RAW COPY

AI FOR GOOD GLOBAL SUMMIT
BREAKTHROUGH PITCHES
MAY 15, 2018

Services Provided By:
Caption First, Inc.
P.O. Box 3066
Monument, CO 80132
1-877-825-5234
+001-719-482-9835
www.captionfirst.com

***

***

>> Coming to the last part of the official programme, there's still a few seats here, you don't have to stand in the back. There are seats available here.You don't have to stand.

So we come now to the pitching session. The heart of the event, our four tracks that will run in parallel tomorrow, so we have a track on AI, satellite imagery, an, track on AI and health, smart cities and we have a track on trust in AI. Those four tracks run in parallel from nine to 6:00 or so, 9:30 to can, maybe a bit longer tomorrow. They're competing with each other for audience. Now we have each of the leads coming up and presenting the session. I had mentioned this morning that a lot of work went into that, curating the tracks, I thank all those team leads and the colleagues who helped prepare the track for all the effort that went into that.

We just go one by one, you come up, one after and then we finish one and then the next one comes up.

The first one --

>> Stuart Russell is the first lead and he'll explain it and others will explain their own track, what is important to remember, is that it did take two months to think about, to gather context and projects to explain to you so tomorrow when

you join this track, you have a chance to make an impact by choosing projects that will be impactful for AI for good. The presentation tonight, it is an opportunity for you to get a sense of what tomorrow may be in this track, what the projects will address and what the challenges that they'll try to take on and whether you are going to be the person that could help them.

Without further ado, Stewart, you're the first one.

>> STUART RUSSELL: Thank you. I should say I'm lead in name only. The work was done by the other people on the list here that you can see. These are the people that know about at satellites. I just know that using satellites is a good idea.

When you think about AI for good, right, obviously when people do good in the world, they're using the intelligence to do that. It goes without saying, if we had more intelligence we could do more good, we could also do as Michael Moller pointed out, more bad. AI can really expand both the scope and scale of what you can do with your mind and it can do that quite dramatically. We saw that with an amazing worldwide effort that involved 20,000 research teams from 150 countries coming together to solve a really important problem for the human race which is which movies should you watch on Netflix. You have to ask yourself, okay, we achieved an amazing global coordinated collaboration to solve that problem, but is that really the pinnacle of human ambitious to improve movie rating predictions? I think not. I think we can do better than that and if we put the same effort into solving problems that really matter, then we will really make a difference.

Let's look at satellite data. Right now there are hundreds of satellites sending image data down to earth and roughly speaking, these are ballpark figures, every day we imaged the entire earth to approximately 50-centimeter resolution and 50-centimeter, it is small enough to see individuals, individuals, you name it, you see it, that's 2 times 10 to the power of 15 per pixels a day.

Coulding down to earth. 2 billion -- equivalent to 2 it billion one megapixel photographs. That sounds like an incredible amount. It is actually not that different from the amount of data your eyes process in a single day. What's the difference? The difference is that the world is far too complicated and big for your eyes to understand. If we could take that satellite feed and make you watch a movie of it, you see the whole world that that resolution in a single day, you wouldn't be able to make sense of it because it is too much. There is too much variety, there is too much complexity, there is just too much and the brain cannot possibly comprehend it. With AI, perhaps you can.

So if you think this is something you can do on your

laptop, it comes to 100 million a year just to buy the disks to put all of that data into so it is not trivial but it is also not -- these are.

Ins that can be easily accommodated in the budgets of major corporations and countries. For a long time we have recorded the whole world. We have had satellites and they're sending down data and we can store it. Recording is not the point. it is seeing. That's the point. not recording busying. AI, it is the only thing that can let us see the whole world at once. Really what we're talking about is turning -- creating a sensory capability for the human lions to see the whole world at once and turn it into a global realtime database of everything that matters that you care about. This is not oh, let me see if I can do a project where I get some satellite data for some particular city and I have a bunch of people that write a bunch of code, then we can track the movement of camels around the city or something like that. That's not what we're talking about. We're talking about creating database where all of that information would already be available. You wouldn't start a big project and write software, you just happened to the database. You can do pretty much anything you want with this kind of information.

The world is incredibly complicated. The challenges for AI here are enormous. Right. For a long time people have done things like trying to classify pixels in satellite data to say, you know, is it forests, agriculture, city, is it water. Right it. That's fine. Just think about agriculture. You think of agriculture as it is just fields. How complicated can that be. Here are fields in Florida. They have a long, skinny shape. Here is some fields in China, these are rice Patties and they're completely different. Here are fields in Spain, they're under glass. Here are some fields in Saudi Arabia, they have a very funny circular shape. Here are some fields in Indonesia that happen to be underwater. This is a sea we'd farm in Indonesia.

An AI system that's trying to keep track of what's happening in agriculture, it needs to understand all of these things and another thousands varieties of things that agriculture happens in the world and these are not straightforward classification problems that we already have solutions for. There's lots of interesting AI work to do.

If we can do this, right, there is a number of roles for this kind of data in managing the world and solving the SDGs. The simplest one, which people have used for a long time, it is visualization, to see for the first time things happening on the scale, oen a global scale over long periods of time with climate, with forest fires, with vegetation movement, so on. We can also do state estimates, if you're trying to build a system

that manages the world, makes progress on an SDG, you want to know if you're making progress and we can use satellite data to do that and provide control feedback. If things are going in the wrong direction, you can change your plan and move in the right direction. You can also look at past events. When something happens, you can look at the satellite records, so just like people do when a crime happens and they look at the closed circuit television, they see who was there when some major global event happens, you can look at the satellite record and try to understand what was going on. Similarly, if you understand how the system works, you can use the data to predict the future. Learning models of the earth's system using satellite data, which is a much vaster source of data than we have ever had access to before is possible.

So just a simple example, this is from UNICEF, the surveillance of a refugee camp in Jordan at its peek was 150,000 people living in a small area. I think it is 5.3 square a kilometers. It is a dense population area in the world at its peak. You can look at time-lapse photography and see how the camp filled in overtime and people put their dwellings and gradually made it denser and denser and denser. The U.N. actually counted all of these buildings but the AI systems surprisingly were not up to scratch. They had to count these by hand, one-by-one, on very grainy -- temperature is a high resolution image of the same thing. If AI is able to solve Netflix challenge, then AI should be able to do this. Please get our act together, do things as simple as counting houses.

So they have started on that. There are a number of collaborations already out there. Digital globe, a big provider of satellite data is working with several other companies that do AI like Facebook, there is one on counting houses and ground truth is important. They have ground truth it had all of the area and then train the vision algorithms to recognize houses. Last time I checked, they were only able to get 50% accuracy in finding houses. That's apparently not good enough to do, you know, real population estimates and Examples. You can look at changes overtime. This is a map of South Africa showing the population density in different parts of the city which was done using an AI system applying it to the satellite data. You can do things on scale, but accuracy is still lacking. Just to give you an example of why it is not trivial. Hue other are some houses in Madrid as you see with the street in between. Fortunately, these are not houses, but they're tombs, they only contain dead people. It is not completely obvious how an AI system is supposed to figure that out. It does have to if it is going to be accurate in encounting living people. Our track is going to -- we'll introduce the challenges and the

opportunities, talk about what satellite data is, how you use it, and then we have three sessions dealing with three particular tasks, poverty, deforestation, the affect on the environment, and then agriculture. We're also open to other suggestions and we really want to have a lot of brainstorming during the day and we're trying to actually result in real ongoing projects that produce general purpose tools that all of you who my have any intersection with the areas would be able to use to solve your problems. So this is not going to be just a day of talking, it will actually be resulting in doing and some projects with money, data, resources and hopefully the contribution of the AI community to solving the problems. I also wanted to mention that there is a more general effort that others are putting together called an AI and data commons, it is an environment in which as solutions -- as technological solutions are developed, we put them together with the people who own the problems, the people owning the data with resources, training, in particular to solve some of the legal problems that inevitably come up when you try to do real work out in the field and make it easy for the solutions to get out there into the world and to actually solve the problem quickly. That's all I have. Thank you very much.

>> Thank you very much.

Next -- is it just you? Marcel Salathe, professor at EPFL. The team comprises of Marcel Salathe and Rameth Krishnamurthy and Sameer Pujari from WHO. It is Marcel Salathe presenting AI and health. AI and health 101, it is a primer, a two-pager if you go to the website, it -- you will see that two pages written by Marcel Salathe on AI and health.

>> MARCEL SALATHE: I would like to-ing thank most of the speakers of today. They have actually already done a lot of the pitching for us. You have heard that AI and health is sort of a pretty natural map I'm going to I have it again. It is an exciting day ahead of us, I thank our colleagues from the WHO, Rameth Krishnamurthy and Sameer Pujari who helped put together this programme and we're excited. Health and AI are natural companions. They are two of the most exciting fields converging if you think about it. The health sector, it is in many countries the largest economic sector or the fastest growing, if not the largest. At the same time we have AI which we have talked about as we heard this morning since about 50 or 60 years and we have now in recent years reached a new level of maturity for many different reasons and these two fields now intersect and as you can imagine, there is a ton of energy that's being released. We're excited about this also, we're here after all also for the SDGs, health, good health, well-being is obviously a Sustainable Development goal. It ties very naturallier into

all of the other Sustainable Development Goals. All of them affect health and well-being and vice versa without health and well-being, all the others could be potentially quite challenging to achieve.

We had a good pitch this morning I think by the Director General of the WHO who also tweeted and said Artificial Intelligence plays a key role for health, something we obviously ahad agree with, and that it is already playing a big role. The role is only going to increase. What are we going to talk about? Health is so multidimensional connecting to so many different fields and in fact almost all of the other sessions as well so we'll focus a bit on technology, obviously, we're talking about AI at the end of the day and we're talking about trust, because that's going to be essential for the uptake of these kinds of technologies and of course, we'll talk about policy issues as well. Technology I think we need to understand may be a bit deeper, the end user of health applications at the end of the day don't care I think deeply about the technology, but we as technologists, as medical experts, as policymakers, we need to understand it a little bit more. Only if we understand the technology do we know also its pitfalls. We'll focus also a lot on that.

I show you three examples. These are just three out of many so on the top, you see the skeleton that was talked about earlier, we see, you know, very scalable technology, for example, the skin disease recognition that if it works accurately can be deployed on 6 billion smartphones we expect by 2021. Talk about impact and scale. We'll talk about medical chat bots and this is broader domains, visual understanding, language understanding and interpretation of complex biological, social, environmental data, this is very broad because health is very broad. Think about dermatology, think about radiologist, any other medical image analysis and AI is playing a role and will play a central role of autonomous vehicles as we have heard could potentially be one of the major Public Health breakthroughs in the next 10 # to 20 years. Language understanding, understanding the vast body of EHRs, understanding the entirety of medical literature, what it means if you can suddenly have a bot on your phone that has read the entire medical literature and keeps reading the entire medical literature in realtime. These are the kind of questions we'll talk about.

Robotics, genetic data, behavior data, it is unlimited. We should also talk about the other side, I think it is very important, again, we're talking about the major economic area, so let's not kid ourself, there's a lot of different interests and we know from a historical perspective that if the incentives

are not properly aligned, systems, they tend to not go foo the directions that we ideally would like to. How do we actually create the right incentive structures? What's that mean? What about open medical data, what about Open Source algorithms, is that possible? Is it easy to do? Who wants the medical data out in the open for AI for good. There are a lot of tricky questions and we have to think innovate actively about new models, cooperatives and citizen science and so on.

Also, what are AI for good models that are financially sustainable? We'll also talk about that, that's important in the health context I think.

Finally,, we want to talk about barriers to pimp men takings, I think most of us share the excitement around AI and health and most of us, we're also painfully aware we can't go out and do something because there will be many barriers we have to overcome, how do we do this collectively.

The last point I want to mention, obviously also talking about the ethical issues and in particular, about issues that are completely new in AI. Modern AI is by and large deep learning based, not all of it, but by and large, these are neural networks, and these are essentially buy and large still black boxes. So if a black box gives you an accurate medical diagnosis, what are you going to do with that? How do you hold them accountable, how do you explain it to the patient and so on. What's the data that these algorithms have been trained on, have they, foreign, in the skin cancer example trained on tick later data that were mostly, you know, white skin, textural corporate, we understand that there is a lot of sexism, is racism in these areas that were generated in the past 100 years. If we now develop AI that learns on these datasets, will they be biased as well? The answer is yes. These are serious problems that we have to address. We'll do all of that, of course, we cannot dot this in one day but we hope to have selected a good group of people that will discuss this with you, that will also show project pitches that you can eventually buy in and support the way you can. We're looking forward to seeing you tomorrow.

Thank you very much.

>> Thank you very much. We come to track number 3, both of you, smart cities, we have Renato de Castro here and also Alexandre Cadain.

>> RENATO de CASTRO: We'll do your style. Yeah. Good afternoon. We'll go here.

I was listening to the pitches before and it is exciting to talk about AI that I was almost giving up -- my track to work with Marcel Salathe. Actually everyone is here, everyone that's here, they really believe and we're trying to do something for good. Last year I got to 100 global events, smart city global,

smart city, it is a popular topic now, in Barcelona, we got almost 20,000 people from 120 countries to discuss about smart cities last November and SDGs were exactly the main point. before talking about our track tomorrow, I would like to frame better what we define for smart cities and when we're invited to do this together, we have two expertise, the smart city, the AI part, the first thing we decided, it was to frame the problems and opportunities in the smart cities and AI. I think we have the presentation, do you have -- please.

Okay. The first issue, we have a immigration of 75,000 people. This is, of course, making worse our uber problems, and if you say a smaller number, analyzing, in some cities, we get in 201485 people per hour. So, of course, this is the biggest challenges that we have in our cities now, and it is driving us to all of the problems in health, in mobility, in education and so on.

>> Thank you. Again, this is one of the examples of the challenges we face in the area of cities and the past years myself, I was lucky to help organizations leverage AI to impacket society for good. One of the opportunities we have today and these days to discuss the global opportunity of AI, it is obviously regarding the city, the amounts that are there that we have access to. This is just reminding us of the explosion that we have towards IoT and the fact that if we remember 10 years ago in 2008 was the date in which basically we humans were exceeding in numbers from the numbers of objects connected to the Internet and we expect now that by 2020 we'll have 50 billion connected objects, IoTs leading the way to create some kind of sandbox in which we may be able to use AI to identify specific problems within the city and also participate in helping us to design and assist us in designing solutions towards this. This is the reason why we wanted to work together on this idea of confronting those challenges with this new opportunity that we have, all of this, all of us together here, and the reason why we want to do this, it is to try to redefine the area of smart city, just asking ourselves what do we want to build as a smarter city tomorrow.

>> It is exactly the point, to frame what we're going to be tomorrow discussing or at least what are the main tracks in the smart cities that we're going to be addressing. I had several meetings today and the first question is how do you define, how do you define smart cities? It is very difficult to have just one definition, but I think there are five main elements that should be embedded in the smart city pitch nowadays. Smart city is not new, we have already talked about smart cities in the states, but the first one, smart city, it was and is and it will be forever about technology. The big difference of technology

is not more, not anymore the end, but the big driver. Most part of the smart city projects that I saw that are successful now, they're citizen-oriented. Then, the finer result that they're trying to seek with smart city projects, it is to have a better quality of life. These are the three main principles behind successful smart cities nowadays. What changed it from the 80s until now, it is that we have shown new competence. We have to understand that our societies now is influenced by what we call new economies. We have creative economy, we have sharing economy, we have a very strong behavior called cocreation that citizens and people want to cocreate in all of the process. When we talk about immigration as one of the key problems in smart cities, we must understand that there are two sides of immigration, cities getting a lot of people every day and cities that are losing people every day. When it comes to India, smart city, mission, national project, they talk much more about smart villages than cities, they need to avoid 300 million people to immigrate from the village to the cities. So Brazilians, they should be embedded in our talk about smart cities, this is exactly here with all of these data that we have, that our AI can play an important role in bringing resilience to our cities.

In the design of the session, we wants to express a point about community, it is a session called smart cities and communities and we want to see how AI can help us in both sides. We want to remind this kind of quote from Shakespeare who said what is the city, but the people. We want to make sure that we're going to focus on the idea that people, citizens, can actually be the first source to help us identify problems to be a part of this bottom-up approach we believe in, and during the session, we'll have in the different sub tracks that we designed different examples, different subjects related to violence and abuse, but also subjects and projects trying to help homeless people empower themselves through AI towards Entrepreneurship for instance. For every project, we'll have the same kind of design and presentation that was starting with the global issue within the smart city that they are tackling, how they're designing a solution with AI and then we'll try and see how the design can basically impact society way beyond the SDG11 that is about cities, we discovered that all of this project basically has an impact on all different SDGs we know. Then we'll try to see altogether how we can collaborate, which is the SDG17 to make sure that this project can develop better, maybe scale in a better way, duplicating other cities. We have here really the glass with half water. The issue about social inclusion, digital inclusion, it is very strong nowadays. We saw the picture in the morning with the farmer with the cellphone and a very, very -- let's say basic condition. I visit some of these

in my hometown with a group from Boston two year years ago and the thing they got most impressed with is the people the poor people, the children, they were in the very, very basic conditions, and the government didn't reach them with basic issues but they had already mobiles and cellphones. These guys, they're socially excluded but we're totally digitally included. We are using AI, this is a very good reference for us, because we can track them. We know where they are, what they're doing. If you start finding the patterns, you can really scale faster ways to solve their problems. We have decided to divide our it day in four tracks. For these four tracks, we got to invite 16 speakers from 1 # 1 countries, countries covering four continents. We have people from Brazil to Japan, South Africa, London, we really tried to -- it is really the issue about multistakeholders and trying to incorporate together for this.

>> To prepare you for the day of tomorrow, we have four different panels and the whole purpose of the day, again, it will be to discuss the different subjects, then go with really precise and practical projects that will be presented to you so that we can build on it and cocreate and imagine further. The whole idea is to see between all of the projects what needs we see in common in order to work towards the data and comments for the third day. The first panel, it will be about discussing the AI in cities to make sure that within the group with he share the common background of where is the research with AI and the intersection with the cities.

Al.

>> For panel two, with he selected some global projects from different areas to go deeper how solutions now are being applied in already established projects and projects to come.

>> And C, AI fostering smart government, the idea was to see how AI could be a powerful tool to assist government officials to not only react faster but to predict challenges.

>> The last one, exactly the main ideas of the whole event, be how AI will empower people to be smart citizens going from education to the issue of homeless people in order to use technology to really try to make something for good. We want all of you to know that we're looking for an interdisciplinary group to make sure we can alling with at no particular time, we're talking about cities and real different knowledge are necessary to build this and we're lucky to have you all here and hope to see you tomorrow. Arrive soon, because we have fewer space.

Thank you very much. .

Last but not least, the last presentation, it is started by Stephen Cave. This track is slightly different. It is called trust in AI and is more horizontal and covers a number of

different fields. Up to you.

>> STEPHEN CAVE: Thank you. Wonderful to be back in Geneva. Wonderful to be back in the Summit and see how it has developed in the last year and it is an honor to be asked to run this important track on trust in AI. This is our team, five of us are going to tell you a bit about what we have planned tomorrow now. I get to start. Why is trust in AI important? Well, let me give you a few examples. Imagine a major technology company developing an app and this app is for farmers in regions that can evaluate how much moisture a farm gets and direct their irrigation system to make sure that every square meter gets exactly the water it needs and no more. They don't use it. It could save money, save the most precious resource in the region and they don't use it. Why not? They have never met the developers, don't know who they are, they don't know what's being done with their data, they don't know who is making money out of this app or how so they don't use it. Let me give you another example. Imagine two neighboring countries, developing AI strategy and the first one says okay, we want to be a world leader in AI by 2030 and the other says we want to be the world leader in AI by 2025. The other stops collaboration with the other country, they don'ts want them to get ahead and the other country announces new break threws and the first country is terrified they'll be left behind. So it invests all of the more in developing the technology as fast as possible, closing down the ethic programs, stops the AI for good programs and so on.

A third example, like those we just heard on the AI and health track, imagine a medical AI system that diagnoses a certain kind of skin cancer with 95% accuracy and doctors only manage 80%. So deploying this saves thousands of lives, the system uses a form of machine learning that's opaque, not interpretable, not transparent. What doctors don't understand is why it makes the diagnosis it does. Then the doctorses see it making a mistake that they never would have made. They don't know why it is making this mistake, they don'ts know when it will make that mistake again. So the confidence in the system collapses. Now what each of those three examples has in common, it is trust and the breakdown of trust. In the first case, the lack of trust between the company developing the technology and the stakeholder community that will be effected. In the second case, there's a lack of trust between two countries driving forward to develop that technology. In the third case, there is a lack of trust in the technology itself. In each of those cases real opportunities to use AI for good are lost. Now we believe, as I'm sure all of you do, that AI has enormous potential for good and has enormous potential to accelerate the rate of which we can achieve the Sustainable Development Goals

and it would only do so if firstly those developing the technology can earn the trust of the stakeholder communities that would be effected by it. Secondly, only if trust is built across boundaries or between the companies and the countries who are developing the technology. Thirdly, only if the AI systems themselves are demonstrable trustworthy. We together with our colleagues have been thinking a good deal about trust and trustworthiness in AI systems in recent years and so we were delighted to be asked to come and run this track on trust in AI. We were particularly delighted to be asked to this forum. I think this particular forum, the U.N. AI for good global Summit, it is uniquely placed to address this challenge. We're all aware that historically the development and deployment of this technology and the debates around this technology have been dominated by one particular geograph cam group, an ethnic group, Agenda, starting to change, other regions are investing in the technology with great success and there is a huge amount of work to be done to build for trust between the many, many different communities who would be affected by the technology so that we can use it for the greater good. For example, so much work still to be done in addressing the enormous gender imbalance in AI, so much work still to be done in reaching marginalized communities and so much work to be done in addressing the structural interqualities that are themselves the legacies of historical injustices that continue to shape people's lives between regions and within countries. A project we may call decolonizing AI. We're not going to achieve this in the next few days a the this Summit and we can make a start. We can leverage all of the excellence work that's already going on. Now, I'm fortunate that we have worked with a fantastic team, these people here in recent months to develop this track and tomorrow we want to tell you about nine concrete practical projects that divide up into the three themes that I mentioned, the three kinds of trust. We want your support and your ideas. Please come and join us and help us tomorrow and to whet your appetite a little, I'm now going to invite on the stage the three themed leaders that will tell you a bit more about those three kinds of trust and the projects we have in mind.

I am a researcher for the future of the leader of intelligence and a theme leader for the project. The first theme that we'll address focuses on building trust between developers and various stakeholder communities if AI is to be applied for good it must be applied to improve the well-being of all communities and given the current trend in the development of AI we recognize that there are groups that have been historically excluded. We -- it is these communities that we are centering in our work.

The the trust will be crucial to implement the implementationings and practicing how one builds trust, it will vary widely from community to community. Different factors, such as age, educational background, cultural variations will be relevant to building trust around AI technology. The projects in this theme will start to map out some of this complicated territory and by exploring methods of building trust and also identifying where there are potential benefits of AI. The first of these three projects, they'll work to establish trusted networks between AI applications and mental health patients by first understanding where trust has been broken. The second project will look at engaging with a community of farmers in East Africa using AI to help increase food production for that region and the third project, it will work towards mitigating the affects ever AI in different countries and transition economies. For each project we have collaborated with champions and representatives from diverse communities so that we can work on building trust in AI.

>> I'm the theme manager on building trust across boundaries. So the general aim of the theme, it is to explore practical ways to build trust and collaborations across culture, nations, corporate boundaries on issues relevant to the use ofs AI for good, we're proposing three concrete projects under the theme. The first is cross culture study with different stakeholders involved in building trust for beneficial AI. For instance, relations between AI technology itself and the potential users of the technology and the technologist and the policymakers and for instance also researchers and collars in humanities and social scientists. Different people in the development of AI, we must understand each other well and work together on issues concerning the impact and the challenges of AI so that the aim of this particular project is to identify and build an international network of experts and organizations so as to develop a roadmap, so to speak, on discussions on the dimensions of trust in a cross culture context. The next project,ish it is interesting, it is called a global narrative, so not surprisingly, different cultures see AI through different lenses, a different philosophical culture, traditions and have their own different conceptions of AI and to build trust we must understand the different AI the way how they're being perceived. So the aim of this second project, it is to explore how narrative traditions in different parts of the world both has hopes and fears for AI and how the influence, the development, the implementation of this technology is.

The third project, it is a case study of AI technology and it is a regulatory strategy. In this case we're focusing on autonomous vehicle in particular. The aim of this project, it

is to explore, again, in the cross culture context what are the different regulatory systems and practices that are involved in different countries and what kind of enter developer trust that's required in order to successfully produce autonomous vehicle across different regulatory systems.

You're welcome to join us tomorrow. Thank you.

>> Thank you. I work at the royal society, the U.K.'s national academy of sciences. I have been developing with colleagues the third theme in our session tomorrow which is about trustworthy AI systems. So our first two themes look at you who we build relationships across boundaries of different sorts and this third one is taking a slightly different tact and is looking at the capabilities of the AI systems with he develop and the context in which we deploy them.

This is because it is a fascinating set of projects for what happened in the interface of science and policy where our scientific capabilities shape the responses we require and vice versa, so in this third theme we're asking how do we create AI systems which are not demonstrable trustworthy. We have three interesting projects that are showcasing, firstly, on AI and policymaking, we'll be asking what information to do policymakers need to make judgments on how reliable an AI system is. How can we start a dialogue between policymakers and technologists that can create this common language. Our second project, it will be looking at the accuracy of the data on which -- from which we develop AI systems asking how we can manage bias in the systems and whether there is a role for new ways of managing data and new data repositoriys in eliminating bias from the AI we create. A third project is looking at the values that underpin the AI we develop. We'll look at concepts like fairness, exploring how those play out for different communities. Across these had projects we want to develop a better understanding of how and when we should hey ply judgment in using AI systems. So we understand when we should or shouldn't put them to use. In doing so, we should also develop better AI and more applications for social good. We'll leave it there and hand it over to you for the bigger picture for the theme.

>> FRANCESCA ROSSI: I work at IBM research and it was really a pleasure to work with this wonderful group of people to put together this track on trust in AI. I think we are convinced that AI needs to be trustworthy and we need to build this system of trust because otherwise AI, we all agree that it will be very, very beneficial and we cannot without the system of trust really get all of the advantages and exploit all of the beneficial potential of AI and this is really -- you're seeing that in the wonderful example that was given at the beginning,

the example, we tried to examine trust from the three different angles and they should very well -- without trust, there is a risk of not being able to get all the benefits of AI. We decided -- that's why we decided to build the three teams with projects within because we need to build trust of the end users towards the system with which they're going to work, we need to create trust between all the stakeholders, the communities that are impacted by AI and those that are going to produce the AI system in particular communities and to trust across cultures, cross corporations, across countries, and these three angles have been essential in our view and they could have other ones and tomorrow we'll discuss together. Here I just want to tell you, what will happen tomorrow, it is going to be hopefully very, very interesting, and very fruitful in terms of the nine projects that we'll present and discuss. However, we think and we hope there will be a starting point of really a long term effort to build trust in AI. We're going to announce, we'll give you more detail about that in the day after tomorrow in the session that we'll summarize what will happen during our track and also move forward and we'll announce what we call the trust if aic tri and we believe that the trust factory will be like a global incubator for projects related to building trust between AI and communities between AI and different stakeholders between AI and then user and we really think that this is something that's needed. We all talk about trust in the applications and domains, so it is important to have a place where people can go and work together in a Mullfy disciplinary, multistakeholder way to build trust with small projects, very specific, like the ones we're going to present tomorrow that are going to build trust and show how we can build trust from various angles.

Think about really an incubator with the resources and work and people coming from different stakeholders. This will be -- we'll be funding for projects and we hope that many of you will join us in supporting this initiative, not just again for tomorrow, for a short time, but for the very long time. I hope that you all will join us tomorrow, we will see many of you in the discussion and I really hope, and we all hope, that you will join us in this idea of the trust factory. There is a website already for the trust factory, it is called trust factory.ai and you can go, look at the main idea, but later on, we'll give you more information the day after tomorrow.

Thank you.

>> Thank you, all. Can I have a show of hands, who plans to attend the satellite session tomorrow? All right. Okay.

The health one? Okay. A little bit more maybe. Yeah.

The smart cities? Okay. The trust? I think that looks like the majority. It is pretty reasonably equal.

We'll start at 9:00, the sessions start at 9:00 tomorrow. We don't use this room, we use the two other rooms here close by and then the two rooms in the building where you came in this morning. So that concludes the official programme of day one. Now we have a reception and the reception we thank ACM for sponsoring that, it will take place in the compound of the United Nations. We have a little walk there, which is maybe 5 minutes perhaps. We meet at the building where you registered this morning. We go into two teams, one team is leaving at 7:00, and one team is leaving at 7:15 so we go across this area with this Chair with three legs so that takes maybe 3, 4 minutes. Within the compound, it is maybe 3 minutes or so to walk. The reception, it will take place on the top floor of the United Nations building when the weather is nice, they have a fantastic view of Geneva and the mountains. I was told it is not raining right now, but the probability is 40%, I'm not sure what that means. Maybe you want to take umbrellas with you. You will not be able to enter the United Nations on your own, you have to join one of those two groups, one group is leaving at 7:00 and the other at 7:15. Looking forward to seeing you at the reception.

Thank you.