Raw file.

May 15, 2018.

1345 CET.

ITU.

AI for Good Global Summit.   Transformations on the horizon.

   >> Dear captioner, please can you give me five line breaks, please?

   >> Hello, this is an audio test for the captioning.

>> We will resume at 1345.

>> This is an audio test for the captioner. Please, can you give me one break line..

>> Welcome, everyone the to the first plenary of this exciting summit, I'm the Director for disarmament at the United Nationed based here in Geneva and I have the honor to be moderating facilitating this sessions about transformations and artificial intelligence based technologies over the horizon. We have a stellar panel with us today, I will quickly go through just the names and then introduce them as we go along, professor Dr. Wolfram Burgard, first one here, Terah Lyons, Celine Herweijer and Wendell Wallach. We will have a very diverse but somehow we weave it all together discussion from the evolutions and transformations in the technical space and AI and robotics into the ethics space and the governance space but also some of the areas of impact including on climate and other SDGs.

Since this is an AI summit after all, I will start out with some tech issues, I know you have all been informed that we can engage interactively through using pigeonhole where you can ask questions throughout the presentations by the speakers, we are going to make the session very interactive, we are going to go back and forth between former presentations, Q and As, interaction between the various speakers, so please use pigeonhole. You should have it on your devices now. Ask questions. I'm paying close attention. I'll be using these questions when I interact with all of the participants, sorry, all

of the presenters coming up here, one after the other, and also in our engagement down at the podium here.

I'm not going to spend a lot of time, we have limited time and we are keen to hear what our expert presenters will bring to us of their insights and knowledge. With no further ado, I'm going to hand over to the first speaker of the day, which is professor Wolfram Burgard. He is with the department of computer science of the university of Freiburg. He is most groundbreaking work has been the area of what is called probabilistic robotics where he developed a efficient means of estimating the position orientation of a mobile robot efficiently even without knowledge of its initial position. He will explain to us what that is all about. But this work has had significant impact on the engineering development of robots, and in addition to several of his prizewinning research he's also involved in other projects and I will be curious about this one particularly, he also developed the first robotic museum guide for a international museum project. Clearly, very versatile applications of the technologies you are working on. We are excited to hear your take on lay of the land and AI and robotics.

>> WOLFRAM BURGARD: Thank you very much.

(applause).

Thank you for inviting me. It's a pleasure for me to be here, and particularly I like the topic of this conference, because I think AI we need to transform the way we think about AI in Europe and this

topic is actually the right one, we need to give it a positive attitude and get AI to technology that is going to help our societies in various aspects. What I'm going to do in this talk is briefly tell you about a few things that we have been developing in the past and which address some of the, from my point of view, needs of our society and we will talk about other works from people from others, other work from other people that actually look into particularly application areas of AI.

Let's see as to whether this is going to work. This is the very first challenge, first is does AI work and second does the presentation work. We will have a look at this. Here we are. The second one is to whether this device is also going to work. Wonderful.

In AI we are interested in and this is the view of AI building systems that can interpret sensory data, create internal models, and then develop activities out of this or reason about this and create the next best action to take.

This cannot only be software agents, it can be physical agents like robots. This is the basically, the technology that we are working on in AI. In general what we are interested in, is creating machines that maximize their performance. Doing this in various tasks, when you look at our daily life, you see AI in many tasks, many situations, for example, when you talk to your cell phone or intelligent assistant at home, or when you start your information gation system in the car -- navigation system in the car, basic AI algorithm, you interpret

text or convert text into or pixels into text, or when you classify your images, and try to find your images in your photo library, or on the web.

Finally there is recommend systems and website ranking, these are all famous AI algorithms, well-known AI algorithms that we use in our everyday life.

But there is going to be more, like in games, for example, where AI has taken over somehow, and where people actually talk about the AI as kind of like an agent by itself, but actually it is a science that wants to develop intelligent systems that can play human like or even better than humans.

Interestingly, like AI has in certain cases even started to outperform humans, and we call this superhuman performance. One of these examples is Watson system that a few years ago actually won in the show Jeopardy, quiz show, against the so far two most successful Jeopardy players.

This is an interesting aspect of AI, that when we focus the topics strongly enough, we can actually achieve performance that exceeds that what humans typically show. This started with games, but has now turned over to tasks that we think require true intelligence like being a game that we regard as people who win this are really intelligent people, also a area where AI can win, but there are more application areas where we, until now, think that only experts can handle this. One of the most for me surprising experiences was at

the win of the big data and deep learning technology in Go where because on our slides in the AI course we teach the students, we have the slide, Go is probably one of the games where AI in the next 20 years is not going to win against a human, or the super players. Within two years, actually it turned out to be not true. We had to remove the slide from our lectures.

I'm wondering, I'm curious what else is going to be there in the future. The tool that we are using in AI and this is also important, are search and optimization, for example, searching the closest route logic, representing information, probabilistic methods and uncertain reasoning, classifiers, statistical learning, and nowadays artificial neural networks and evaluating progress, having developed a lot of tools to make sure that we can measure our progress that we have made.

I want to talk about the few examples where I think AI will be a, play a major role and definitely you heard about those this morning, but I will look at this a little more from the research side so that you can actually see where we are, and that we haven't solved all these things. But there is still development going on.

The first one is logistics and manufacturing. Turns out that nowadays, robots are mostly bolted to the ground, and they are therefore highly unflexible, but specialized to a specific tasks. Humans on the other side are flexible, and also can produce very versatile, form very versatile tasks, but they are not as efficient as robots.

One approach to solve this problem is actually like make robots mobile. And in this way, achieve that what is called transformative production. In this way you can for example completely reassign a factory floor, and start producing something completely different at the entrance of the production line than you are right now doing at the very end.

This is a key enabler for the future industry, in Europe, which is heavily based on manufacturing and production as well.

One example for this, and this was mentioned before, is when you want to have moving robots, or one key task is navigation. The robots need to be able to get from A to B. This is an example that you see here on the slide of a robot that actually can perform highly precise navigation. These black dots on the ground are to demonstrate about the precision of these vehicles.

This capability is not only important for robots on factory floors, but also for safe driving cars. Imagine you have a car that needs to make a left turn, you want to make sure this car is on the left lane from which it can make a left turn and not being on the right lane. There is a high need for highly accurate navigation.

This can affect the transform to more complex systems like dynamic environments as well, like you see here, this is a laboratory experiment where we modify the environment to measure the accuracy, we can transform this to more production-like looking robots. One second. Like this one over here, where that actually produces all the transports

more than 11 tons and we have a bigger version with 20 tons, and this also is the prototype of a small scale robot and the big version of this, and you can see the accuracy which is in the range of milli meters by which those systems navigate. You can turn this also into larger systems and this is an application that has been built for Boeing and algorithms actually are used in these platforms for the construction of the 777 fuselage, which is moved autonomously through the factory floors of Boeing.

This is one of the key enablers for production, and logistics, but also in fact for self-driving cars. There are people at these companies that right now is thinking about two applications. The first one is parking and the other one are taxi services. This is a quote that I took from Chris, statistically the driver is the most unreliable component of a car, which means that if you can get rid of the driver, we are facing a world with fewer accidents and fewer death casualties. In order to make this world a better place, actually what we should do is take the driver out of the car.

This is what companies are working on, like here is one of the self driving cars in mountain view after a successful taxi ride from the car train station to the campus. In this case the car was successfully able to transport me between those two places. This is a experiment that we did a couple years ago with valley parking where you start a car autonomously in front of a big parking garage and let it autonomously navigate through the parking garage and park

there safely on the roof in this case. The impact of this is many fold.

You can have tighter parking in those, you do not need to have people which means that you can have way more cars in these parking structures and many other advantages that does have. In fact, you can get rid of the cars in the city and we do not need parking in the cities anymore.

The potential for healthcare is many fold. On one hand you have big data and on the secondhand we do have right now tools for learning from the data. Here is one experiment that we performed recently about in the context of Parkinson's disease, where the goal is to develop a automatic system that can tell you the effect of the brain stimulation for Parkinson's disease patients. Basically, usually what they do right now is as far as let them walk and measure around a obstacle and measure the time until the patient comes back.

This is the basic test that they are doing in hospitals right now. What we are doing is using AI technology to classify the movements, in this case with a motion capture suit, there is a little math behind it, I will skip this, but you should know that AI is heavily based on math, and based on this we can actually achieve a performance that is on par with standard Parkinson's disease rating scales and at the same time has fewer variance between individual doctors and also between individual, different assessments of the very same doctor.

We also got rid recently of the motion capture suit which doesn't look like really ready for clinical applications, and we can now do this with cameras and track people with cameras, and the next step is in fact going towards cell phone. You use your cell phone in the pocket to track your movements and assess the health status of patients.

This is the application to Parkinson. Here is another support, assistant for paralyzed patients and I love this video, because of the one scene that you will see in a second, and what happens here is, a robot that detects the face of the human, that wears a BMI brain machine interfacing hat to interpret the thoughts of the person and whenever he thinks about the right arm or right leg movements, then the robot interprets this as a go signal and moves towards the face of the patient in order to serve a drink. In a few seconds you will see what that means. In fact, you will also see why I show this video. We are not yet done. But it is actually the case that such systems are envisioned to help heavily paralyzed patients, for example.

And at the very end, I will talk about a project that one of my colleagues, Sebastian, has done, and it made it to the front page of Nature, we had a similar presentation or task this morning about skin cancer detection. The idea there is that you use your cell phone to take a picture of your skin and then the system tells you as to whether this dark spot on your skin is cancer or, dangerous cancer

or not. There are a few facts about this but in the end this is a deep network that has been applied, and to the outcome of this approach can be seen here. This is the performance, when you compare this to these red dots which are the expert dermatologist performance on the same data set, you see that with these systems, you can even achieve expert level performance, and we are not at the end of this stage, so in the end we can actually get probably better than experts, which means that we can provide many more people with access to medical care.

At the very end, another aspect that we are working on is precision farming, in the end, we somehow need to feed the world. In order to do this, we need to build, one approach for doing this is building robots that can support agriculture processes. There are several advantages of that. In the end, I will give you one example of what that means. So here for example is what robots can do, if that works. Yeah. So there is two pictures missing but in the end what this robot does is extracts the tiny little plants from those pictures and estimates where the rows are in order to perform this navigation. As you know, we are all talking about, and this is in a video, showing this robot in operation, and moving on the field where we can actually perform precision targeting of specific plants in order to achieve specific treatments over there. This is done by two of my students, two female students who love working with robots.

This brings me to the end. This is a video, you see this in action.

This is something I truly believe this is going to be important for our future. At the end, AI and robotics from my point of view are super important for the health and the well-being of our society. We should actually take the chance and invest in this and convert our society in one in which we have better opportunities for all people. Thank you.

(applause).

>> ANJA KASPERSEN: Thank you so much, professor, for that enlightening speech. If I may, we got quite a few questions from the audience already. I will distribute them across the various presenters, because they fit with different presentations and skill sets, in case you don't hear your questions right up front, if I may, you spoke a lot about the developments in robotics field. In your view, are we aligned in where we are in robotics and where we are in the broader disciplines of AI, depending on what subset you are looking at, is there a misalignment or is one racing ahead or is the other one catching up?

>> WOLFRAM BURGARD: There are different application areas. When you think about computer vision, it has been successfully applied in various domains like image analysis or even skin cancer detection, for example. But we think about physical agents, you have this additional problem that this agent should not harm anyone. So false negatives are actually enormously dangerous, meaning that the robot doesn't see that there is a person crossing the road, and we had

such a case recently. This is something where a lot more investment needs to be done, and where these algorithms need to be developed very much further, so if one of the pictures sort it automatically in your photo library is from someone else and not from you, then it doesn't really matter.

But if you hit a person with a car, then this is a fatal accident.

>> ANJA KASPERSEN: To build on that, one of the A Is that cause a lot of concern and also some excitement a couple years ago was called libberATIS, poker playing AI that could operate with imperfect operation in imperfect environments. What is your take on those experiments and what are the potential applications and some of the concerns you might have when you hear about this?

>> WOLFRAM BURGARD: In the end when you think about vehicles or robots acting in the real world, there are so many cases, that it is impossible to programme all of them. This is one of the key problems with the development of self driving cars, that it's hard to catch all these cases and we need to have methodology to actually investigate this. This is why these type of research is so important, because we need to have agents that can operate in a world that they do not know exactly.

>> ANJA KASPERSEN: Thank you. We will come back is to other points. You mentioned a lot of interesting things, in agriult you are, medicine, autonomous vehicles -- agriculture. We will do a round once we go through the presenters. The next one up for sharing her insights

with us is Terah Lyons, Executive Director of the partnership on AI, this is a partnership that was established back in 2016. Right? By Amazon, Apple, deep mind, Facebook, Google, IBM and Microsoft. There is quite a few representatives from the partnership on AI in this audience today as well. The focus of this partnership was very much advancing the benefits and addressing challenges of machine intelligence. She is also the former policy advisor to the U.S. chief technology officer in the White House. She was working on the Obama strategy on AI, and also formally with the Mozilla foundation, technology policy fellow.

Terah, we are excited to hear more about the partnership, what you guys are doing, and also maybe share some of your experiences from your previous career working on the governance side.

(applause).

>> TERAH LYONS: Hello, everyone. I'm going to make sure my AV is working as well. It's exciting to be here. Thank you for having me. I'm Terah Lyons, Executive Director of partnership on artificial intelligence to benefit people in society. It's a chance to reconnect with other colleagues who worked hard to bring the partnership into its existence and worked towards our goals since its founding. There are many here who haven't heard about the partnership. I'll talk about it at a high level from a goal perspective and provide insight into early work that the organisation has undertaken in the five months since we started becoming operational since I joined.

The last thing I'll say here too is that a lot of this work deeply intersects with the work being talked about here at this conference and identifying practical solutions for AI for good. It's exciting to be here in Geneva. The partnership on AI is a community that draws together a diverse set of global voices to push for the development of technology that benefits everyone.

This slide behind me is hard to see, I apologize for that, but it's a depiction of a lot of logos of our current partner community. We were established in 2016 as was mentioned by a group of AI researchers, representing six of some of the largest technology companies in the market. Those were the ones mentioned before, including Apple, Amazon, Facebook, IBM, Google, deep mind and Microsoft.

Though these companies were and are still fierce competitors, they were driven by the need for collective understanding and exploration and interrogation of a lot of the benefits and challenges of machine intelligence. That includes benefits and opportunities. Our board today is comprised of half for profit and half nonprofit institutions including the six founding companies and also six nonprofit board representatives, representing some organisations like the American civil liberties union, the McArthur foundation and open AI which is a nonprofit research institute, and as of today we have evolved into a partnership of over 50 member organisations with a diverse and fairly representative combination of perspectives.

We see as our goal to convene leading technologists, academics,

civil society organisations and advocates to create an open platform for inclusive AI debate, to incite critical self reflection in the AI governance community and technology community to develop best practices, to advance public understanding of AI and also to foster aspirational efforts of AI for societally beneficial purposes.

Everyone in the partnership is united by a core belief that artificial intelligence technologies hold great promise for raising the quality of people's lives as was enumerated by the last presenter and can be leveraged to help humanity address important global challenges from climate change, food security, inequality, health, education and more. We will continue to hear more about the applications through the rest of the programme this week.

We are united at the partnership by an understanding that on tapping AI's highest potential requires active understanding with and grappling with the challenges associated with its development and use.

This is a picture of the first ever event we hosted in Berlin last October. At a high level, our goals are enumerated on the slide but I'll get into that in a moment. Part of the reason why the partnership was created to begin with was because we are at a critical juncture for AI development and its applications. As the promise of AI materializes, its real world impacts are increasingly felt and its promise is matched by challenges and questions about its development and its use.

We confront and address these questions now, is crucial to ensure that we develop AI for all. The societal questions implicated by AI are sufficiently complex and cannot be determined by, we cannot determine solutions alone, excuse me, neither as singular entities nor as a singular discipline. Sore thing that unites the partners in the partnership is an interest in collective multistakeholder engagement with a lot of these topics.

AI has far-reaching consequences for domains and constituencies which need to be involved in the development of technology and technology governance which is another thing that motivates the work that we do.

The partnership is currently in the midst of forming working groups around six of our thematic pillars. These six areas are enumerated on the slide, and are the motivating questions for the field right now as determined by the partnership as an organisation. These working groups that we are establishing are the primary multistakeholder engines of the partnership and they are led by our partner organisations and supported by our staff and focus on research and developing best practices for the testing, fielding and responsible development of AI technologies.

In keeping with the spirit of open collaboration between industry and civil society organisations and the rest of the nonprofit sector each of these working groups is co-chaired by one for profit and one nonprofit representative from our partner base.

Last month we held the first meeting of our AI labor and economy working group in New York City, and that group is co-chaired by the center for Internet society in India and McKenzie global institute. The group is focused on bringing a set of differing global viewpoints together to explore the capabilities and challenges brought by AI to the economy and to labor markets to measure progress and ensure that AI is developed in responsible manner and that its benefits are generally speaking broadly shared.

Last week, we were in London at deep mind to help launch our fair accountable and transparent AI working group, co-chaired by deep mind ethics in society and Princeton center for information technology policy. That group's focus is on important questions around justice, fairness, bias and other related topics in artificial intelligence.

We will soon launch our safety critical AI working group co-chaired by the electronic frontier foundation and Microsoft and that group is focused on the safe robust and responsible development and deployment of AI in safety critical environments which span a lot of different domains, but include healthcare and transportation.

At our heart of the partnership is a deep commitment to AI for social good, and we will soon stand up a working group specifically on this topic actually. Our work will leverage the partnership's unique multistakeholder nature and its capacity to look at how to maximize AI's potential for supporting the public interest in areas such as education, public health and sustainability.

There are already incredible applications of AI in public interest domains, many of which we will hear about again this week. The partnership's specific value proposition in this area is an incentivizing innovation and investment in application areas that are not currently incentivized by traditional market forces in bringing groups together to make sure we apply the full force and function of the different types of stakeholders that we have in our community to those particular global challenges.

As with everything we do, we will seek the research community, industry and add vocky organisations to seek a multidisciplinary approach to tackle global challenges -- advocacy organisations.

Lastly, our work will extend beyond these groups. There is more to say later. But what makes it important is that they like this room actually represent a very broad constituency of stakeholders who are united by their interest in leveraging AI and ensuring that it's responsibly developed. Our aspiration is to grow the partnership over time so that we represent an even more diverse constituency than what we currently do right now.

I'll close by saying that in everything we do, it's critically important for us to remember that technology is merely a series of decisions made by humans. This seems to be a theme already of some of the presentations I've heard today, but it's worth us all keeping in mind. The other thing we are keeping in mind is that we are those humans in most circumstances, and I for one am very much looking

forward to seeing the positive impact that the people of the partnership on AI and those gathered here today have on the future of AI and humanity. Thank you.

(applause).

>> ANJA KASPERSEN: Thank you so much for that, Terah. One of the questions that popped up right now, I'll read it out but slightly rephrase it, not to put you on the spot, now the technology renewed itself. We are witnessing a AI technology race where countries are aggressively investing in AI, countries and companies investing in AI. How do we make sure that, how can one make sure that regulations adopted by a few countries will be applied at a global scale? My question to you is, given your experience, working at the intersection between not any industries but the very companies and industries that are driving much of what we see as the AI developments in our contemporary time, and your experience working for a country that has for a long time been known to also be spearheading some of the technology developments, where do you see the regulatory space, the feedback loops that you work with to make sure that what you gather from gatherings like this, goes back into the developer side, responsible innovation side? Since you have the jewel experience, I'll be keen to hear your views on this.

>> TERAH LYONS: Absolutely. It is a really good question. It is important and complicated one. The first aspect of it is in making sure that we understand that technology policy has to be developed

with more voices than just policymakers or technologists, as they are traditionally considered.

One interesting learning that came from my time working at the White House on these topics is that this inherently has to be a multi disciplinary process and one which is multistakeholder in nature.

A interesting thing to think about when you look at globally scaled policy is differences in jurisdiction and the way in which different governments operate based on the constraints that they face and the priorities that they have.

Inherently governance will be different based on those attributes. But there is also the concept which might be undergirding which is the notion of a shared value system or shared priorities in a global community of creators, and the notion that there is something that commonly unites companies across borders that commonly unites even governments across borders, in the priorities that they undertake in research and development and ethics and value systems that support the development of technology systems.

At a high level, that is the answer I have to give. Obviously, this is something that will evolve over time, and is going to be different based on the application area that I think we see AI applied to as well.

>> ANJA KASPERSEN: You would say that you have a good feedback loop, where you bring these insights into developers?

>> TERAH LYONS: Absolutely, yeah. That is the whole purpose of

the partnership on AI, in fact. The one thing I had glossed over in the presentation but mentioned briefly is the fact that we were created by technologists. That is significant. The AI researchers or the heads of research at the respective companies that started the partnership to begin with, a couple of years ago, were the voices that were the ones that were loudest in advocating for the fact that this needs to be a collective endeavor, and since it's grown to the very, the multiple interested parties that we have represented in the organisation today, which are all necessary to the type of conversation that we want to facilitate, but at the heart of what we are doing is impact back to the research realm and back to developers who have the power to affect the way that technology platforms and products are put in consumers's hands and impact people every day.

>> ANJA KASPERSEN: Thank you for elaborating that. We will move to the next presenter. I'll be engaging with questions that you put up on pigeonhole as we go along. The next one I would like to hear or share her insights and I find this is a great to have somebody with a climate experience on the panel is Celine Herweijer, doctorate in climate modeling and policy and lends her expertise on climate policy, finance, resilience and disaster risk reduction and international development as a partner in PricewaterhouseCoopers, long name now. Sustain ability and climate change team and advises governments and private sector across different regions on these issues. Celine, welcome. The floor is yours.

(applause).

>> CELINE HERWEIJER: Thanks, Anja, it's great to be here. I'm looking forward to spending the next few days here.

My role now sits across both sustainability and innovation. That is because we have seen the opportunity of bringing those two worlds together. Talking about AI, boy, does the earth itself need AI at the moment. That is where I'm going to focus my ten minutes or so, is talking about some of the work we have been doing specifically on unleashing AI applications for the earth's major challenges.

A few sobering facts to set the scene, this is not a happy story but this is where we are. The earth has never been under so much strain, for those of you familiar with the concept of geological epoch scientists say we have moved from the holo scene which is a period of relative stability in the earth system where human civilizations flourished for 10,000 years to another scene. Let's look at biodiversity. It's looking like we are in the middle of a mass extinction crisis, one out of five species are facing extinction. If we carry on, that is looking like 50 percent of species on earth by the end of the century. Chemistry of our oceans has never changed so fast in 3 million years. It's profoundly disrupting ocean ecosystems. 92 percent of population live in areas where air quality is unsafe. In water in ten years we will have a 40 percent shortfall of the water we need to power, give water to our population. Climate change is exacerbating all these risks. Even though we had a great

outcome a couple years ago in Paris of the global climate agreement
if you add up all the commitments our national governments made to
decarbonize the current climate policies that are being put in place
since Paris, we need to get to 2 degrees and the national panel is
about to release a report saying even 1.5 degrees with serious risks.
That is the backdrop.

It is not a good news story. The good news story is that at the
same time, this is all happening, we are at the beginning of the
fourth industrial revolution. The fastest period of innovation ever.
AI obviously is one of the most pervasive of the four AI technologies
we are focusing on here these days. It's a exciting time for AI.
We are at the point where start-ups and industry alike can start
using the technologies that a decade ago a handful of research labs
had access to. A few data scientists can make a big impact using
AI. At PwC we spend time working with our corporate clients and
business clients on how to create strategic business value out of
AI and think about what that transition means for them and how it
could disrupt business.

We are spending time working with governments on the economic
implications of AI and focusing on investing in AI. We cover both
ends of the spectrum. What we try to do and this is what I mention
about bringing sustainability and innovation together, we try to
embed responsible AI. Thinking up front about impact, social
environments impact and all the advice that we give to our clients

when we do work in this area.

Coming back to AI for the earth, because we are talking about what is transformative solutions AI can offer, if we don't address the challenges we have no future. We don't have a future for humanity in the next century let's say. So AI for the earth is an absolute must, it's a must that we think about and we work with technologists and governments and industry and others to think about how we can apply AI for positive sustainability outcomes.

There are two key defining factors of AI that are helpful when we think about addressing earth challenges. The first one is AI delivers productivity gains. It optimizes our systems. Let's think about the need that we need to optimize our water system, optimize our energy system, our mobility systems, urban systems, industrial value trains, agricultural systems, and our use of raw materials. AI enables optimization and come back to the fact that in the west alone, we are at present, we are using ten times the amount of resources that the planet has, we carry on with the trajectory. There is a huge challenge in terms of optimization and improving productivity as we grow our populations further into the future.

The second key characteristic of AI which is super important for the earth is the characteristic of what it is doing for scientific discovery. It is speeding up scientific discovery. It is doing so in terms of new techniques, unsupervised learning, those of you familiar with reinforced learning, the infamous story of alpha goes

zero which has been talked about especially now where we have AI coming up with moves, humans haven't yet. In a way, accelerating the natural selection cycles of intelligence in humans, but using machines.

But it also, big data driven discovery is absolutely important for all sorts of things within the spheres of science. When we think about key earth challenges where we need breakthroughs in scientific discovery, climate impact monitoring, energy fusion, advance material generation, or AI used with biotechnology breakthroughs which will purify water or air or make more resilient crops, these are absolutely crucial breakthroughs.

I want to spend a couple minutes flagging an important initiative that we are working on that we kicked off at the U.N. General Assembly summit in New York last year, with the World Economic Forum and Stanford University we are looking to work with as many partners as possible across this field. It's called the fourth industrial revolution for the earth initiative. It is a multistakeholder initiative. We have three key aims. The first one is to develop a research base around fourth industrial revolution applications for the earth, and artificial intelligence is a core workstream for us in that, so building up a database of AI applications for the earth.

The second key area that we are focusing on is supporting a few break through projects in the space, so there is one on building a earth bank of codes which helps with biodiversity. There is another

one on ocean data platform which helps on transparency and around preserving ocean biodiversity. There is a number of these different project examples that we are supporting which bring together partnerships between big platform companies, research centers and governments and policymakers.

The third area of the programme that we are working on is building an accelerator. It's recognizing that we need more of a flow in the space, need more partnerships, need more projects and new ventures, specifically addressing earth, AI for earth related challenges. We are working on the design of that. The idea is the accelerator will provide that deal flow for interested venture for patent and other financiers and technical and commercial support to those organisations working on AI for the earth type applications.

Conscious of timing, but I did want to give a bit of a flavor for what this actually means, because I've been talking quite in the macro concepts. We released a report which looked at AI for earth opportunities. The real ones are out there today, existing in emerging applications. We collected 80 or so. We are probably now at over 200 applications, so if you are interested, have a look at that report.

We look to identify game changing AI for the earth applications, and we look to take those forward in new partnerships. To give you a flavor of some of those macro game changer applications, Wolfram already mentioned autonomous vehicles but more broadly, the smart transport systems, AI smart transport systems are fundamental to

addressing climate change and air pollution, and that is a variety
of things including the rollout of autonomous vehicles which enable
mobility on demand, but before that things like eco driving, improved
navigation in cities, platooning of cars, all which create
efficiencies in our transport systems. Precision agriculture was
mentioned, another big game changer which can produce inputs in terms
of water, fertilizers, pesticides, all of which harm biodiversity
as well.

An area close to my own heart, because I started my career doing
a PhD essentially and as a fellow using machine learning in the early
days, simple statistical machine learning but in high performance
computing is an emerging area which doesn't get much air time, called
climate infomatics. This is the use of AI within the climate and
modeling space. It is something like a infomatics lab, NASA is looking
at this in the United States and many other universities.

But the issue is the climate models are one of the highest users
of the super computers we have. We have a huge amount of
high-resolution data collected super high frequency, and if you think
we have that all the way very far back in terms of climate records,
they have complicated underlying physics. We can use AI, the use
of AI is not only helping to improve the accuracy of those models
to improve the understanding of more complex, more scale weather
events and climate extremes, but ultimately we are going to be able
to start using deep learning to make, once we calculate the model

with super computers, we will be able to do them on home computers which will improve scientific productivity.

Those are a few examples. To wrap up, a couple of leaving thoughts. The first one is to realize this opportunity, we do need collaboration. Is it not the climate scientists on their own who will get there, we need technologists and many others in any of these areas, it's experts working with policymakers and technologists and industry.

That goes back to education and how we design our education systems around these challenge areas as well.

The second one, we have to be purposeful about responsible AI, which is not going to happen by itself. We have to embed sustainability principles into the AI transition and revolution, whether you are a company or a country looking at your national strategy. But also you have to invest hard to realize the more game changing solutions that we mentioned.

I had a third one but I've now forgotten. I guess I have to leave that one. But I look forward to talking more.

>> ANJA KASPERSEN: You will have time to gather your thoughts.
   (applause).

Is there somebody called Abhishek in the audience? Can you ask the question, it's getting top votes.

>> Celine is the perfect person to respond to it initially.

>> I think so. Thanks for giving me the opportunity. My question was how small organisations or start-ups can basically think about

ethical considerations for their business when they are resource
constrained versus contrasting that with larger organisations or
companies that do have let's call them spare resources or capacities
to think about that, because when we are looking at really small
organisations with nonprofits or start-ups, they are very much driven
by having to work on whatever their mission is, and then for them
to be able to dedicate resources to this might be hard, so how do -- are
there tools, are there ways that they can do that? Think about the
ethical considerations for their business.

   >> CELINE HERWEIJER: There is a number of different ways to answer
that, maybe my fellow panelists want to comment. But it is difficult.
You focus on product market, on the growth that you need as your
start-up. One thing is that for certain start-ups when actually the
kind of underlying product or service we are developing is giving
rise to some positive impacts, you do have an opportunity to seek
or to attract impact finance by demonstrating that impact that you
have, so it's almost a new start of, a new source of more patient
capital, that if that's your purpose of your business, there is a
entry point to that finance which is more patient and often on better
terms which can be helpful.

   If that is not the case, and it's more around more broadly trying
to think up front around unintended consequences and social
environmental impacts, I think one thing to think about is that that's,
considering those things up front will add to your long term value

creation of your business, and showing that you understand that as

when you grow and speak to investors and the talent on your team

is going to be very important. I think that there will be a different

point in the lifecycle where it becomes important through the grace

stage of a start-up but it's always important to think about that

up front and obviously things change, because a lot of people when

you are designing things up front, what your product or service ends

up doing a five or ten years down the line, look at Facebook for

example or the Worldwide Web is different than it was conceived.

>> ANJA KASPERSEN: We can come back to that. I wanted to kick

off with this one. Don't be afraid. I'm not going to call out everyone.

Thank you for raising that in person. One of the other questions

been voted up is AI, the risk mitigation discussion, how AI can deepen

inequalities, you talked about this new initiative, using sidelight

imagery and doing this advanced modeling of the earth's resources,

and whether data can be used in terms of deepening equality and

integrity of data, how do you advise governments and industries on

this, what is your thoughts around are you worried about creating

these tools that fall into malicious hands or use. How do you protect

the data? Also how do we in your view and when you advise governments

and companies to use AI in such a way that it does not deepen inequalities

or cement certain biases.

>> CELINE HERWEIJER: Again unintended consequences, and trying

to think early about those. One of the easy to understand examples

in that space is there are new AI tools out there that get used by those working in the conservation field that help to understand where rhinos are because they are a endangered species so they can protect them and have controls around them but if you put that in the hands of a poacher you are doing a huge amount of damage. As we come up with interesting innovations, thinking about Cybersecurity and more broadly around data privacy issues and what happens with malicious use of the technologies is absolutely critical.

I think what you have to do is a case by case look at the solutions you are developing and try to do lots of scenarios around what some of those unintended consequences might be, but also be flexible in that some of the ones that come back to bite you will not be the ones that you thought about at the beginning so you constantly have to monitor the use of that which is why working in collaboration with NGOs and multistakeholder groups is important as part of that so you set up the right protocols in the beginning.

>> ANJA KASPERSEN: Thank you for that. We will come back to that question which is an important one as we proceed.

The last speaker on my list is Wendell Wallach, who is, both Celine and Terah were speaking about importance of interdisciplinary multistakeholder engagement, and he calls himself the chief silo buster and trying to make these fields come together. He has played a seminal role over the past 15 years in underscoring the societal ethical and governance challenges posed by the deployment of emerging

technologies particularly AI and biotechnologies and neuroscience.

He is a consultant and ethicist and scholar at Yale University and an advisor and he has professional affiliations, he has been working for a long time in the field of how to create agile and adapted governance. He published two significant books, one called the dangerous master, how to keep technology from slipping beyond the control which is appropriate title for discussion here, and also moral machines. Can we teach robots right from wrong, which also fits nicely with the discussions to be had here. Wendell, if you can share your insights with us at this stage.

>> WENDELL WALLACH: Last year's AI for Good Global Summit was truly an inspiring event. Yet at times it seemed a bit naive, as techno optimists wax poetic about what would happen once we conglomerated all the data, health and otherwise of everyone in Africa or Asia, we could reduce poverty if we sicced deep learning algorithms on that data.

And those particular ideas I think sent a shudder up and down those of us who have human rights backgrounds, or were cognizant of the rights of people to be used without informed consent in research.

This is not a new problem. It's been a problem in public health all along. When is it appropriate to trade off human rights for major societal benefits? Most of us have concluded looking back at history and man's inhumanity to mankind, that opportunities for doing that are very few and far between.

Now, a lot has gone down since that last year's meeting. We have witnessed the weaponnization of AI, use of AI to manipulate our behavior, suddenly bias in algorithms or the lack of transparency in algorithmic analysis has come to the fore. There is an array of issues around data ownership, data rights and responsibility of those who control data to those of us whose data they control. So we perhaps aren't quite as subservient to the technological idealith as we might have been last year.

I think that in and of itself is good news. It's exemplified in a maturation of the second gathering already. But there is more focus on ways we need to protect each other, there is more focus on what can go wrong. My primary question these days is how, as we are reaping the benefits of artificial intelligence and other emerging technologies, how are we going to also mitigate the risks and undesirable societal consequences.

I'd like to leave you with a simple distinction during this talk. I'm going to make a distinction between outwardly turning AI for good, and inwardly turning AI for good and perhaps accentuate the latter a bit more than it will be in other presentations at the summit.

What I mean by outwardly turning AI for good is this focus on discrete benefits that can be accrued with specific applications. That by far is and rightly should be the primary focus of this conference. Particularly in our ability to come up with applications that will mitigate the lives of the many billions among us who are truly

vulnerable in the way that most of us in this room are not.

I'd also like to suggest that in that area, perhaps some of the most powerful applications will be very simple ones, will not be these more complicated ideas that we are speculating about for the future. A few days ago I met Rose galinga, who founded a insurance company called Pula. Rose is a do-gooder who surprised herself when she founded this insurance company. What she does is provide insurance to 611,000 tiny farmers in Africa.

These aren't people who normally got insurance because just selling them the policy, let alone administering the policy when something went wrong, there was no business model for that. But what they came up with was that the policy itself would be packaged in each bag of seed and fertilizer purchased. And given the fact that cell phones are now ubiquitous, in fact, increasingly we have people with cell phones who do not have electricity, they have these tiny little solar jury rigged solar factories that charge their phones for them, they made it possible to register that insurance policy directly from the cell phone.

Now, if a drought should occur, again, they are not asking the people who registered those insurance policies to put a claim in. What they are doing is, they are looking at satellite data, particularly of cloud formations, and they are analyzing that with deep learning algorithms, and they are deducing how much rainfall there will be over different parts of Africa. If they deduce that there will not

be sufficient rainfall for crop production in particular areas, they immediately send a message by cell phone to those farmers that says that along with a certificate, that says they can go back and get another bag of seed or another bag of fertilizer.

A very small application, but one tremendously important when you think of the fact that actuarially, once every 15 years there is crop failures across these regions.

So this is one of my favorite examples of outwardly turning AI for good.

When I look at the Sustainable Development Goals, all 17 of them, I can see ways in which emerging technologies can help us meet every single one of them. I don't think technology alone is a solution to many of them, but it can certainly help us address them. And given the fact that artificial intelligence is touching nearly every facet of life, artificial intelligence itself may not be central to the solution of each of the Sustainable Development Goals, but it can probably help amplify the use of other technologies or bring sufficient, bring efficiencies to bear that might not be there otherwise.

But then I look at Sustainable Development Goal 8: Decent work. The promise of decent work for everybody. I look at Sustainable Development Goal 10, reduced inequality.

These are good examples of whether, where tech progress may actually exacerbate the problem. We have all been hearing about technological unemployment, that was John Maynard Canes' term for the long standing

fear that technology would rob more jobs than it creates. I'm among those who believe that we are already seeing a downward pressure on job creation, wage growth, because of automation.

Regardless of whether you think in the long run artificial intelligence will create more jobs than it robs, we are going to see major periods of disruption before that is realized. Reduced inequality, well, we have already, that already came up in one of the last questions, so I think we all understand that that's deeply problematic. In fact, so problematic, that economists are coming to understand that the tech conomy is designed so that productivity growth flows more and more towards us that are best off in the world, particularly the 1 percent at least as we like to metaphorically point to them but it touches all our lives in one form or another.

Then I look at goals 1 and 2, no poverty and zero hunger. Well, the lack of work and inequality, who do they touch the most? They always touch the same most vulnerable segment of our populations.

So, the problems with goals 8 and 10 are that it's going to exacerbate our inability to meet goals 1 and 2. We are confronted with this fundamental problem that in all areas, progress and artificial intelligence is not going to serve the sustainability goals.

And furthermore, there are ways in which this progress will dehumanize people, rob people of meaning in their life, particularly if that meaning comes from jobs that they can no longer find, and interfere with basic human flourishing.

So, here is where I come to inwardly turning AI for good. Inwardly turning AI for good is about mitigating the harms. Mitigating and attending to what can go wrong, and looking closely at the societal impacts of these technologies to be sure that we aren't taking our communities in directions that we don't want to evolve in, that don't have meaning for humans in the long run.

Now, those harms come in many forms. There are system failures that are due to bad design, that are due to incompetence, that are due to underestimating the likelihood of low probability events that can have a large impact. There's the appropriation of these technologies by rogue actors and by elites for self-serving purposes.

So, our charge is to look at both sides of AI for good, and attacking the second side particularly we need to look at the ethical and governance concerns.

Now, I don't really make a distinction between ethics and governance concerns, because I'm not using the government word, I'm using governance. I'm using the word of how we bring appropriate oversight to the technologies we are developing, and how we, in that oversight, look for gaps and ways to address those gaps, and what are the best mechanisms for addressing those gaps. So those mechanisms might be technological solutionism, they might be better corporate oversight, they might be in other forms.

Rather than go on, I wanted to leave you with that, but let me just leave you again with the final point, let's have a distinction

on both sides of what AI for good means, outwardly turning AI for good and inwardly turning AI for good, and if somebody else has got a better phrasing for making those distinctions, please let me know.

(applause).

>> ANJA KASPERSEN: While you were speaking, Wendell, you actually got a direct question raised to you from one of our participants, do you imagine a future where human robots and other types of AIs will have rights and responsibilities like paying taxes, in order to conform a new type of society merge with AI. If this doesn't fit within what you call the philosophy fiction, it's a good one.

>> WENDELL WALLACH: Do I have this turned on? I'm not sure. Yeah, good, okay.

I can imagine absolutely anything. I do not pretend to know what is and isn't possible. But I tend to get highly skeptical of some of these more fanciful features, at least skeptical about when or if they might occur.

I don't always find it helpful for us to indulge these fantasies, it may be helpful in the form of, yes, we want to begin research on whether we can control super intelligence, just in case we accidentally invent it, without being prepared for it. That I think is a meaningful agenda, but it's agenda for researchers.

So I don't find it particularly helpful that we talk about AI rights or the contributing taxes because I suspect that it's still a long ways off, we have pretty primitive technologies so far, and

I get deeply concerned that those topics distract us from looking at the real challenges that we have today.

>> ANJA KASPERSEN: One of the issues you raised in your presentation relates to one of the questions that has been voted up, which is we tend to overestimate the effect of technology in the short run and underestimate the effect in the long run. How do you see this applying to where we are in, as you said, in sort of the caught between the hype and the tech naivete and something in between?

>> WENDELL WALLACH: Another excellent question.

I don't think this is easy. I don't want us to be overly concerned with what goes wrong in a way that we obsessively cannot take advantage of these technologies.

And I think the basic point is true, we tend to underestimate up front what the impact can be and overestimate in the long term. But I know, I think that is no longer true when you listen to the techno optimists. If anything they are overestimating, overpromising, and hyping technologies that aren't fully realizable.

Where are the self-driving cars that we were promised that we could buy a few years ago? We still have technological problems that are in the way of whether we are going to have those in the next few years.

But we are in this universe where sometimes we are now overestimating the technologies that will be realized. But that said, in the long run, there are these impacts, yet I think we should be focusing more

on the impacts that are truly feasible, those that require attention, and not too quickly jump to impacts based on science and technology that we don't yet have.

>> ANJA KASPERSEN: Thank you for that. Let us now go back to just a dialogue based on some of the interventions that has been made through pigeonhole from the audience. Professor, one of the questions is around AI assistants and developments we are seeing and some say this may be the potential game changer to take deep learning into a new space because learning data set is us, our individual behaviors, and with the launch of Google's AI replicating a human, with the phone calls and we all read the news report and criticism against it, where do you see this technology stand and how do you feel about it, as a technologist, and at the intersection of the development of it?

>> WOLFRAM BURGARD: Technological, this is really hard problem, and I mean we all know, I find my navigation system in the car a little bit boring because it always sounds the same, and would love to have it a little more lively, to be honest. Maybe talking like us, yeah.

The same applies to other problems as well, if you think about the mixed traffic that we envision for the near future, where we will have some self driving cars and human driven cars, it's going to be extremely hard to programme these cars to behave properly in different areas in the world. Right now most of the companies are

in Mountain View where traffic is completely different from Paris or Palermo or New Delhi, so you can imagine in those places cars need to behave completely differently. It is going to be impossible to programme these cars and adjust them to all the cities they are going to be used in.

The only way around this is using learning and learning from other people, and see how people drive in the cities within the limits of what is allowed, which isn't always the case. But this is basically the way to go. We experienced, you mentioned this in our experiments with robots in public spaces, that people actually challenge robots.

So they are not always cooperative to robots. It's not only the case that robots might make mistakes but people are actually also somewhat noncooperative to robots. That might be fatal in the context of traffic, when people challenge cars, then it will be really really problematic.

That is something where we maybe need to develop cars that behave like human driven cars, in most of the cases, being safer at the same time. But that cannot be distinguished from human driven cars.

>> ANJA KASPERSEN: You are not worried about that we are moving into a air of anthro more Phissing these devices and washing out the lines between us and them, you don't see that coming soon.

>> WOLFRAM BURGARD: I do see that coming soon in some aspects, like interaction systems but also in the context of navigation, this is going to happen. People are also working on robots that perform

legible motions, so if you think about a human robot collaboration task, you do want the robot to move in a way that is understandable by humans and humans find appropriate.

>> ANJA KASPERSEN: Thank you for that. Terah, one question, many people ask themselves in the audience when you put out the slide of who is the partners and who does the partnership work with, is of course the sort of news that is on everyone's mind when we learned by the poor management of our data for those who are on the Facebook platform and also the poor control over third party applications.

Can you comment on that? The involvement with Facebook through the partnership initiative and also sort of more broadly how you view data integrity and how that is being addressed by the partnership initiative?

>> TERAH LYONS: Absolutely. The partnership really, I mentioned this before, but it was really created as a environment of critical self reflection, and the partnership, when it was created in 2016 did not anticipate the problems that it would be encountering. But it was predictive in the sense that a lot of these major industry players understood that the technology they were creating was one which had to be handled with care.

One thing that we are trying to do actively now is to create a environment of accountability, which allows for the whole spectrum of voices in our community to be heard. It's worth mentioning that 60 percent of our partners or nonprofit or civil society organisations

many of which are actively in opposition in some cases to work that technology companies do, which infringes upon data rights or data protection, which might be found problematic in certain application realms which might provide a host of other concerns in the public's eye associated with technology development which is nothing new, but with AI in particular, there have been some concerns incited for reasons that all of us have touched upon on the stage here today.

A lot of what we are trying to do is make sure that those honest conversations can be had, and that partners when they are in discussion with each other are making sure that those who are not technology companies' opinions are voiced on an equal level, and position of power equal to the tech companies and the tech interests represented as well.

I'm actually heartened by, I'm a skeptic, I'm a former regulator myself, never worked in the technology industry myself, but I've been a technologist, and I'm actually deeply heartened by the seriousness with which tech companies have taken a lot of the concerns and challenges associated with AI in the context of the work that the partnership is doing. I anticipate this issue will be one which we will handle in some form or fashion. We just started the work that I described kicked off a couple weeks ago in the form of our working groups in the presentation that I gave.

But as these conversations evolve, I'm sure that these topics will be considered in the context of our mission, which is to create

best practices and promulgate them, and that partners will equally take seriously the work that the community develops and hopefully ascribe to it themselves and implement it in meaningful ways.

>> ANJA KASPERSEN: Thank you for that. Celine, we spoke a little bit about the, one of the questions that have been raised is the skin in the game of the PwC in this, and your work, and I know you work across a lot of different fields. Where as a professional services company and offering advice do you see, I mean when I speak on these issues I often speak about the mismatch between the adoption rate of technology, not necessarily the scalability of the technology itself, and our maturity to manage it.

Do you feel that we are in a space where our adoption rate is matched by understanding what it is exactly that we take on board and also managing the impact on it, be it improving the supply chain or doing advanced climate modeling models to preserve earth's resources.

>> CELINE HERWEIJER: If you want a detailed answer, my colleague Rob will discuss it in detail in the session after the coffee break around responsible AI. Parts of that up front is assessing adoption maturity of the clients that we are working with both public and private sector and working from that context of maturity.

More broadly, what is the skin in the game for PwC, we do annual CEO survey to our clients that we always launch in Dallas every year and the last two years you see every, you see it out in bold very

clearly that tech disruption and in particular AI is front of mind with all of the CEOs right now. It's a board level issue, it is one everyone is going through now. It varies a lot across sectors. Look at the energy industry, there will be far more involved in the actual application, whereas other industries are early on in that journey.

But absolutely, as I mentioned, the starting point of not just thinking about digital and tech disruption and the impact of AI on disruption of opportunity for business but starting from the point at the beginning of making it a responsible transition is how we are trying to go about, and we are putting a lot of time and effort into developing work in collaboration with partnership, partners around things like explainable AI, audits for AI, ethics of AI, AI for the earth initiatives, a lot of work on the economic impacts of AI as well. We are trying to think about this holistically.

>> ANJA KASPERSEN: Thank you for that. Kenny Chen, can you ask your question, don't spend time elaborating, ask your question. It's a good one. I ask you to raise that to Wendell.

>> Hi, there is conversation about regulatory environments and standards and those kinds of things, but on a fundamental level, there is still very critically different ways that people are even talking about this domain. The lexicon, the goals, the priority areas, how do we reconcile essentially these language differences across sectors and industries.

>> WENDELL WALLACH: I feel like I planted you in the audience

or something (chuckles) in some cases we have translations across. We have bodies such as the IEEE and ISO that are creating standards that are recognized internationally. But it becomes more problematic in other areas. It even becomes problematic within our own societies.

But the greatest difficulties are coming up even with definitions that everybody understands what the essential issue is, we have that for example going on here at the U.N., in these discussions about lethal autonomous weapons, and how they should be managed, or whether there should be treaties or bans or so forth. The conversations have been circular because people aren't using the languages in the same way.

Part of this is again silo busting, breaking down the walls and talking to each other and seeing if we can understand each other's languages.

But the other difficulty is that the existing governance frameworks we have are totally maladapted for the speed of technological change. That's not good news. That's not good news in an environment where you are having the development of AI applications at such a rapid clip, perhaps one every day in a different domain and there is no time for legislators to make decisions about what should or should not have standards or be regulated.

So plus the fact that within the more Democratic countries, there is real reaction against regulation and overbureaucratization of scientific development in the deployment of beneficial technologies.

That means we are in an atmosphere where we need to be creative and come up with totally new approaches for the oversight of these technologies. First of all, you are going to have to get over the language barriers within your individual countries or regions, and then the next step becomes even harder, to internationalize that.

>> ANJA KASPERSEN: Thank you for that. Celine, going to jump back to you for a second. One of the questions is around, you were talking about how you are collecting a lot of data. One of the questions goes how do we collect problems and metrics data sets from AI research to track progress and it relates to one of the questions you raised in your presentation on the validation and testing of data and systems. I want to ask both of you to reply to this. I'll go to Celine first. But do you in addition to looking at applications, you mentioned 200 plus applications that you have found in your research, also look for what they are doing in medical research for a long time, what didn't work, why didn't it work, what was the problem and how do we get the metrics better to get the designs better.

>> CELINE HERWEIJER: I'll let the technologist answer the more technological aspect of the question.

But in terms of what we are collecting, yes, very much is on use cases, because use case is an experience can teach you a lot. In terms of those use cases, for what reasons were the motives behind the development, what stage of maturity are they, who are they focused on, what specific challenges, how are they financed, all that type

of data, also what are the availability of APIs and tools and algorithms that others can use, in a way to democratize those applications as well.

Absolutely, what we are trying to do is collate this database. What we focus on at the moment are earth related challenges and specific example, but you can extend it to the entire SDGs and that is something I know a number of actors are keen to do. There was a more technical question in there around explain ability potentially. I'll pass on to my colleague.

>> ANJA KASPERSEN: Professor.

>> WOLFRAM BURGARD: One aspect to measure progress is basically based on publicly available data sets, which is something that is extremely important and what the society is also doing right now, we are creating big databases with data sets labeled data sets, for example, where we can test different systems on and also measure progress, as something that the AI community has developed over the past, and also, there's been a lot of progress in how to measure progress and so on and so forth.

When it comes to explainability this is more complicated, because right now many of these deep learning architectures are actually black box systems, so it's hard to look into them. On the other hand we need to be aware of the fact that even other complex machine learning approaches are hard to explain, simply because you need to transform, for example, the data in a high dimensional space that cannot be

imagined by humans, in some sort of reasonable fashion.

So we do not actually understand what is going on there in high dimensional spaces. So you need to have to live from my point of view with the fact that systems might not be explainable in the future, as long as they serve the purpose and we can take the risk of potential faults, in the worst case we need to do a lot of statistics in order to figure out how risky those systems are, I think this is probably the way that we have to go, and in order to measure this, we need to have data sets, maybe simulation environments in which the systems can be evaluated in order to verify them.

>> ANJA KASPERSEN: More investments in the early, design phases.

>> WOLFRAM BURGARD: Exactly.

>> ANJA KASPERSEN: There is fundamentally a series of decisions made by humans so we have to translate into the validation of the systems. Mindful of time, I want to give Clements the chance, I don't know if it's his or her, looks like it's a him, your question, and I'm going to ask all the panelists to then all reply to your question as their final round before we end this session. You can go ahead, it's on education.

>> Thanks a lot for the opportunity. My question was about if you look to artificial intelligence and the impact it has done and does in the future, which are good or bad, what kind of education do we need to provide and how do we probably need to adapt our educational systems today to prepare students and scholars for the future? Thanks

a lot.

   >> ANJA KASPERSEN: Thank you for that question. I thought that
was a good one to end with. How do we build literacy, how do we bridge
gaps due to silo busting, make it multidisciplinary. What is your
advice? Start now with you here.

   >> There needs to be a substantial investment into the education
system, universities at all levels, maybe even at high schools. We
need to start educating students, in order particularly in Europe,
in order to create these needed amount of AI engineers in all different
fields. That will make us ready for future developments and for the
need of AI technology that we are going to expect.

   >> ANJA KASPERSEN: Terah.

   >> TERAH LYONS: I agree with the point about investment in education.
I also think that a lot of the solution has to do with the nature
of education as well, and the notion of ethics by design has become
increasingly popularized especially in the technical field in recent
years. I read a statistic as recent as last week, that over 80 computer
science programmes globally now have an ethics course that approaches
either computer science ethics or AI ethics specifically, which is
incredible.

   But more needs to be done to integrate considerations for unintended
consequences and some of the other concerns that we have talked about
today into curriculum, and I'll let Wendell make the point as the
silo buster, but also to make sure that education is inherently multi

disciplinary and interdisciplinary as well.

    >> ANJA KASPERSEN: Wendell.

    >> WENDELL WALLACH: Well, without repeating what has been said
so far, let's just presume that, but I think it goes beyond education
and it's not just for students. It's also for all of us, that we
are not having a public conversation that is effective and includes
everybody about the world technology is creating for us, and whether
we even want that world. So we really need to find public ways of
engagement where everyone talks about what the trade-offs are, what
norms we really want to put in place, and whether we are headed toward
futures that they want to bequeath their children and grandchildren,
and that is not happening.

    So my main proposal for education, I would like to see the problem
of the week, you know, the problem of the week is, are we willing
to make the necessary trade-offs to speed up the adoption of
self-driving cars even if it means that sometimes they will kill
people, that an attentive driver would not. It is a complicated issue,
but it's an issue that we need to have a conversation about.

    Those issues about, do we want to go to a transhumannist future,
do we want to have massive amounts of jobs replaced just because
they can be automated, and we get efficiencies from them, and if
we do, then what are we going to do for those for whom those jobs
are robbed.

    We are beginning to have those conversations, but only beginning,

and I think rather than education, we are still in a world of disinformation and sometimes obfuscation.

>> ANJA KASPERSEN: Thank you. Celine.

>> As a mother of a 3-year-old, it's on my mind, apparently at 5 they already start cutting classes, so the days of technologists are on their way. Number one, kind of the future, making sure we have enough of the data scientists that we need as AI becomes ever more widespread tech, right. Number two, intir disciplinary, whatever domain you focus on and you become an expert and you have the ability to think about the pair of big data and AI more broadly, it is mainstreaming some level of familiarity with AI and as part of the mainstreaming on the technology side there is also mainstreaming consideration of the ethics side of the values, of the governance, as part of that educational process.

This is starting at school age and all the way up into university education, and the final bit about it is let's not forget about the fact that jobs are going to change, they are already changing and they are going to change quite rapidly in the next ten, 15 plus years and there is a huge amount, reeducation we need to do of our older workforce and those are going to be in the more automated industries who are going to be hit hard, so how we re-educate those as well.

>> ANJA KASPERSEN: Thank you so much for the comments, and just as a final before we conclude the session, I think what we learned from this session is that the biggest transformation is us, how we

adopt, how we evolve, how we govern, how we edge kailt ourselves.
I like to say we have -- educate ourselves, I like to say we empower
to engage and we engage to empower, and this is a big issue in the
AI space. There is hype, there is people making knowledge about this
issue into their own business model, that enables or somehow
discourages others to interact with it, to fully engage with it.
That is one bridge we have to make sure is actually built, how do
you enable people to access information about this issue to fully
engage and control the transformative impacts on our life, which
is why it's great that this summit is happening.

Back to the SDGs which is the framing for why we are here and
we talk about education as a final question, and every one of the
presenters touched upon, how do we make sure that innovation, our
drive to transform and optimize and do better and read data better
is responsible, SDG9 talks about growth and innovation, how do we
make sure that that growth and innovation is not just sustainable
but also responsible.

How do we make sure it's informed, that it has a testing and
validation and investments in those environments that we need it
to have. 17 around partnership with Terah talking about how we make
sure we have the right partnerships. Doesn't have to be the big
industries but they have to be involved, because they are the key
drivers of the enabling technologies that impact on the rest of us.
Education, ethics, how do we make sure these are not electives but

embedded into the courses and trainings as you were speaking about on education. And strong institutions, and I know PwC for example advises governments how do you build institutions, build systems, be it from the minimal thing to macrothing, but how do you build strong institutions that can cater to making sure the alignment between SDGs and transformative impacts on AI actually reaps benefits.

That was my sort of poor effort of trying to summarize what we heard today. But I would like to encourage all of you to, I say thank you to all of those that posed excellent questions. I rephrased some of them to make them flow but we addressed some of the top ones. I thank the presenters for really sharing their insights and for great presentations and Q and A session.

>> Thank you very much for the panel. This was amazing. Before you go to break, some notes of logistics and information. As you heard this morning, we are presenting the sessions tomorrow, tomorrow we have four tracks and wish you are going to be participating in new projects that will be expressed to you and you are going to be participating in building those. In preparation for those tracks tomorrow, right now after break you will have a choice of four different, five different panels in the programme, that are consisting of show and tell. Examples of AI for good that have been implemented by different agencies, UNESCO and others and the list goes on for that. In this room we have innovative applications of AI in data science in the field, popoff room, in CI, application of AI in education

and health, in room C2, which is next door as well, safe and secure AI, all the applications and examples. In room K, where you got your badge, data for good. And at the same building you have socially inclusive AI. All these five tracks are mostly show and tell to give you a sense of what projects and innovative ways of AI for good have been demonstrated, explained and what obstacles they have. That will give you good scoping and good preparation for tomorrow. Thank you very much. Enjoy your break.

(end of session at 1536)