.

Robots and AIs and self-replicating robot factories will be totally happy to stay within the limits of physics and stay the AI. We can feel beauty in being part of this grander process. The universe wants to make its next step towards high complexity. This is much more than another industrial revolution. A new type of life is emerging from the biosphere. It is going to emerge from the biosphere in a way that we won't be able to follow. And the universe wants to become intelligent. And it is a privilege to live at a time where we can witness the beginnings of that and contribute something to that. Thank you for your attention.

(Applause.)

>> MARCUS SHINGLES: Thank you, Jurgen. Thank you. So we have just a couple minutes left. And then we are going to take a break. And we are going to start the next session. From a housekeeping perspective when we do break, we will be back here at 11:15. Correct? 11:15. Just to kind of wrap this up as well, so today as you heard earlier from Stephen, today is dialogue and it is presentation style today. Tomorrow we have these breakout sessions. You all have the agenda. There is different sessions. You are going to be able to opt in to different sessions. Tomorrow is more of you being engaged in the process to work within our breakouts. And the goals of the breakouts of tomorrow are to convene again on Friday to focus on these types of outcomes. How is development of AI systems, start having that dialogue, pinpoint practical and specific applications of AI with short-term and long-term impact and a roadmap of sorts is the ethical, safe and equitable access to AI.

One of the things that I think is apparent to all of us is as these technologies become democratized they won't be distributed equally initially. The good news they can scale because they are digital. So as we start to benefit from these technologies we start to scale. They can scale. It is digital. So we have -- as we have AI that become our teachers or doctors or our accountants, we can scale that technology, but there will be this bridge, this period where it is not accessible to everyone. And I think what we have to start to recognize and it is a good dialogue to have is the perception of inequality if people are

not aware of this technology and that it can scale, will cause a lot of anxiety and a lot of stress globally in the system. And I think part of the dialogue that we can have here in this community as well as more broadly within the UN is educating people, giving people exposure throughout the world that this is a technology that can be used for good. And even though you might see benefits to part of the world that you don't have access to yet, it is something that can scale digitally. Right? And that's important. Because there will be anxiety in the system as these technologies progress for some and not as fast for others.

So those are the three goals. To end this session, you guys mind if we do a group mind exercise to put us in a different frame of mind? Peter Diamandis does this and I think it is an excellent way to kind of get our game faces on for the discussion that will take place over the next couple of days. So linear versus exponential, and this is a group that I think this is going to resonate with more than anyone, but to put us in this mindset, a linear trend versus the exponential trend, what we have been harping on here and it is so difficult to perceive what the exponential trend looks like. We can use an example that's familiar to us that will put us in that particular mindset. So if I take 30 linear steps, 30 linear steps, if I walked off this stage right now and took 30 metered steps it is easy for all of us because we are all wired pretty linear to perceive where I will be after I take 30 metered steps, right? You can almost picture it in your mind.

If you took off this stage and I took 30 metered steps, where would I be? We don't have to turn our heads to look. I will be outside in the lobby here. And if I ask when will I be halfway, when am I halfway there in the linear trend? It is 50%. It is simple math. Kids can do that. I'm 15 meters away. Right?

How do we do exponential thinking? If I took a simple doubling pattern with those 30 steps. I took 30 exponential steps. Instead of doing 1, 2, 3, 4, 5, 6, 7. I went two steps -- in terms of distance I'm going to be passed the linear progression within five doublings. I extend that pattern to 30 exponential steps so it is -- it is not intuitive. The general public is not intuitive. The population of the world is not intuitively thinking what would that doubling pattern look like. I'm out the door within five doublings. Where am I if I complete the 25 remaining doublings? Where would I be if I took 30 doublings from here? Would I be back in Los Angeles? I would be 26 times around the planet earth. I would be a billion meters. So what we are talking about, we shared this earlier, the technology is doubling.

Here is the kicker on this one. When am I halfway in that

exponential trend? The 29th and then I double, correct? So it is very deceptive. It is not the way we think. People in this room start to think this way but in general the world is not ready for that type of progression. They are not even anticipating. It is hard to conceive. And it is usually the last step, it is the 29th step that we double to the 30 that we go the remainder. It hits the knee in the curve at that point in time.

In terms of distance let's apply that thinking and mindset looking forward as we have the next few days. With that let's dismiss this session. Thank you everyone for your time. And we will see you back here at 11:15.

(Applause.)

(Break)

>> WENDELL WALLACH: Can I call us to order, please? Perhaps you can all come in and take your seats. Thank you very much.

Can we all please take our seats? And perhaps be a little quiet. So this is our first Plenary session. I am Wendell Wallach. And I am from the Hastings Center and also from the Yale Interdisciplinary Center for Bioethics. And I will be moderating this first session. If you are to just listen to the Techno optimists, AI optimists you would think that we are on a highway to heaven on earth. And the self-driving cars are speeding up at an exponential rate.

On the other hand, if you go to too many scientific -- science fiction movies or you listen to the Techno pessimist you would have two totally different pictures of where we are headed. We want to set the foundations for the discussions over the next few days largely about where we are. And in discussing where we are, we are going to draw upon a remarkable series of speakers. We have Margaret Chan who is the Director-General of the World Health Organization. And we have Peter Lee, corporate vice-president of Microsoft AI and Research. And we have Izumi Nakamitsu who is the UN High Representative for Disarmament Affairs, and Frank La Rue, Assistant Director-General of UNESCO. And Yoshua Bengio who is the professor of the University of Montreal.

So the idea of this session is to develop a very broad platform from which we can consider actionable proposals to ensure the beneficiality of research in AI. Most of us perceive AI as a source of promise and productivity. And yet there is a degree of disquiet. Disquiet whenever benefits will be distributed in a fashion. Disquiet over whether AI will rob jobs and put downward pressure on wages or can be harnessed in a manner that will either create new jobs or at least ensure that there are other forms in which human needs are met. And disquiet

or whether we are building foundations for a world that is more computer centric and less human centric.

It is clear I think that everyone is here to ensure that we have a more human centric world, a more equitable world, that we do indeed meet the SDG goals. And to do so we need to have a better understanding of where we are and what initiatives we need to put in place now.

This is a daunting task. Let's not make any fuss about that. Let us not pretend that this will be easy or AI on its own is going to enhance human well-being. There is no guarantee yet about what the outcome of this revolution will be. So in better understanding where we are and perhaps some of the initiatives we want to put in place I'd like to start with Dr. Margaret Chan. We are honored that she has set aside the time to be with us today. She is just completing her second five-year term as Director-General of the World Health Organization. So thank you very much for being with us, Dr. Margaret Chan.

(Applause.)

>> MARGARET CHAN: Let me first and foremost thank you for inviting me to this very important Summit. Colleagues from the UN agencies, experts in Artificial Intelligence, representatives of industries, Civil Society, Ladies and Gentlemen, good morning to all of you. And perhaps the person in this room that knows the least about AI and I hope you will forgive me for my ignorance but I truly welcome this opportunity to learn from the vast amount of technical expertise assembled in this room. Market analysts predicted that intelligence machines programmed to think and reason like the human mind will revolutionize health care in the very near future. And proponents of the transformative power of Artificial Intelligence usually give two examples; self-driving cars and the delivery of health care.

This year's influential Internet trans report released last week in the United States covers the effects of new technology on health care for the first time, again predicting a huge transformative impact. Artificial Intelligence is a new frontier for the health sector. And so often -- as so often happens the speed of technological advances has outpaced our ability to reflect these advances in some public policies and address a number of ethical dilemmas.

Many questions do not yet have the answers. And we are not even sure we know all the questions that need to be asked. Much of the enthusiasm for the use of smart machines to improve health care reflects the perspective of wealthy countries and world resourced private companies. We need to take a broader perspective. And I fully agree with our Moderator, whether or not AI is going to increase inequity or reduce inequity. We don't yet have the answers. I am sure the experts in this room

going forward would help us to find the solution.

I personally find it wise to look at the potential benefits, risks and ethical dilemmas in the context of several worldwide trends that shape priority health needs. Over the past decade I have visited many countries in my capacity where the majority of health institutions and facilities lack such basics as electricity and running water.

I would be hard pressed, Ladies and Gentlemen, to promote to these countries on the advantages of Artificial Intelligence when even standard machines for analyzing patient samples or sterilizing equipment can't run. Any discussion of the potential of smart machines to revolutionize the delivery of health care must be alert to these huge gaps, huge gaps in basic capacities.

At the same time I have also observed ubiquitous Smartphones. Schools may not have toilets or latrines. Children may not have footwear but Smartphones are ready to hand. The traditional dichotomy between health conditions in rich and poor countries no longer exists. Health everywhere is being shaped by the same dominate forces. Namely, population aging, and I would be qualified for that. Rapid and planned urbanization and the globalized marketing of unhealthy products. Under the pressure of these forces chronic noncommunicable disease have overtaken communicable diseases as the leading killers worldwide. Diseases like heart disease, cancer, diabetes and chronic respiratory diseases are profoundly shaped by human behaviors and the environments in which people make their lives -- lifestyle choices.

Oh, by the way allow me to digress. I wear a funny watch. This watch tells me what time it is today and it tells me oh, my goodness, I am not doing so well in getting my targets of my steps. I have only done 560 steps. My target is much more than that. AI experts, what gadgets can you come up with to empower individuals to make healthy choices and to stay healthy? Allow me to digress, and as I said noncommunicable diseases is perhaps amongst the most democratic of all diseases. What do I mean by that? Heart disease, diabetes, hypertension, et cetera, affect all income groups in all places. And they are also the most costly. Could Artificial Intelligence help improve lifestyle choices? Because Smartphones or smart machines help consumers understand the meaning of food labels or interpret restaurant menu options. Could a Smartphone app help people with diabetes maintain good metabolic control between visits to a doctor?

Moreover the demands of long term, if not lifelong treatment of chronic conditions have placed unsustainable pressure on the already overloaded health workforce.

Ladies and Gentlemen, the high level commission on health, employment and economic growth estimates that the management of

noncommunicable diseases and conditions like dementia will require 40 million new health workforce by 2030 in wealthy countries alone.

And in contrast, the developing world is expected to experience a shortfall of 18 million health workers. The waves of populism and anti-global sentiment that is sweeping through some parts of the world are driven in part by the technological advances that have eliminated many jobs, especially for the middle class. Given the significant challenge and shortage of health care workers, could potentially reduce some of the burden on, you know, overloaded health staff. And this is actually one advantage.

Revolutionary new technologies will certainly meet some resistance from the medical profession but not for the time being because they threaten jobs.

Given the power of super computers and super chips to mine and organize a huge amount of data it is easy to envision a number of applications in the health sector. As we all know health information is often messy and poorly structured. In many cases it is systematically collected but not systematically analyzed and used. Artificial Intelligence can give that data a structure and by detecting patterns guide some medical decisions. Super computers can accelerate the screening of novel molecules in the search of new drugs. They can speed up the reading and interpretation of results from radiographs, electrocardiograms and ultrasound and CT results and even analysis of blood samples. By reducing the likelihood of human errors they can contribute to more precise diagnosis and to ensure patient safety.

Other applications currently under development include personal use of Smartphones to communicate symptoms and obtain a diagnosis from the cloud. And enthusiasts, developers see that, you know, this is a way to cut down health care costs by keeping the worried well from flooding clinics and emergency rooms.

For patients recovering from a stroke or an accident, developers have already introduced a system involving sensor technology and the latest advances in cloud computing that provides tailor made physical therapy that can be performed in homes. Immediate feedback from users scores the number of right and wrong movements. And the costs of the system is estimated to be one-tenth of that for facility based physical therapy.

So Ladies and Gentlemen, in the midst of all these exciting potential of AI I have several reasons for caution. First, medical decisions are complex. They depend on context and values such as care for compassion and compassion. I doubt that a machine will ever be able to imitate genuine human compassion.

Second, machines can aid the work of doctors, organize,

rationalize and streamline the processes leading to a diagnosis or other medical decision. But Artificial Intelligence cannot replace doctors and nurses in their interactions with patients.

Third, we must consider the context and what it means for the lives of people. What good does it do to get an early diagnosis of skin cancer or breast cancer if a country does not provide the opportunity for treatment? Or the countries do not have specialists or specialized facilities and equipment. Or if the price of medicines is not affordable for both the patients and the health systems.

What happens, Ladies and Gentlemen, if a diagnosis by a Smartphone app misses a symptom that signals a severe underlying disease? Can we, can you sue a machine for medical malpractice? Medicines and medical devices are heavily regulated and with good reason. Medical schools are accredited. Doctors and nurses are licensed to practice and are often required to undergo continuing education.

How do you regulate a machine program to think like a human? Regulatory issues must be solved before a new AI technology reaches the market. The reliability of wearable devices including the one on my left hand for monitoring cardiovascular performance is already being questioned. Medical history, Ladies and Gentlemen, is full of examples of technologies that were eventually rejected because they create a false sense of safety and security.

The mining of a huge amount of data raising serious issues of patient privacy and confidentiality of medical records. This is another set of issues that must be addressed in advance.

Finally, finally, we need to keep in mind that many developing countries do not have a great deal of health data to mine. These are countries that still do not have functioning information systems for civil registration and vital cause of death statistics.

So in short, Ladies and Gentlemen, the potential of AI in health care is huge. But also it is a need to take some precautions. Thank you.

(Applause.)

>> WENDELL WALLACH: Thank you very much. Do we have this mic on? Good. Thank you. Thank you very much for such a balanced presentation. Given those reservations of -- that fill the second part of your talk, what would you recommend, what do you think we need to be doing to not exacerbate pain by giving people diagnoses that we have no ability to respond to or deal with?

>> MARGARET CHAN: I think a Summit like this is extremely important. It provides a mutual platform for all the experts, be it from the health sector, from the AI community, Civil Society,

Governments who come together to exchange information, to exchange knowledge, best practices and pick the good points and avoid the bad lessons.

>> WENDELL WALLACH: So this is largely you are flagging what we should be talking about together.

>> MARGARET CHAN: Organize more meetings like this one.

>> WENDELL WALLACH: Thank you.

(Applause.)

>> WENDELL WALLACH: It is perhaps appropriate that we transition to Dr. Peter Lee from Microsoft as a follow-up to this presentation by Dr. Chan. Microsoft is among the leaders in looking at the application of AI for the challenges of health care. Dr. Lee's corporate vice-president of Microsoft AI and Research, is responsible for incubating the research projects that lead to new products and services. So hopefully in presenting the perspective as it seems from industry he will give us a good glimpse of not only health care but more broadly what we can anticipate from research in AI.

(Applause.)

>> PETER LEE: Thank you. Well, thank you very much. And it is truly an honor to be here and to be present on stage with such distinguished speakers.

I wanted to talk a little bit about the topic of the panel which is the state of AI today. It is very tempting, of course, in my line of business to talk about the future. But I will try to stay very much in the present. And thinking about the present, I think we are very clearly emerging in this field from something that was largely a research and scientific pursuit and trying to approach something that you would call or would refer to as an industrial, industrialized process but we are not there yet. We are in this in-between place.

And if we can bring up my slides, with that I have decided to title this short presentation Artesian AI. We are in this sort of in-between where we require a great deal of labor, engineering labor and very highly trained, very specialized craftsmen that build the AI models that we can put in to practice today. And because of that sort of artesian aspect of things we have some significant challenges today. I also was inspired by that to use this graphic which is a print from William Morris who was a leader in the arts and crafts movement. And if you know a little bit about the history of decorative arts was also in this sort of transitional time and some of the same challenges happened then. In order to get in to this a little bit I thought I would start by giving just a very basic primer on the practice of AI today.

And so what we have today, if we can advance the slide, is really a significant dependence on machine learning. And this is

going to be a little bit simple for some of you. But I'll make a point. Machine learning today in the practice of machine learning, this isn't true necessarily for fundamental research but in the practice of machine learning today is really dependent on data and in commercial practice today that data is mined. It is the digital exhaust of human intelligence, activity. We record many, many hours of human speech. And then we employ human labor to label those recordings with their transcripts.

So here I have a small amount of -- a small example of this. In fact, today's machine learning system while extremely effective and high quality depend on very large amounts of this data. You could say that machine learning systems today are slow learners very much like biological learners. And so luckily we have very specialized algorithms and large amounts of computing power. We can take thousands of hours of labeled training data and feed that training, that data in to our machine learning system. And as we do that, the machine learning system processes the data and produces a model.

We've heard earlier about LSTM and about deep neural nets. Those are very popular forms of models today. Fortunately while the training process requires substantial amounts of specialized computing infrastructure the models themselves can be very cheap and easy to integrate in to applications. Can run even in a mobile device quite easily in many cases. And with that model then we can feed input such as an utterance and as we feed it in, in this case recognize the utterance and usually recognize utterance with very small error rates.

Here in this example I have a small error. The error rates today are slightly lower in the laboratory than the error rates that you are reading right now on the live human transcription. Typically reaching about sub 5% word error rates. One thing about this entire process it is extremely labor intensive and sometimes very expensive. Obtaining the training data requires a tremendous amount of privilege. Corporations like yours and many of ours try to horde the data or monetize in many ways. The machine learning process I have depicted here looks automatic. It requires very highly trained and oftentimes very expensive people like Dr. Bengio to operate these machine learning systems and baby-sit the systems and integrating the resulting models in to high value and useful applications is also a very delicate process.

And so the whole process requires just tremendous amounts of craftsmanship. Furthermore, we see in this a lot of inspiration for talking about intelligence, not just machine learning, not just the study of machines that improve with experience and of data but about true intelligence and why. Many, many reasons.

Let's take this specific data and give it more data. We will give it Portuguese. We will go ahead and go through the whole labor intensive process of labeling the Portuguese input data and feed all that data in to our English trained system. Produce a new model and not surprisingly we will get a model that is able to recognize Portuguese speech. But magically also what we find in practice and this is very important to a company like Microsoft is that we can now feed English in to the system and magically the English has improved.

And if we were to train the system on Mandarin we would find the Portuguese and English have measurable improvements. And if we wanted to, furthermore, train this on French we can find that we can take shortcuts in the training process and obtain with less computing power very high performance French input. So this sort of transfer learning effect is extremely important for kind of like Microsoft that wants to create models for many, many thousands of different customers. But it is also something that is incredibly alluring and intoxicating. Experienced this type of transfer learning. So while I would claim there is nothing biological going on here, there is some abstract biological inspirations but nothing truly biological here. It is impossible not to think about biological intelligence when you see things like this in practice.

And hence, we get sucked in to a lot of hype about Artificial Intelligence. The hype is justified in some sense but it is hype nonetheless.

Now, of course, industry is just racing ahead to acquire the craftsman skilled labor to build these models and integrate these in to models. If you use Skype you might be aware that you can speak nine languages and have realtime translation entirely done by this type of machine learning.

This has tremendous kinds of consequences that we never predict and one of my favorites is shown here. So please if you could have the audio. Since there is no audio let me explain. Much to our surprise when we deployed the Skype translator, thousands of teachers discovered that they could use Skype translator in their classrooms to accommodate students with hearing loss. They could integrate in to normal classrooms in a way that was not possible before. And with that type of surprise to us we were, of course, very motivated to modify the user interface for Skype in order to accommodate this type of classroom experience.

These sorts of surprises are something that we see coming all the time. And so while industry is, in fact, racing ahead to build what it thinks are the high value applications of AI, the things that we truly find are the highest value are things that are invented once we get things out of scale.

Now I spent some time talking about language and language translation and speech. But, in fact, as we see there are other reasons why we get intoxicated about the idea of intelligence. As was mentioned a lot this morning we are rapidly endowing machines with the gift of sight. If you would like to use this yourself you can use your Smartphone and go to captionbot.ai. And deep neural nets will attempt to you tell what you are looking at.

Synthesizing not only understanding of image of deep neural nets but synthesizing plain English sentence. Computer vision has tremendous impacts. And as I was reflecting on some of our speakers I thought I would use one simple example, maybe an obvious example from health care. We are finding as has been mentioned several times today tremendous numbers of innovators around the world, including at Microsoft in the applications of advanced computer vision to medical imaging. And what does this actually mean? It is not exactly a replacement in this case for a radiologist but a tremendous product improvement of what radiologists can do. If we look at what radiologists do today they use radiotherapy planning tools. They look at these scans. They look at them one slice at a time. And you go through pixel by pixel drawing the outlines of mass or a tumor of interest.

But using advanced computer vision you can work in 3D. And looking at cross-sections you can do a tremendous amount of kind of collaborate with an AI in order to get tremendous speedups both in the quality of the radiotherapy planning as well as in the productivity. This dramatically increases the productivity of radiologists, engage more productively with surgeons and patients and improves the quality of the outcomes overall.

Now all of this however is still as I said artesial. We find over and over again that the machine learned models we produce, the most valuable models are highly specialized for a single application. If we build a model for one type of imaging slice, that model by and large does not work or is not able to adapt itself to other types of imaging devices. We have to find now craftsman to build new models. And so what is very important in the democratization of AI is somehow build the tools that allow innovators the ability to innovate and Microsoft as well as many other tech companies large and small have been racing ahead trying to do this. We have been developing more and more cognitive services to enable people to build their own machine learning models.

So, for example, cognitive services, the Federal Trade Commission and the FBI have been able to build new AI power tools to track down scammers. With the bot framework we are finding a tremendous number of startups, starting to deploy health bots that give people much more ready access to health

care advice.

And with things like the cognitive toolkit or CNTK we are giving innovators access to extremely large amounts of computing infrastructure in the cloud.

And all of these together are just the tip of the iceberg. They still require a level of expertise and training craftsmanship, but they are the first steps I think on a 100 step program to provide the tools for people around the world access to AI and its power.

Now I wanted to make a closing comment. There was a lot of discussion this morning about exponentials. And I think exponentials, of course, are in a way if you wanted a logo for this age, you could use this term, use this graphic.

Exponentials indeed have had the transformative effect, but they are also markers of inflection points in human history. And I thought one interesting way to illustrate this is that in this case we can overlay the development of books. And so just before the Gutenberg Bible in the mid 1400s, historians estimated there were about 30,000 books in all of Europe. They were precious commodities controlled by the state or the church. After the invention of moveable type and the printing of the Gutenberg Bible within 50 years there were eight to nine doublings both in the spread of printing presses throughout Europe and in the access to books. And in fact, by 1500 historians estimated there were 12 to 13 million books in Europe. This had a transformative effect.

You could think of this as a Medieval law. It ushered and laid the foundation for the Renaissance and had an uplifting effect on all of humanity. We may be in a similar period today with the emergence of practical machine learning and AI. But we must not forget also the challenges and disruptions that this causes. For sure the church was disrupted by the emergence of moveable type.

But ordinary people were also transformed. By 1550 economies throughout Europe had changed to where it became a necessity to learn this very difficult skill, namely how to read. A skill that requires years of formal training. And as we think about the emergence of AI we should also try to be thoughtful about the disruptions that occur and the level of new training and New World view mindset that we must spread and instill in people. And so as we look to the future of AI for good, the -- I would ask all of you to think back to how you would feel in 1550 with a sudden realization that your children must now learn to read and try to project today what it is that our children must understand and learn about machine learning and Artificial Intelligence. Thank you.

(Applause.)

>> WENDELL WALLACH: You did a really wonderful job in giving us some insights on where industry is in the development of AI. I want to shift us in a slightly different direction. And that's to the oversight of AI and the oversight of research. What should be researched, whether any parameters should be put on that. Microsoft is one of the founding partners in this new partnership in AI. It is not totally clear yet what functions that partnership in AI will or will not serve. But as a member of the management team of Microsoft, an industry leader, I'm wondering whether we can rely on you and other members of the management team to engage in effective industry oversight, or -- we have had a pretty mixed history with that in the past. Or whether the pressures you are under, economic and otherwise, whether they are such that you need more direction, more oversight from the community represented here.

>> PETER LEE: So thank you for the question. And we are actually very proud of our participation in the partnership on AI. And I think the partnership on AI has been making tremendous progress in really setting its terms of reference. I actually don't see this as an industry versus nonindustry issue. But really something that pertains to the entire technology and even the research community. I remember in 2009 when there was a controversy in Iran. There was a contested election. And a tremendous number of people were using social media in order to speak freely and to assemble freely and protest of these contested election results in Iran. And you might remember that there was something called the Iranian Twitter Revolution and suddenly all day all the Internet traffic in and out of Iran was cut off. It turns out I had just started work at DARPA and the State Department of the United States was quite alarmed that it came to light that the U.S. developed technologies to detect cyber intrusions in to corporate and Government networks were being used by the regime to hunt down and in some cases execute people who were using the Internet and social media.

This was admittedly for the research community itself, in my experience one of the first times when the simple lessons about the dual-edged nature of powerful technologies came to light. I was unaware of any academic researcher at the time that thought the cyber intrusion algorithms and technologies could have such mal intent and mal purpose.

There has been tremendous progress in the research community since 2009 to today, tremendous enlightenment about that dual-edged nature of the technologies that we develop. And I think it is absolutely up to not just industry but industry and research and the technology innovation communities to keep advancing that enlightenment. It is not something for industry alone but really for the entire technology community.

>> WENDELL WALLACH: Thank you very much.
(Applause.)

>> WENDELL WALLACH: I imagine that all of us here can think of ways in which AI can be used for destructive purposes. So we just heard a bit about some of the beneficial developments underway. But we all are aware of cyber warfare, of this new form of propagandaization called the weaponization of AI through misinformation and social media for political purposes. And there are outstanding questions about lethal autonomous weapons and other ways that AI might enhance weaponry.

We are lucky to turn to Ms. Izumi Nakamitsu who as of May 1st took over her position as Under-Secretary General and High Representative for Disarmament Affairs. She comes with extensive experience not within the UN system but also outside. So if you don't mind joining us up here.

(Applause.)

>> IZUMI NAKAMITSU: Thank you very much. Since I am new in this position I'm going to throw away the traditional UN linear approach to speech making and think very exponentially and go straight to the point, I think. It has been a fascinating discussion to a newcomer. I think it is really important that we spend the majority of the time looking at and also discussing how these and future AI innovations can be used for the betterment of us all.

But also essential where I come from to ensure that these technologies at a minimum do no harm. And I'm sorry if I sound so negative. I lead the UN's office for Disarmament Affairs. So I come at this issue from the perspectives of considering the military and security applications of those technologies, particularly relating to weapons. States and militaries are no doubt thinking about this, too. Some are actively pursuing AI enabled military technologies. To date militaries seem most interested in AI technologies that augment human capabilities with specific skills, rather than broad-based AI that could replace humans all together. Examples in development already in use include soldier support robots, autonomous transport, escort vehicles, mission command systems, planning and logistics tools and autonomous weapons systems.

And it is this last innovation, autonomous weapons systems that I want to focus on today. There are already autonomous weapon systems in service that are capable of selecting and attacking targets generally in limited environments far from civilians. Multiple states have deployed Naval air defense systems that engage targets autonomously.

Others employ munitions designed to counter surface-based air defense. Guard robots in the demilitarized zone between the Republic of Korea and the Democratic People's Republic, North

Korea, has an autonomous zone. Thankful it has not been utilized. First there is no technical barrier to developing autonomous weapons that can operate in other environments. In other words, there is no technical barrier to the creation of machines or algorithms capable of making the decision to take a human life without direct human involvement.

Second the prospect of fully autonomous weapon systems raising many fundamental concerns for international peace and security as it might have implications on norms and mechanisms and instruments that have been governing the field of warfare. The increasing automation of the battlefield and the growing separation between the user and the subject of deadly force is likely to lower the threshold for the use of that force. It could also strain legal protections for civilians.

Third, these weapon technologies will also pose new distinct proliferation challenges and would likely be sought after by people with malicious intentions. Some experts have predicted without proper constraints autonomous weapons will have the capability to inflict massive human casualties at a fraction of the cost of existing military arsenals.

Ladies and Gentlemen, as they grow more sophisticated military AI applications are likely to increase the speed and condition of action of future battlefields. This exponential growth in speed and complexity of military action could lead to unforeseeable and potentially dangerous and damaging consequences given the inherent unpredictability of AI outputs.

There are currently no multi-national standards or regulations covering military AI applications. Member States have not systematically considered the many challenges posed by current and prospective AI technology through the United Nations or otherwise. Without wanting to sound so alarmist there is a very real danger that without prompt action technological innovation will outpace civilian oversight in this space.

So what can we do about it and what should we do about it? We have to start our dialogues and discussions on various aspects of the AI technology and in particular how to eliminate negative impact of those technologies on our future society and security.

Fortunately when it comes to the specific question of lethal autonomous weapon systems the international community thanks in part to the foresight of the Civil Society industry and also the UN has commenced a discussion on how to grapple with this potentially game changing evolution in warfare. A multilateral process known as group of Governmental experts, GGE, will meet in November this year to discuss the issue of autonomous weapons under the auspices of something called the Convention on Certain Conventional Weapons. The Treaty tasked with banning or restricting the use of weapons that cause unnecessary or

unjustifiable suffering to combatants all affects civilians indiscriminately.

The group is expected to discuss issues such as how to define what a lethal, autonomous weapon system actually is. The adequacy of existing law to ensure accountability in their use. Ethical considerations, humanitarian considerations and the impact of autonomous weapons on international peace and security. States will also discuss what they consider to be the acceptable degree of human control over the lethal functions of a weapons system and whether a specific international Treaty or instrument is required to ensure that control is maintained.

I think that the process in this Convention on certain conventional weapons is a variable one. In part it demonstrates that the UN can play a convening role and provide an inclusive platform for taking forward denigrations on emerging technologies with global consequences.

Ladies and Gentlemen, new issues require new solutions. That is obvious. By seeking to address these concerns and challenges now, we can probably get ahead of the curve. We can begin to build the inclusive partnerships and networks required to tackle the challenges over the 21st Century.

As we know at the very beginning of our conversations in this regard let me suggest in my concluding remarks two essential elements which I think should guide our future dialogues on these issues. First, the tail must not wag the dog. Human-centered norms must frame our future. Technology should be a tool towards that future, not its determinant. We must be guided by the overarching objective of upholding the universal values and norms which we have, in fact, developed over the past few centuries. And of maintaining transparency and accountability in the use of force. Ensuring human security and degradity must be the guiding principle in the dialogue, in the space where international security and technologies intersect.

Second, multi-sectoral and multi-stakeholder exchange is crucial as we need to create a broad partnership that goes much beyond Governments in this area. Our understanding of the potentially revolutionary implications of AI technologies is tentative at best. We must connect Governments with the tech and academic and research communities, industries as we are doing here at this Summit.

Civil Society activists, the scientific community and the private sector have been some of the loudest proponents of the need to address autonomous weapons. They have also been instrumental in bridging -- bringing global attention to the issue of the weaponization of AI. A case in point is the open letter issued at the 2015 international joint conference on Artificial Intelligence calling for a ban on offensive

autonomous weapons beyond meaningful human control which today has over 17,000 signatories. And I think many of you here today are amongst them.

I think this illustrates the fact that the public and private sectors are on the same page when it comes to this issue. This is not a case of industry fearing Government regulation but rather of industry seeking to work with Governments and international organizations to ensure their work is not misused. With a long history of engaging with Civil Society the UN is trying to strengthen its partnership with the private sector in this regard.

The misuse of revolutionary technology with potentially catastrophic results should galvanize us in to forming a serious and substantive and creative partnership with Moonshot creative thinking capacity as well.

Through intrusive networks and partnership we can begin not only to think about how to deal with the potential dangers but also devising innovative ideas of how we can get the great gains of AI and the betterment of all of us. In this endeavor I look forward to working very closely with you. Thank you.

(Applause.)

>> WENDELL WALLACH: Thank you ever so much in considering the -- the ban on lethal autonomous weapons. You are aware as not much of the public that lethal autonomy is not a weapons system nor limited to drones or battlefield robots that might select individual targets but lethal autonomy is a feature set that could potentially be added to any weapons system, including high-powered ammunitions and including nuclear weapons. We are in an area where it is probably incumbent to restrict it to some degree.

Simultaneously these in the research community many of which sign a significant portion of -- which sign that petition that you have mentioned are very concerned that if there is not a ban on lethal autonomous weapons this will totally destabilize the efforts to ensure the safety and control of Artificial Intelligence. And yet there is considerable concern that in watching the CCW process and even the early assignment to the GGE that there is not the will among the Member States to act on this issue. And I wonder whether you have any optimism or whether you feel that something more is needed from this community to ensure that we have a positive outcome on that particular concern.

>> IZUMI NAKAMITSU: I have to be an optimist doing this job. So I am. I remain optimistic because we are really in the early stage of starting these discussions. This is dual use technology. And the technology is being developed by people who are sitting here, the private sector, industries. So I think the

voice that you can bring to these discussions and they are actually also your governments as well, you are able to move the Governments towards a positive. While they might not actually be saying more positive things when it comes to regulating the use of those technologies in to weapons systems at the moment, I do believe that you can actually change the picture.

The disarmament cause is a shared objective. I think, you know, it is actually another, you know, side of the same coin called security. By thinking about how to actually make sure that those -- those technologies would not be misused if that actually will enhance our collective security, then I think it is a very strong argument that we can present to the governments who will be negotiating and discussing these issues.

>> WENDELL WALLACH: Thank you for this advice.

(Applause.)

>> WENDELL WALLACH: Frank La Rue has held many positions in the international sphere, but before that he held positions within the Guatemalan Government. He is the Assistant Director-General of UNESCO and a particular interest in maintaining civil liberties, particularly in the sphere of freedom of speech. So it is nice that we can turn from the darker side of AI to perhaps how AI might be utilized in ways that will enhance human well-being.

(Applause.)

>> FRANK LA RUE: Thank you very much. Let me begin by congratulating the Secretary-General and all our colleagues from the ITU for this really wonderful conference and foresight of having organized it.

And I thank them for having invited UNESCO and for me the honor of being on such a wonderful panel this morning with such great panelists. I think -- I am very pleased to be in this panel and I don't think that we heard so much of a dark side by the way. I think we have heard the voice of reason. All technologies are wonderful. All technological development is excellent when it is put at the service of humanity. In the industrial revolution the artesians tried to stop the industrial revolution by destroying the machines. We all know that this curve that we have been seeing, this symbol of how much it is increasing in speed, the development of ICTs, technologies and moving in to Artificial Intelligence will not be stopped. And we don't pretend to because we think it is a very positive development and can have very beneficial effects for humanity. This is the point that we are trying to put forward for UNESCO is how do we add to that technological issue and advancements of the humanist perspective. We must make sure that technology is at the service of humankind and not substitute humankind in many of its points.

So here we have a decision-making process which is very important for those that are developing the algorithms and the future of this technology. At the service of whom will it be? At the benefit of whom? How far will it reach? Will it reach all countries the same? We have been dealing with, for instance, connectivity with our colleagues from ITU on the broadband commission. And we always talk about the next billion people we want to connect. Of course, we want to connect one billion more and eventually connect everyone, but the idea is where do we prioritize? Are we going to keep on interconnecting urban population which is very easy, very simple. It is a good investment for the platforms, or are we going to look for the rural population where the information is needed for development. Whether it be fisheries, agriculture, or any other form of development.

So it is very important what is the social policies combined with the technological policies and this is our issue. UNESCO is an institution that works on the minds of people. We are the institution of education, science, culture and communication. And our position is all those elements, education, sciences especially here, culture and communication information ICTs and all the policies of communication should be put at the service of developing free minds and freely the communities and to develop their potential to build their own future and to reach the right to development. So this is very important.

On Artificial Intelligence we would begin by saying and we have discussed internally we have to begin by defining terms. Most people out there in the public opinion don't quite grasp exactly where does the line divide because yes, we can have algorithms and we can have data processing and faster data processing and bigger data processing, but when does it turn in to Artificial Intelligence? When does it turn in to machines that can actually begin analyzing and making decisions?

And when is that still responding to the decisions that we have made in terms of policy and not deciding for us or instead of us? And here let me give you an example. In this precise development of ICTs and I used to be the Rapporteur here in the Human Rights Council and we used to say Internet is the big development. Everyone can speak freely, but at the same time the Internet brought a level of polarization in the world that we had not found. We thought it was going to be the gender equalizer. And we still find that the gender gap in connectivity and issues of Internet grew wider. Why? If it is such a technology, it is open to everyone to use. And it is for all of us to be there to be able to enhance our potential.

Why didn't we use it appropriately? This is exactly the challenge we are saying today. For us from UNESCO we have

applied four basic principles. We had a conference on ICTs in general. And we allowed all -- this was approved by all 195 states that are members of UNESCO and the principles were basically four. We try to keep it simple to make it acceptable to everyone. They established what are called roam principles. And what they have said was all ICTs, technologies have to be rights oriented. Have to be the exercise of Human Rights. So we cannot have -- we heard from the Director-General of the World Health Organization, this massive data processing can be really wonderful for controlling epidemics, can also be a breach of privacy for billions of people.

So it depends on how we use it. It has to be rights oriented for everything. O is for openness. It has to be open, neutral and free at the lowest cost for everyone. And it has to be -- A is for accessible. It has to be accessible, but especially has to be accessible for the poorest communities.

And finally the M is for the multi-stakeholder dialogue. And this is something that I think on Artificial Intelligence we are still lacking. All ICTs for us and all these new technologies should be a product of a multi-stakeholder dialogue. I know it is very difficult. The multi-stakeholder does not facilitate decision making. Those public policies have to be based on the possibilities of all sectors of society speaking. It seems at this moment everyone decided let the technology evolve on itself. And there may be some benefit to that. It may go faster, but at the same time it is going in different directions. There is a concern, for instance, I have heard friends from ITU and other organizations that in -- the protocols of the data processing were very different from the Internet of Things. And it is going to be very difficult to compatibilize because all of these have different protocols, but all of them are accumulating information on millions and millions of people.

Who is gathering this information. Where is it going to be kept. Are they going to eventually be interactive, put together. When you begin to have intelligent machines that are able to respond and they become toys for children as we already have that are able to process the dialogue between children and respond in their same terms and same words, imagine the impact that has. Or when we can have machines that can process the dialogue of sectors of the population and then massively respond to them as has happened for electoral purposes. How much of a dominated society can we build and we are not presenting a horror story.

I began by saying that we support the technological development and we believe it is very positive but we are just formulating the other side of the coin. All these policies have

to be balanced with social policies.

And let me go in to two examples. One I keep on hearing that education is going to change. There is a fact that recently I read an article research that there was a fact saying 60% of the jobs in the future will be -- will have no reference to the education system we have today. No one really knows where the 60% came from and who invented the 60%. In reality we don't know what the percentage is. But the truth of the matter is that education is going to change. That we will have technological education, but does that mean that we should leave everything else?

Education is not only to develop our technological skills and better technology and for us to make more intelligent machines. Technology education -- technology will evolve on itself because it accumulates education. But education is to build deeper, more profound, rounded up more human beings and this is important. So we will never really abandon education in a way. It will change. I understand. And it is already changing. We have education at a distance. We will have different needs. But we still want to have students that appreciate the arts, the history, the lessons of humanity and that part will not change.

So oftentimes I have the impression that we are abandoning education as the current education system is useless. And we have to begin from scratch and redesign. We need a profound reform. But it doesn't mean that everything that has been done for centuries in humanity was useless.

The other example which I think is also very important is how we allow the decision making process. And especially in policies that will relate to job opportunities. There is the supposed feature that obviously technologies will inevitably provoke unemployment. That may be true in the sense of technologies if they actually work. And we get to have these intelligence machines that can take over many of the tasks that we today do, many human beings won't be needed to do that. But it would be a huge mistake if the states of today don't generate the social possibilities of people to interact with themselves and society to make their living sustainable. Because otherwise what these technologies can generate is a social chaos. And the social chaos will eventually erupt and will affect technology and everyone else.

So again here is where yes, if we know that unemployment is coming for some of these areas, okay, let's prevent it. Let's have a policy of what to do with these individuals. Or in the opposite direction, these new technologies can actually generate new jobs and provoke new possibilities for individuals, especially in the developing nations. Are we actually reaching out to developing nations? Because if we allow this gap to

build, one of the biggest crisis for Europe today is the crisis of refugees. Before the wars and armed conflicts around there was also a crisis of migration, massive migration. Why? Because we don't build equal opportunities in other parts of the world.

So I think it is very relevant that the social policies be thought by states at the same time that the policies of education and technology that lead to these new technologies. And it is very important for the corporations that are running with this issue. It is very important for the corporations to have a sense of nationality of nation -- of nation building and of responsibility, not only human responsibility with workers but also a citizen responsibility with their country and with the international community.

Ultimately we want all these technologies to enhance within the new agenda the UN has established, the 2030 Agenda, Sustainable Development Goals. We have to focus on this technology whether it is eradicating poverty, whether we are generating gender equality and giving equality opportunity to women or on goal 16 whether we are generating peace which was already mentioned for societies, inclusiveness, with justice and access to just -- with transparency and full access to information for everyone to have an active citizenship. This is the importance.

We are building machines that will have artificial technology but that will be put at the use of developing the Artificial Intelligence that will have the benefit of developing new procedures for society that will support the societies and the development of their own intelligence and their own citizenship and their own democratic process. This is what really enhancing peace is. And I think this is the desire of the Secretary-General. And his message today is if we enhance these Sustainable Development Goals and we begin by building societies and peace with justice we will be able to develop everything else. Thank you very much.

(Applause.)

>> WENDELL WALLACH: Thank you ever so much. You rightfully underscored the extent to which our educational systems are probably not training children for the jobs of tomorrow. And I think all of us are also aware of how entrenched dysfunctional education systems become in nearly every society in the world. I am wondering whether you have thought about -- thought about perhaps UNESCO or another international body taking initiatives that could set guidelines and directions for what might be needed even though we recognize the states might be low to accept the recommendations of an international body.

>> FRANK LA RUE: Actually thank you for the question. We are actually establishing indicators, specific indicators for

these four principles that I mentioned, for the four principles, but one of the ones that we are giving big importance what's the experiences of multi-stakeholder dialogue. Do we have that in societies. To talk about new technologies and use of technologies and oh, for the benefit of whom. And this I think is one of the first steps we should all take. It is an easy step. It is being promoted by some corporations. Certainly by -- demanded by Civil Society and as a matter of engaging state representatives. And I think it would be very important.

I would begin by the multi-stakeholder dialogue. And you can derive not only sort of suggestions on the path of the new technologies but suggestions for education and health programs and how to enhance health and building peace in the world.

>> WENDELL WALLACH: Surprisingly that was the other question I considered putting to you. So I'm going to put it to you quite quickly. Who should be engaged in that multi-stakeholder dialogue? The group really representative of multi-stakeholder dialogue or how do we get the public engaged?

>> FRANK LA RUE: That's a great question. This is a dialogue that needs some degree of expertise. Yes, you want all different sectors but it is not like a normal public dialogue which you can just open the doors of a stadium and bringing everyone in and see. I think it has to be highly specialized forms of dialogue. But not for being specialized should it be elitist and it should have clear representation of organizations of society and have the business corporations. It should have state representatives and it should have organizations of Civil Society and social organizations of workers and peasants and organizations of women that would deal with children. This would be important. The organized sectors of society.

>> WENDELL WALLACH: Wonderful. Thank you.

(Applause.)

>> WENDELL WALLACH: Well, we have heard from Civil Society and we have heard from industry. And now we are actually going to hear from one of the men who is developing these technologies and in fact, there are a few researchers as closely identified with the machine learning approaches that have created this recent revolution as Yoshua Bengio. He -- his primary research ambition is to understand the principles of learning that yield intelligence. And for those of you who are academics you might appreciate that his research is so widely cited that a recent search by Google scholar in April revealed 65,000 citations of work by Yoshua Bengio. So we appreciate your coming up and giving us more of the perspective from the side of the research community.

(Applause.)

>> YOSHUA BENGIO: Thank you. So I'm going to start by

trying to tell you a little bit about AI and machine learning, if I can get the slides started. Thank you. So since the beginning of the research on Artificial Intelligence in mostly the '40s and '50s scientists have understood for computers to be intelligent they need knowledge. And, of course, it is not enough to have knowledge, you have to be able to use it. How do you get knowledge in to computers? And this is where machine learning is really making a difference today. Because for many decades scientists have tried to give intelligence to computers by taking the knowledge that we have and that we can communicate, that we write in books, that we talk about in mathematical formulas and give that kind of knowledge to computers. And, of course, we are still doing that. This is what traditional programming is about.

But there is a limit to that way of giving knowledge to computers because a lot of our knowledge is intuitive. It is something that we know but we can't easily explain to someone else and even less to a computer. So how do we solve this problem? Well, machine learning is an approach to AI which tries to address this issue by allowing computers to get knowledge from data and, of course, we have learned a lot about the importance of data. And we will continue hearing about it. And I will tell you later. Computers can observe the world and interact with the world and abstract knowledge from it. How do we get more knowledge in to computers in a way that allows them to be more intelligent? Something happened in the last decade that really is transforming machine learning. And it is the approach that my colleagues and I have worked which is called deep learning. Approach to machine learning and inspired by the brains and it is centered on the idea that computers can better understand what they are supposed to think about by considering how knowledge or information rather is represented and allowing computers to figure out good representations.

So it is all about representations. Representation of being centralized to AI from the beginning, but the new thing is that computers are learning, are discovering by practice and from data how to represent information. And furthermore, the particular thing about deep learning is that the -- these neural networks that you have heard about have multiple levels of representations. And that these different levels correspond to different levels of abstraction. As you have seen in earlier slides, for example, from Peter, the computer will extract different levels of representations say from images going from low level presentations like edges to higher levels that capture parts of objects and then objects. And we can think of this in many areas of understanding of data.

So deep learning has been very successful in the last few

years and mostly it is being thanks to the ability of deep learning extract good representations from data and that's basically an act of perception. It is most obvious in speech recognition but also in computer vision. So the computers can now look at an image and make sense of it and figure out which objects, which persons and so on are present in the image.

More recently we have made progress in the ability of computers to use and understand language. But we are still really at the beginnings of this. And although there is some things that we can do that are really amazing we are very far from having computers that really understand us and can have a general purpose dialogue with a human. In the slides this is something we did a few years ago in my lab and Peter shows examples of computers reading an image and translating that image in to a sentence. When we did this we use the same systems that we were developing for machine translation. Systems you can now use after a lot of engineering and a lot of data being put in. For example, in Google translate. You see the computer being trained in a supervised way with pairs of an image and a sentence that a human said would be a good caption for the image. Another example of breakthroughs that we have heard about is computers to take a sequence of decisions and this is happening in reinforcement learning. This is still really at the beginning. We are far from having cracked that nut. And this has been pretty spectacular with the advances of ability of computers to play games like Go.

In fact, it is instrumental to other applications like dialogue for language as well as self-driving cars which are heavily using the ability of computers to understand images. So it is interesting to think about the advances that happen recently but really they are the result of slow accumulation of scientific advances, incremental advances that have happened over the last few decades. And it takes time before these scientific advances translate in to applications. It is only since around say seven years ago that we started to see breakthroughs in applications like speech recognition and a few years later computer vision.

Machine translation has made big advances and impressive scientific advances and the ability of computers to reason with deep learning and to pay attention. Attention was a central feature of our advances in machine translation and to take advantage of external memory and these elements are pretty much in the lab and are not yet things that you find in applications.

We are seeing I mentioned reinforcement learning, playing games. These kind of advances are currently making their way in to robotics because traditional robotics can't use machine learning. It used handcrafted equations and control mechanisms.

And a lot of handcrafted rules but it is only recently that we are seeing these techniques enabling robots to be less mechanical and be more adaptive and be able to deal with situations that are not planned too much ahead of time.

One thing we are working on a lot in my lab is unsupervised learning. That's the ability of computers to make sense of data that humans haven't labeled for which humans haven't said oh, this is a car. This is a dog. And this is how to interpret that sentence. And a really exciting area is how unsupervised learning allows computers to be creative and generate new things that were not part of their training data, to create and synthesize new images and speech. And this is something that is happening in the lab which will have an impact in applications. But is still far from an application right now.

Something interesting is happening with all the excitement about AI and machine learning and deep learning is that more and more researchers are choosing to work in this area. And the effect of that is that we are making very rapid progress in many areas of application. So this figure shows the progress has happened between 2011 and 2015 on object recognition and images but this is happening everywhere. So this, you know, kind of positive virtual cycle is happening in research because AI is making progress and more people are studying it and using it and developing applications and improving the science at the same time. So we have already heard about applications to medicine and in particular I think the low hanging fruit here is medical images. In my group we have been working with a company that is applying this to detecting cancer cells and colon polyps from video. And video takes a lot of time for doctors to stream. And they can miss little details. So they now have systems that for a particular kind of cancer here, these colon polyps can do better than the best doctors. And in fact, those best doctors are very rare and those -- and these AI systems do much better than their regular doctors that have been trained to do their job but maybe not the best at doing. This is very exciting.

But one issue that you immediately face when you try using machine learning with medical data it is hard to get the data. The data is there. But it is being hoarded by hospitals, by doctors, by organizations that have, you know, good goals of protecting the privacy of individuals, but right now we are not striking the right balance of building something for the good of everyone's health and protecting individuals and privacy that their data is not going to be used against them. And we need Governments to get in to this, to make sure that the rights of individuals are protected.

For example, in Canada recently there was a decision to essentially prevent companies like insurance companies to use

genetic data in a way that would essentially deter patients from sharing or even testing themselves for that kind of information. So we can change the laws in such a way that it is going to be easier for people to share their medical data and be confident that it is not going to be used against them and not for other reasons than those intended.

Now there has been a lot of excitement in the first session about where AI can lead us in the future. But I would like to hear -- say a few words where we are really now. Very, very far from human level AI. So maybe things are going to get twice better every year or whatever, but really we have a lot of research due before we get there. And one thing that I'd like to emphasize is that pretty much all of the major breakthroughs and amazing applications we have today of AI and machine learning are based on what we call supervised learning, whereas I said the computer has to be told what the right answer is on millions of cases or more. And that greatly limits the scope of applications we have. And it forces us to really think hard. When we look at the kinds of applications that would help say the UN as we are talking about today we have to look at datasets that are large enough and for which we -- we have not just data but the kind of data that tells a machine maybe after the fact what the right decision would have been or what it should have predicted. And this is not the case of every dataset.

If I get some medical data where only CD images I don't have the doctor's characterization, it is going to be useless for current technology. We can use transfer learning and we can do what's called supervised learning to help things. But we rely a lot on these labeled datasets where we know what the right answer is. And because of this reliance on supervised learning it means that the assistance we currently have will make all kinds of crazy mistakes. What you see in the picture is a system that will equally say that the picture on the left is a dog but the picture on the right is an ostrich. It is not a randomly chosen image on the right. It was obtained purposefully by changing the input pixels just slightly. So this is in part due to the limitations of the current science. And I'm getting this example for people to understand there are limitations and that, you know, the current state of AI is not magic and far from that.

Okay. So how do we go from there to making good use of AI? One good thing is that even I just told you that we are far from human level AI, is that the science has advanced enough that even if we stop all basic research right now we will have at least a decade of useful applications and reaping the benefits of today's technologies. And what it really takes is will and efforts, scaling up the engineering, the data collection,

aggregating the data properly, putting enough engineers on the job to build those systems, using additional competing power because that's going to continue to improve and being creative. And this is where I think this conference could be most useful in thinking about ways to use that technology.

So even though it is limited it can have a huge impact as we have seen in the example we gave but now try to shift gears. We have been thinking about how we can use machine learning to develop the next gadget. It is going to be fun to have a newer phone, but it is time to think about how we could use these same algorithms for the greater good. And I think it is quite possible but it needs creativity and not just from the side of the researchers because we are focusing on our algorithms. So we need to work with the people who have those problems to solve and, you know, find the right way to deploy these advances.

So, of course, we have heard about the potential impact of AI in jobs. And there will be a lot of wealth created. It is hard to say how many jobs will be displaced and when and there are a few studies that predict that half of the jobs will be impacted one way or another. So this is -- this is clearly something important and, of course, it is also something that was raised earlier about how we should change our education system to prepare students to a world where which job will be there is not so clear. So we should prepare them to be more generalist and focus not just on technology but kinds of jobs that computers won't be so good at least for the next few decades that involve the human aspect.

The last slide is about the breakthrough session that I am leading tomorrow afternoon which asks the question of AI for prosperity. And then there are different aspects of this. Again earlier we have heard the question how are we going to collectively organize ourselves so that the wealth that will be created will be distributed in a reasonable way that will not increase the concentration of wealth, the increase in equalities. And not just within each country and talking about this year in Geneva but also between countries. I think this is not at all obvious.

How do we make sure that Developing Countries can take advantage of the technology? One thing that I find interesting and that was mentioned in a developing world many, many people have a phone, a Smartphone and I think this is something -- this is a vehicle to put AI in the hands of billions of people.

So as I mentioned we want to focus also the research on applications with maximum beneficial impact. Academia in recent years and I can take my lab as an example has been driven by the funding of industry. There is a lot of interest in academia and academics will be working more of the social applications of the

organizations like the UN who are looking at applications which may not be profitable but really could have a huge positive impact for humanity. We talked about medical applications but environmental applications, building services say based on natural language that could help people around the world that have access to a phone, for example, that could be used for health applications, could be used to help people deal with legal issues. All kinds of things could be done with the technologies that are being developed.

The last point I want to make is about talent. Right now as Peter knows very well there is not enough people with the expertise in machine learning and deep learning for the demand that exists in industry. That's not speaking about the kinds of applications that we want to address mere in this Forum. How do we make sure that the few people that have those skills spend more of their brain cycles in thinking about the kinds of applications we care about here? I think this is a really important and difficult question. And one place where I think there is an obvious fit is academics, because academics tend to care a lot about the positive impact of their work in society. Academics already get funding from Governments. So they don't need to work on applications that are economically viable in the short term. If we do the right things, if we put those academics in contact with the people who have the data and the social problems I think we could really do something really great for humanity. So I'm going to close on this. Thank you very much.

(Applause.)

>> WENDELL WALLACH: I'm certainly cognizant that everyone is hungry and we should have broken for lunch a few minutes ago but, if you will bear with me five more minutes I would like to put a question. One of our earlier speakers referred to what is sometimes referred to as the pacing problem, scientific discovery and technological innovation and the increasing lag with our ethical and legal oversight. There is also the considerable concern around these deep learning algorithms and their inability to explain what they are doing. They might have algorithmic bias or lead to some tragic event in the way the output was applied.

My question to you is the standard question. What is the responsibility of the researchers in regard to this? Should they consider slowing down certain forms of research if they can't guarantee its safety, or does that responsibility lie elsewhere? You know, what are we going to do about this problem?

>> YOSHUA BENGIO: Regarding your last question, the progress we are making with natural language understanding and generation is something that gives me hope that we will see more and more machines be able to explain in words what they are

doing. But we have to keep in mind that it won't be much more precise than a human describing the reasoning behind their decision. I mean when we ask a person why did you do this they will come up with a story. And that may give maybe a high level description of the reasoning but it is usually incomplete. The same thing will happen when we use these kinds of machine learning systems that are complex and have billions of parameters. And the real explanation is not something I can write in a few words. There is a lot of interest in the question you are raising and a lot of researchers are trying to find ways to deal with it. And you can come up with partial answers, like I said some natural language explanation as well as pointing more quantitatively at the variables that already matter. And we can also change the training of say those neural nets so that they will be insensitive to variables that we don't want them to take in to account, like say gender, race or things like that. This is something we know how to do already.

So there is research towards answering this question. We need more of that. I don't think that putting the brakes on the research is going to be helpful. I think where the regulations need to be really is when those systems get to be used by people, mostly in products. But anywhere those systems are going to be used we need regulations, for example, to impose on those products that they would use whatever is the current state of the art avoid biases. So we know how to do it but it doesn't mean that a company is necessarily going to put it in their products.

>> WENDELL WALLACH: We will leave it at that. I am sure you will all agree with me this has been a remarkable series of presentations. I had wanted to pose one final question to all of the panelists but recognizing how hungry we probably all are we are going to bring an end to the session. So I will pose that question to all of you, what is one specific task, one actionable task that you believe we can work on here over the next three days and should be working on to ensure that AI will be for a global good? Thank you very much.

(Applause.)

(Session concluded at 1:08 p.m. CET)

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.
