RAW FILE

ITU

AI FOR GOOD GLOBAL SUMMIT

GENEVA, SWITZERLAND

DAY 2, 16 MAY 2018

09:00

FOUR AI BREAKTHROUGH TEAMS - THE HEART OF THE SUMMIT

ROOM K

TRUST IN AI

Services provided by:
Caption First, Inc.
P.O. Box 3066
Monument, CO 80132
1-877-825-5234
+001-719-481-9835
www.captionfirst.com

***

***

>> Good morning, everyone. Welcome to the Trust in AI track. My name Huw price, from University of Cambridge, involved with the Leverhulme Centre for the Future of Intelligence.

>> FRANCESCA ROSSI: My name is Francesca Rossi with IBM Research and the University of Padova.

>> HUW PRICE: Let me tell you something about the structure of our track, for those of you who missed the introduction at yesterday afternoon's session. It's clear that one of the most important things to ensuring that the AI revolution goes well and is widely beneficial is trust, but trust in many dimensions. We need to ensure that the machines and AI systems themselves are trustworthy, of course, but that's not the only respect in which trust is relevant here. We also need trust and well-earned trust by humans of humans. For example, the end users and other stakeholders who stand to benefit most from AI need to trust the developers and other people who are offering them these solutions. And developers and others who need to work together collaboratively to ensure that AI is beneficial need to trust each other across the various kinds of boundaries and borders

which conventionally divide people when they try and cooperate. I am thinking not just of cultural and national borders, but commercial borders between countries and disciplinary borders between people, say, from technical disciplines and people from the social sciences and humanities. So we need to find ways of building trust across all those kinds of borders.

So the way we've approached this track is to divide the problem of building and earning trust in the AI space into three themes corresponding to three ideas we just talked about, and under each of those themes we've -- the wonderful team of people that we've been working with have put together three proposals, three project outlines which we want your help in working on and in refining, and which we hope will go forward and produce some concrete results over the next 12 months. In each case, results which will do something to build and engineer well-earned trust in the AI space.

I hope everyone now has a copy of the program booklet, looks like this, has a picture of our wonderful team on the cover. If not, there are more on the desk at the back where you came in. The day is going to be divided up into three 90-minute sessions, and in each of those three sessions we'll hear from the three project leaders on the three projects which fall under that team.

At the end of the day, we have a fascinating panel at the end of the day, and then after that, breakout sessions which are all encouraged to participate in, where the nine projects individually will have an opportunity to workshop the project ideas, to bring in

new ideas, and to bring in new partners, possible funding, things of that kind. We are trying to ensure that all of the projects go forward from this meeting and produce some concrete results by next year's Summit.

Francesca, what have I forgotten to say?

>> FRANCESCA ROSSI: Well, and also as you know, as we said yesterday, these are nine projects that we really would like to push because I think they are just a starting point of building trust in AI along these three dimensions. But these are just starting points. So if you have any other idea for another project which is different from these nine ones or for expanding any one of these nine projects, you can just get to us during the day today, or you can go to this website, TrustFactory.AI, where there is a contact form, and you can give us any idea and you can get hold of us and start a discussion about possible future projects or anything related to that.

As you know, this day today is a starting point of something that we hope is going to be much longer term, which is already of this big incubator, global incubator for building trust in AI, which is called the TrustFactory, and so between today and tomorrow, hopefully we are going to have very good beginning of this initiative.

>> HUW PRICE: Thank you, Francesca. Now I am going to hand over to my colleague, Stephen Cave, who is going to introduce and share the first session this morning. Thank you.

>> STEPHEN CAVE: Thank you very much, Huw and Francesca. Good

morning, everyone. I am Stephen Cave, Executive Director of the Leverhulme Centre for the Future of Intelligence, or CFI as we call ourselves, at the University of Cambridge.

This first session, we are going to talk about building trust with the stakeholder communities who we hope will benefit from this extraordinary technology. It goes without saying, I am sure you all understand that if those relationships of trust are not built, then we will not be able to use this technology to its fullest. That trust has to not only be built, of course, but also earned.

Now, we have three very different ways in which that can happen in different kinds of communities that we want to present to you this morning. We hope they will inspire you to get involved and to contribute or to develop project ideas of your own.

The format this morning, we have about 25 minutes to half an hour for each of the projects, so I am going to ask each project leader to come up and tell you a bit about their idea, and there will be plenty of time for questions, which you can submit either via this very sophisticated app or the old-fashioned way by putting your hand up.

So moving on to the first project, which is on building trust networks in AI for mental healthcare, we have Dr. Becky Inkster, a neuroscientist and specialist in digital psychiatry and AI and ethics, and her collaborator, Rafael Calvo. So Becky, Rafael, over to you.

>> BECKY INKSTER: Good morning. Can everyone hear me? Excellent. Just wait a moment for my brain to kick in and my slides to load. But thank you very much for this opportunity to speak with all of you today. The idea is that we are going to be sharing our far-reaching and very innovative -- and this is not just in the AI space, but we need to think in all creative spaces.

I will keep talking. I will show you a slide, we call this freestyling in hip-hop. I am going to freestyle. You have to have a lot of interest. Throughout my career, I have dedicated my time to mental health, and I can guarantee that I am going to continue every day to dedicate the rest of my career and my life to mental health. This is so crucial, and I do have some slides later to show some scary statistics, but we do need to think very creatively about these issues.

In my experience, I am passionate about mental health innovation, and I am also passionate about crossing sectors. I see that there are gaps in opportunities and accountability, and we must explore this if we are going to gain trust and to make some progress here.

So some of the work that I have been involved with is the human genome project. I got my hands dirty, sequence to genes. I have also used laser to extract cells in post-mortem to examine proteins in the brain. Building further, looking at neural networks in the brain and how the pathways connect, and again going further, I've been very passionate about exploring culture, so the other end of the spectrum here, using music such as hip hop to dissect lyrics and

understand mental health from a very different perspective, to use decades of hip-hop lyrics to understand what the drug messages might be and other clues from an epidemiological perspective.

Okay. We are ready to go.

This is my Instagram. I touched on the genetics, the genomics, the neuroimaging, and then I sort of swayed more to the cultural and social side of my work, which is also designing bespoked jewelry with an ethical focus. As we move in an AI direction, I want to make sure people are aware of issues and things like autonomy. So that example of the jewelry is computer code for the word "privacy," so really exploring ethics is important to me. But going back to the music example, we, as scientists, have a responsibility to reach far and wide, not just in academia, and one of my proudest pieces of work is on the Pulitzer Prize-winner Kendrick Lamar and dissecting his lyrics, which is ranked ninth of all time on the Cambridge University website for most views. You wouldn't expect this in an ivory tower. It reached a new demographic. As scientists, we have to do this. We've got to up our bar.

Some of the other examples of work that I do is I volunteer in a homeless shelter, and I have gone inside prisons to teach psychoeducation about mental health, and I battle rapped at the same time, and I won because no one wanted to challenge me. I am not very good either. But so I think it's important that throughout all of my experiences that I am sharing with you now, the take-home is that we have to be innovative in the work that we do and push ourselves,

especially as scientists, to reach out because trust won't be handed to us on a plate, and it won't be necessarily in one compartment in a tidy way. We need to go out and earn it, and we need to roll up our sleeves, and we need to listen.

In the last two months, I have had some really fortunate experiences across sectors, and it's really revved me up to continue to cross sectors to get my hands dirty within sectors, but to cross. So I was invited to go to Menlo Park Facebook the exact day that the Cambridge Analytica scandal broke. Shortly after that, I was invited to speak at the UK parliament, a launch for an all-party parliamentary group on social media and young people. Then I recently came back from New York speaking at the American Psychiatric Association, so really making strong connections with the psychiatry community here.

And again, just to illustrate that we need to bridge these gaps and these divides. Every summer I run a conference, and I can mention that a little later on.

So for those who aren't familiar with the statistics on mental health, you will be soon. And I am surprised if you haven't already. It's on the agenda of almost everyone. I am just going to share a few statistics here today, just given the time limitation, and I will lead them out for the people in the back. Depression is the leading worldwide cause of disability and ill health. And in just over a decade's time, by 2030, mental health is predicted to be the leading global disease burden. So this is very serious indeed.

What's also very alarming is that these symptoms are emerging during adolescence, so we see 50% of these cases that we could start to address below the age of 15. 75% of the cases below the age of 18. And it really is an opportunity for technology here.

What's very saddening is that the social situation with poor children and poor households really bearing the brunt of the mental health impact and being much more likely to experience symptoms and receive -- have problems with mental illness.

Mental healthcare is broken. The systems are broken. There's a shortage of care, and it's a bad situation that's gotten worse. You would think by technology improving, the demand and the awareness and all the campaigns, which is wonderful, we are not really ready for this yet offline, and this is very important to keep in mind. So for example, India's population of approximately 1.3 billion people have 5,000 psychiatrists. So it's quite alarming. And even in places like America, it's very difficult to recruit medical students into psychiatry to practice for many reasons.

There are extremely waiting lengthy times anywhere between 6 months and 2 years just to see someone to talk about your problems at the moment, and that's just utterly unacceptable. It should not be the case.

So when we think about trust, even putting AI aside for the moment, trust in psychiatry is in a difficult place already, so we need to consider things like the influence of big pharma companies and the impact that they have on psychiatrists and how medications are

dispensed and given to patients and the side effects of medications and the treatments and the resistance to treatment.

And then there's stigma. We might have an algorithm that's 100% predictive of a suicidal attempt, and we are unable to reach that person because we are unable to go that extra mile to reach a person. If we are gaining trust, this is one area we have to tackle immediately. And there are no or very few promising biomarkers in society. So this is also a big challenge for us. How do we get people to trust that this is an illness where we can use technology, Big Data, and things like that to create change?

Now, in the UK and other places in the world, we have seen some very disappointing news headlines about NHS trust -- storery -- having breaches and other issues with sensitive data, patient information. And again, it's not just the trust in the healthcare system; it's also the charities who, you know, with all good intentions, are making mistakes, and this is having a huge impact on the public's trust. So for example, monitoring or surveillance of tweets and news feeds. This had a backlash effect on a Samaritan app that was shut down within a days of its launch. The public are not happy even with a charity doing this without consent, not asking about this process. Trust is truly at the core of what we do in psychiatry and what we should be doing in mental health. So there's no need to hammer that point further

But AI brings an extra challenge to the situation. Once you've sort of -- the honeymoon is over and you realize all these wonderful

tools, there are a lot of issues that are, in some cases, unique to the AI space, but we've got a lot of issues. So again, we see more headlines with the bridge Cambridge Analytika. I used this on a zero budget, for good, and involved my participants 100% who were very sick individuals with depression. There are ways we can use these tools for good and use them efficiently, unlike other ways. Now, hashing, again, is a concept that I can explain in more detail, but extracting information, even in its anonymized form, can still be very painfully upsetting and a breach of trust but also of the data when it's paired with other information. So there's a lot of issues of transparency and power asymmetry at the moment. And also this emotional contagion. We have to be very mindful that emotions can travel through networks and that this can happen online, so again, if we are using this as a tool for good, there are opportunities here, but we need to be ethically aware.

When it comes to psychiatry and making claims we can diagnose people with AI, it's dangerous. We need humans to be involved in decisions. Because it's very easy to think of a scenario where someone using algorithm A tells a person they are very likely to attempt suicide within a year's period of time and B which tells them they are fine. I have been commissioned to write a series of papers, one of which focuses on digital ethics, and I have scraped the surface in this talk, but I would love to talk about this in more detail. One is where is the physical examination going, if it's actually a thyroid condition or a tumor, things like that. There are

comorbidities with other diseases, and people are seeing a GP for care, so how do we handle these types of things in the new space we are in?

But I do want to sort of uplift the direction and the conversation that there are so many incredible opportunities for psychiatry, digital psychiatry, and just sweeping across these images below, you can imagine that a germ could be delivering an -- a GM could be delivering an online set of prescriptions for a person across a jurisdiction for which we've managed that legally, and that adherence to the medication, which is a huge issue in psychiatry, could be handled with sensors that are digested, which increases the likelihood of patients staying on their medications and improving their well-being.

We can imagine a scenario where a voice and the face and all of these wonderful AI tools could be used to track the patient's progress and to look at prodromal early warning signs of the onset. As I've mentioned, adolescence is such a key time here.

And also offering opportunity to increase access to therapies that one wouldn't necessarily have the chance to consider. So virtual reality, providing a safe environment, a trusted environment where we can explore revisiting exposures to some dangerous circumstances and triggers that might relate to their PTSD and anxiety and social anxiety therapy. So it's a wonderful opportunity here.

And so there are other examples here, but what's interesting and the last thing I want to highlight on this slide, is that people tend

to really enjoy and trust speaking to chat bots and in this sort of anonymized digital space. There's evidence now starting to back this up. So the solution really is, I think, simple, and it just involves a lot of decision-making along the way. We need to identify where trust is broken down, and where it currently exists. I think that we are neglecting many offline and online situations that we need to preserve and utilize in a better way. It's not just technology that's going to get us to where we need to be. We need a paradigm shift.

So step one, I propose, is at my summer conference to have a roundtable discussion based on the output of a Hackathon, a trust Hackathon, where we go through tech products and we evaluate them from many different perspectives, not just a tech perspective. And this would help form a scoping review which would go out globally to identify where trust is broken and where it still exists. And thirdly, to create or to empower the idea of a paradigm shift. We are in early days of AI mental healthcare, but I think we need to shift our direction already.

So how I propose this is it's called digital social prescribing, which I will summarize in one minute and hand this over. But the idea really is that wouldn't it be incredible to be able to walk around and have interventions at your fingertips in your local community, offline, powered by AI? Of course, this can be opened up to digital community spaces as well. But the idea here is to have a conversation and a trusted alliance with a care professional and a tailored

treatment choice where your preferences, your cultural preferences, everything that you enjoy is factored into your symptomology, your care plan, and the AI powers it behind the scene to decide what is potentially the best decision in your community to help you with your socially prescribed mental healthcare. So there are lots of different choices that may be drawn to you, and they are really literally at your fingertips. This helps build the community.

So in London, for example, sometimes you can stand in the middle of a busy street and not realize there are 300 opportunities to get involved, to be less socially isolated. But it, again, stresses the importance of having a care professional guide you to not be heading to vulnerable positions or places. These are centers of verified trust locations to receive care.

And this is just an example of how the app works, and again, just focusing on technology, either facial recognition or other ways to ensure that it's secure.

And I will just end with this quote, that the most exciting breakthroughs of the 21st century will not occur because of technology but because of an expanding concept of what it means to be human. I love that quote, and I think that we need to remember and preserve and protect our humanity. I think that humanity will drive the future and that AI can power the future.

And these are a list of the partners. There are more partners that I can't name officially, but throughout my work, I have pulled plenty of people together, so it's a much greater group than this,

but these are the confirmed partners for now.

And I am going to hand it over to Rafael to continue talking about trust.

(Applause)

>> RAFAEL CALVO: Thank you, Becky.

So psychiatrists, doctors, nurses, they know a lot about how to develop, build trust with their patients, and they know very well the importance of building this trust. When a patient trusts the doctor, the psychiatrist, they are much, much more likely to follow the treatment they are being prescribed. They are more likely to change behaviors. They are more likely to adopt and follow the instructions or guides or suggestions from the doctors.

So I had a video there. We in our well-being technologies lab, we are developing different technologies that -- and different methodologies to build trust between the tools and the users and between different types of users. So in this video, I wanted to show you, you can access the tool from EQclink.postcomp.org. We use computer vision techniques in a teleconference platform. So let's say Becky here is my psychiatrist, I am the patient. We are using this videoconference tool. We can record the video and then using computer vision, machine learning techniques, automatically pick up the important features.

Now, what are the important features in this communication skills platform? Autonomy, the doctor, Becky, will be providing as much agency as possible to me. This is something that is common in all

the medical training programs. Competence and relatedness. These three factors help build rapport between patients and their doctors or their psychiatrists. And we can build this into the new eLearning and teleconference platforms so the new doctors learn how to do so. In fact, we have done already a number of studies. We have about 500 or so students using the platform in Sydney every semester. We have run cross-cultural -- so one of the themes in this trial will also be looking at cross-cultural trust. And in this platform, we are looking at how, for example, gender differences or cultural differences will change the way we communicate between each other.

So in the video there, you can see this is just a simple effective computing platform. We can track things like in the conversation, signals like listening versus just speaking. You know, when you you have a doctor that asks open-ended questions, all the studies show that helps build a sense of autonomy and a sense of trust. What we've done, for example, in cross-cultural studies is that doctors, when they are talking with people from other cultures, are less likely to be listening and to actually mimic. Mimicry is very important communication approach to building rapport. So when I know generally you are talking to someone else, the other person will know. This is a natural way we have on how we build trust and rapport with the other person we are speaking with.

So we are building these tools, we can look at how much people mimic each other in a communication, and if this is not being supportive of that trust development.

The second approach we are doing is we have built a number of mental health apps. You can download one of them at headgear.org.au. This we just finished the biggest RCT ever for a workplace mental health app. From the beginning, we are following approaches, design and development approaches, that are human centered. So human-centered approaches have been shown to be the best to gain, again, the trust, but the buy-in from the end users. They are participating in the development, in the design of the app, they feel it's their own.

So in all these projects on mental health for youth, mental health for workplaces, we use human-centered approaches that allow us to understand better what the users feel and what will be the most likely to be engaging to them. For example, we run a large study using different Facebook campaigns that were used to attract users to the app. And you can look into how different Facebook campaigns, different languages, different images attract people with different mental health risks. So the person comes to Facebook, and they see an advertisement that uses language that could be about well-being versus another advertisement that is about mental health. Different wording. Different wording will be interesting to different types of people who have different risk profiles. And this allows you to possibly customize the content and build a better experience.

So you are welcome to ask more about these technologies in the break. Thank you.

(Applause)

>> Ladies and gentlemen, we have a few minutes for questions.

Thank you very much, Becky and Rafael. A few questions came in. One, sometimes we are so focused on building trust, and that's valuable, but that's a crucial point and fundamental to how we think about these issues. Another Cambridge philosopher, the great Nora O'Neill talks about the importance of trust and trustworthiness. We need to make sure those are aligned. Today in this session, we are talking about how developers can earn -- not just build, but earn the trust -- of the stakeholder communities. In the last session today, we will be talking about developing technologies that are genuinely trustworthy.

Now, we have some questions come through on the app. One has received four votes, so I will put it to you both. Are there any correlations for people suffering from mental health issues and smartphone use that might hamper the effectiveness of this approach? So addictive tech use aggravating mental health issues? Have you got mics? There might be some.

>> I think what happens with addiction is sometimes technology satisfies personal needs. When we are using technology, we are looking to satisfy, for example, the sense of autonomy or competence or the sense of relatedness. And this need for satisfaction at a certain level can be very addictive. You are satisfying something that is important to you. But it's at the, if you want, interface or activity level.

What happens with addiction is addiction has the opposite effect. Actually, the word addictus means slavery, where the old traditional

Romans who were slaves in Roman times. So the word "addiction" has to do already with being servant to someone else. It's a lack of autonomy because you are losing autonomy in life, at the life level.

At the interface level, you feel autonomy. At the life level, you feel enslaved. And that's what happens with technology and happens with other -- with substances as well.

>> And just to add to that point -- no? Just to add to that point that everything in moderation, as we've all heard before, and but that we do need people around us to be our buddies to guide us, to help us if we are being too extreme in one way or the other. So it just -- it's an issue about monitoring and making sure that we are aware of too much of anything, you know, we have to just be careful about that. There are techniques and ways to handle that.

>> Thank you. Yes. Okay. Hett me take two questions from the floor. It's probably all we are going to have time for. Rumman, and then over there. Yes. They should work.

>> One thing to add to that question is that smartphones are not inherently addictive. It is the design of the application that is addictive. So for example, most social media apps are made to make you feel less than. For example, when you have a particularly popular post on some social media platforms, they will in the future post actually underreport the number of likes you have. To those of us who may not be constantly on social media may seem silly, but for people who are on social media for which that is a currency, you feel less than. These apps are constructed with that addiction in

mind.

I think what the two of you are proposing, particularly yours and actually yours as well, is about creating apps about integration. You mentioned human-centric design. So again, the technology is not inherently addictive. It's how we designed it.

>> BECKY INKSTER: I agreed completely. I think that's why I felt inspired to have this community-based app where you put your phone down and you integrate into society, and you surround yourself with like-minded people or different people, and you literally, your phone is down. But then you pick it up, you provide feedback about your experience, and the AI learns. But yeah, I think we need to preserve and cherish our offline world. And a lot of people are just focused on the tech solutions and the online space, but we absolutely need to maximize because we are underutilizing the offline experience.

>> RAFAEL CALVO: Yes, I think your point is a very important one. Everyone in the design human-computer interaction community now has become very aware of the impact and the responsibility we have when we design any technology. So I think the values of the designer are embedded in every technology we build. If you think about it, every chair you sit on, every keyboard you touch has been ergonomically designed to respect your physical health. How different it will be if every technology we design was designed to respect our psychological health. And this is a revolution I think that is happening, and these groups are part of this. Everybody in the design community, in the human-computer interaction community, is

very aware of this responsibility we have. We need to take into account our values as designers and developers of technologies. Those values will be the ones that change the world in the future. These are the things that we have to be taking into account.

>> Thank you, Rafael. That's a great point. It could be the theme for the whole day.

One more question. Who was it? Who had the hand up? Yes. It was someone over there who I promised. Yeah.

>> Yeah, so just to add to the first question that was asked, I was wondering if the embodiment of the smartphones -- I agree with Rumman's point on the design of apps being the real problem in this case, but if the embodiment of the smartphone kind of becomes a partial trigger for that sort of behavior. Because you know, when you talk about these psychological illnesses, what happens is that you have triggers in your environment that, you know, lead to certain kinds of behaviors. And I am wondering if in your research or otherwise if you've seen just that physical smartphone and the act of picking it up triggers a certain, you know, pathway in your head that makes you go a certain way, and if this would help to -- I don't know if this would be an approach that could help fix that.

>> BECKY INKSTER: Yes, definitely. So we are starting to see some research, but there's always bias in research. We see the good stuff. But there was a study on alcohol misuse, and for a subpopulation in that cohort, it actually increased their alcohol misuse. And this was in a male group. And if we are not teasing

apart our data properly and being very carefully with how we design our studies, we will miss these harms, these health harms that are embedded into the apps. I am not saying it was intentional, but it certainly backfired, where it helped some people, and it harmed others in the same study with the same tool. Sort of hopefully address that a little bit.

>> RAFAEL CALVO: There is a paper published recently that showed just having a mobile phone on the table of the family decreases the connection between families. I think the same thing happens with a TV, just having a TV in the dining room could be distracting to the conversations, reduce the depth of the connection that we build with family members. So it's not just about the mobile phone, other technologies, other instruments or devices that we place in our environment will have that effect.

I think what is different is that the intelligence of the devices can be used to grab our attention in ways that before it was not even possible. So with smart TVs, for example, the new effective computing techniques will allow you to use the camera in the TV to recognize people's facial expression, so then you can target advertisement to what's happening in the room, to the way people feel. If you mention the word "hungry," immediately you might get a fast-food advertisement; right? And this will be target advertisement, and you could take different positions regarding this. There might be a lot of people from companies here that sell these systems. But the effect is kind of what you are suggesting;

no?

>> BECKY INKSTER: Just to very, very quickly follow up on that point. Trigger is such a good word that you used, and trigger can be good or bad. In an attention economy, we can make it a cognitive economy. When you go out and you have these microinterventions at your fingertips, these are triggers for good. So I still think triggers are very important, and that's why the work that I am doing is to trigger social intervention at your fingertips in a microintervention way.

>> Well, I apologize we can't answer all your questions. I saw hands go up and there are many more questions on the app. I am glad we whetted your appetite. Please remember at 4:30 today there will be the breakout sessions, so you can come and join Becky and Rafael to talk further about AI and mental healthcare.

Now we have to move on already, but before we do, please join me in thanking them both for their presentation.

(Applause)

Next up, something completely different. We have a presentation on an app for building trust with poultry farming communities in East Africa. It has been developed by Dina Machuve from Tanzania. She couldn't join us today, but her clap rater, Ezinne Nwankwo, a visiting student at our Center in Cambridge, is going to present their project for us. Thank you, Ezinne.

>> EZINNE NWANKWO: Hi. My name is Ezinne Nwankwo, currently a research with the Centre for the Future of Intelligence, and I have

been working closely with Dina to prepare this project for you all today, so please think of any questions, ways of improving and feedback or ideas that you have to make this app better because it is a work in progress, and we are trying to embed trust in this app for good.

All right. So currently, this is the current situation in East Africa. So modern poultry farmers own anywhere between 200 to 2,000 chicken. Currently there's a community of farmers in a organize on Facebook and WhatsApp, primarily to communicate with other farmers. They share information on advice or advice on how to operate their farms, and in a day they can exchange up to a thousand messages. Farmers will then use Facebook to communicate and advertise their good to consumers.

They also record their data on paper notebooks, and it's all recorded in Swahili, specifically in East Africa. For example, the first item on this sheet refers to the chicken that -- the amount of chicken that they purchased in a day. It also refers to the vaccines that they bought for their poultry. And they record this information about how much they purchased and the date of purchase. Farmers then, they record the information daily, but after a year it can become really difficult to keep track of these records and also to eventually do something with it afterwards to help increase their production.

So we are proposing a one-stop shop for trusted data for agriculture where farmers can get real-time reports and services on

how to increase food production in that region. This app will be used on a mobile phone device, so they can easily record their data, and then with the help of AI, it can be translated into reports and other services to help with their farming.

So why East Africa? So the regions that we are considering are Tanzania, Uganda, and Kenya. And these regions were chosen strategically. So first, the farmers in this region speak Swahili, so that would be the joining language between all of them, between both developers and farmers. And although we are proposing to develop the app in English, the data and communication between farmers and app developers will be in Swahili.

The second reason is that there is a high need for food production in these regions, so the potential for impact is very high.

And the third reason is that poultry data and just in general the average yield of a farm in a year and general agricultural data is pretty scarce in this region, and so with the help of this app, we are also hoping to contribute to that -- in that area as well and providing more data to assess food production and other national statistics for this region.

So Dina has a lot of experience in data management and just building apps like this for other enterprises, so this is an example of one of the apps that she built for a tool for food processing, small to medium enterprises in Tanzania. And this is another example of a banana disease management tool that was used for agriculture extension workers and small Holder farmers to be be able to predict

when a disease would come up in their banana production.

So both of these projects Dina was the project leader and was very involved with building trust among the community to get farmers to use this app. So we are hoping with this new development that we can work on what -- work from these past experiences to see what did not work, where trust was broken, and the usage of these apps to inform how we can continue forward to make sure that these technologies are sustainable and that people are using them and that they are beneficial.

So these are some of the services and just the value proposition of this app that we are providing, so better farm management and recordkeeping, just being able to provide a better way to store data and to keep track of all of their reports in a year. We are also proposing to analyze the reports in order to help with sales and the health of their farm. We are also hoping, again, to provide raw data for future research in this area.

And so for farmers, some of the services would look like -- so tangibly, it would include the mobile app and reports and advertising services for consumers and farmers, and then an online database that contains all of this raw data.

Some of the stakeholders involved include the poultry farmers, the agriculture extension officers, who those are the first line of contact for poultry farmers, the agrabit shops where farmers can then go and sell their goods. The AI researchers. So we are currently collaborating with data science Africa, because in addition to

providing a service for farmers, we are also looking to build on the capacity building on the continent of Africa in providing a project for AI researchers there on the continent to work on and build this technology. And also app stores to sell that.

So some of the key milestones that we are looking to achieve in this coming year, which again, we would definitely love your feedback on, but so first, to survey the farmers to assess where trust has been broken and where we can start to build trust and other problem areas that they see currently, because again, we want to make this sustainable because it's not enough to just build the app and then say here, farmers, use it. We really need to make sure that we are giving them a service that they feel is beneficial to them, that they can trust. Also farmer recruitment, making sure that we engage with the farmers in this whole process. Collecting and analyzing the survey data and presenting the results as well in order to engage the AI community in Africa at the data science conference that will be happening this November. And then we hope to start the initial app development and then starting to get farmers to use it. And hopefully through this process of first assessing trust and building that into the app, we can start to get farmers to use it and make this a sustainable effort.

Yes, so yeah, these are just Dina's affiliations and my affiliations, but again, we welcome your feedback and any support. Thank you.

(Applause)

>> Thank you very much, Ezinne, and a couple of questions have come in already. What ability is there for the farmers to privatize their data and not have it included in the national data aggregation? Who ultimately has ownership of the data?

You have a mic, yes.

Have you got it? Otherwise you can use mine. There we go.

>> Hello. Yeah, so who has ownership of the data? So yeah, so that -- so for this project, we do want the farmers to feel like they have ownership of the data and to actually have -- not just feel like they have ownership of the data, but actually have ownership of the data. So we are trying to explore options and different frameworks to do that. So hopefully with the survey, that is where we would get that system hammered out and get their feedback and understand where it is that they would like to do with the app, how they can feel like they have control, and have control, and what abilities and different services we can provide to make sure that that control is there because that is really important.

In terms of the -- it being allowed in the national system, so some frameworks of this have been an opt-in system or opt-out system, where farmers can decide if they want their data to be a part of this or not. So we are also exploring that because, yeah, that is important to building trust and making sure the farmers have ownership of their data and feel -- and have privacy as well for them. Yeah.

>> Thank you. And a related question that has two votes, what security mechanisms are being designed in the app to protect their data?

>> EZINNE NWANKWO: To protect it. Yeah, so that is also still currently up for debate. Yeah, again, I am not sure the specific security mechanisms, but if you have any ideas, please come to our table and provide them to us because we are looking for feedback. So yeah.

>> Great. Thank you. Would anyone like to raise their hands and ask a question, then please do. Otherwise -- oh, yes.

>> So in your presentation, you mentioned that you needed to rebuild trust. I was just wondering if you could give us a sense as to how that original trust got lost. Like, what is the context of which you are going into this project thinking farmers are already distrustful versus needing to build something for them that they will trust into?

>> EZINNE NWANKWO: So I think from talking with Dina and her experience with the previous apps, I think what ends up happening a lot is that so these apps will be developed, but after time, they are not maintained, and then there's not a continual process of engaging the farmers and teaching them how to use the technology, what they need to do. So they resort back to what they know and is more convenient for them, which is the paper notebooks. So trying to figure out a way to develop the technology so that it is more convenient for them than using the paper notebooks and so that they

feel like they are still engaging with the community of developers versus it kind of just being developing the app, giving it to them, and then disappearing. But feeling that continual engagement and support so that they can continue to use the app and so it's sustained.

>> Great. Thank you.

We have a question here, I think from Rumman, which has two votes. What do you think will be the reach of your app in the population? Do you think there might be a bias toward more affluent farmers.

>> EZINNE NWANKWO: Can you mean what you mean by affluent farmers versus nonaffluent farmers?

>> We are assuming smartphone adoption and social media usage, which might imply that if there is not a reach to a large part of the population that these might be the more affluent farmers who are able to take advantage of social media, technology, et cetera. I am not sure. I don't know the population, so it's really just a question I am posing to you. Might there be some sort of a bias towards the more affluent farmers who might have these types of technological skills and abilities?

>> EZINNE NWANKWO: Yeah, I think that is something to consider. From my understanding of the community so far, it seems like a majority of them at least use Facebook and WhatsApp in some capacity. Maybe it is not through their smartphones. Maybe it's through using a computer at a local data center, other computer lab in their region. But my understanding is that the majority of them communicate and use smartphone data, but that is definitely an important

consideration that we need to factor in and figure out how to then deal with that, so thank you for bringing that up.

>> Thank you. Yes.

>> One of the key issues to trust is kind of transparency and what you want to avoid is that people just use the app because it's more convenient and (?) that's what Facebook and all of these do. How do you build this in so you don't just get the data because it's more convenient for these people to use this, so kind of not a very clean --

>> EZINNE NWANKWO: I think that's why it's important to understand what is it they need and build it into the app. Again, we don't want to collect the data just because it's more convenient, like you were saying, but we want to make sure whatever we do its going to be helpful and beneficial and supportive of the farmers, so definitely collecting their feedback and understanding what it is that they need help with in terms of running their farm. So I think that would be -- that is going to be the key thing to building in that trust and engaging with them and not just kind of developing this app because we think it's going to be helpful, but like what it is that -- determining what it is that they actually think is helpful and need.

>> (Off microphone) -- where you just have that data and people start using it and -- so I don't think Mark Zuckerberg started with the view, hey, I am going to be a Big Data collector, but he never paid attention to the questions and he ended up with all the stuff that brought him into trouble. So you have to build this in right

from the beginning. Don't ask me how because I don't know.

>> EZINNE NWANKWO: Okay. So how to build in trust right in from the beginning you mean?

>> How to Build in a mechanism so you don't run into trust problems later on. People start using all this data, people start doing stuff with it you didn't anticipate, can you prevent that or are you thinking about preventing it?

>> EZINNE NWANKWO: Prevent like third-party users from --

>> Whatever, or misuse or new innovative ideas or whatever?

>> EZINNE NWANKWO: Yeah, so I think we can anticipate some of the challenges that we are going to run into with -- just like by understanding the tech landscape. And yeah, we are working on coming up with solutions to then embed that into the app. So yeah, that is definitely a valid point.

>> And I liked your slide on the stakeholders that are involved in developing this. I guess ongoing stakeholder engagement and transparency about who is doing what with this app will also help to address these concerns.

We have a hand over there, and then one here.

>> I think since the beginning of the morning, we have a confusion because in English, you are very happy, you have two words -- you have trust and confidence. In other language, like French or German, you have only one word. That means you should define or agree that if you are speaking about trust, it's rational, it's with your head. I trust that because you have an experiment, you have a theory, or

I don't know what. Okay? If I am confident, it's maybe with the heart. Okay? It's a feeling. I am confident with you because I know you and we share the same value, so I am confident with you, but it's not trust. Okay?

So building trust or building confidence is not the same, and I like you have a strict definition to understand what you are doing because since the beginning of the morning, you always twist between confidence and trust. Okay? And that is a big confusion. Okay? So if you are saying building trust, it's not the same thing as building confidence. Okay? So I think you will have a -- it's a benefit for you if you do the difference. Otherwise, you will continue with the confusion, and building trust with AI, you will have a collapse at the end because this confusion should be erased.

>> Okay. Thank you. I am not sure I draw the line quite where you do on the difference between trust and confidence. I think what's more important to us is the different betweens trust and trustworthiness, so trust is, if you like, a psychological state, an attitude towards something that's a mix of perhaps the rational and the emotional. Whether that trust is justified, that is, whether what we trust is trustworthy, is a separate question. And you can see if you like the grand project is aligning the two. If you develop a trustworthy technology no one is using because they don't trust it, we will be missing opportunities. And equally, if people over trust an untrustworthy project --

>> You should also understand that (?) in the society, they

confuse the two. If you would like to have one chance of success with AI, you should do the difference between trust and confidence.

>> Thank you. We have time for one more question.

>> Is this working? I am sorry, I am feeling I lost my voice, but I will try.

For me, trust is the first point in sales. So what I think is that you are trying to substitute building in the trust with making sure that people still use the app later on. I am not saying it is not a good way forward, but what I am saying is we need to make sure we know what we are doing, and in this sense, if we are just trying to sell an app, whether it's an AI app or any other, we should not kind of build unity around such important values as trust when we are just trying to sell.

>> Okay. I don't think any of us on the stage are trying to sell anything, but Rafael, would you like to comment?

>> RAFAEL CALVO: Yeah, I will try to bring the two things together, maybe. In my view, there's a difference between the trust that a product makes with their customers -- for example, if you buy, let's say, German car, you will expect it wouldn't break down very easily; right? That's very different to the trust you have with a family member. And maybe that's the difference between the confidence and the trust. And that comes through a process. I think trust is not an outcome; you have to look into the process. And that comes from the point of where the person adopts the technology, why they download it, what were they thinking, what were they expecting, do they feel

it's going to be this hard or this easy, and is it going to contribute to my life in this or that particular way? So that's why the language that you use to attract people is particularly important because it helps them set up their expectations. So the way you set up expectations about your system is very important because otherwise they won't adopt it. That's has to do with the way you track them. But the whole process after they start using it and after they shut is down is very important. So I think that's what human-centered design can contribute to this.

>> Thank you very much. I am glad that these broader questions of what it means to trust AI and when we should are starting to come out, and I hope that they will continue to reflect on this over the course of the day.

I am afraid we don't have any more time for questions on that particular project, we have to move on to the next one, but before we do, please join me in thanking Ezinne.

(Applause)

And remember we will have a chance at 4:30 to break out and discuss any of these projects further.

Now, the last project in it session is something completely different again, which I think shows both the range of the challenges that we are facing, but also the range of opportunities, the range of interventions that we can achieve if we work together on this. And I am going to hand over to Irakli Beridze. Many of you will know Irakli already, he has been working on AI policy for many years.

Irakli words for UNICRI, and is also founding director of the Center for AI and robotics and is going to talk about the effects of AI-induced countries.

Irakli.

>> IRAKLI BERIDZE: Good morning, everyone. Thank you, Stephen, for this introduction. And at the outset, I want to thank ITU and XPRIZE for organizing this AI for Good Summit second time already, and I am happy to be here again, and thank you for organizing this very interesting session with a really good opportunities for us to present projects and see how this is going to develop into implementation as well.

My name is Irakli. I am heading Center on artificial intelligence and roteics, which was established in The Hague under UNICRI, United Nations Interregional Crime and Justice Research Institute.

This project, the title of the project is "building trust in AI - mitigating the effects of AI-induced automation on social stability in developing countries and transition economies."

Do we have slides up? I have my slides here, but you don't see it. It's good if we put it on.

But basically, while we are doing that, I would mention that one of the probably biggest issues surrounding the artificial intelligence is the issue of automation, issue of jobs, whether we are going to lose jobs, how many of them we will lose, in which part of the world, which sectors, which segments will be affected more, and which segments will be pressured more.

There are numerous reports globally which examines the issue of automation. The reports range from percentage to percentage of how many or how much of it will be lost. Some of them range from 20% to 70% or something in between as well. Certainly, this is something we don't know and something we cannot say accurately. Right? So hopefully slides will be up soon so I can continue with my presentation.

If we have a problem, I can do it without slides.

>> You know what they say, power corrupts, but PowerPoint corrupts absolutely.

>> IRAKLI BERIDZE: Yes, but it works out now, so we now have trust or confidence in PowerPoint. Here we go.

So I am talking about basically numerous studies which reflects the issue of automation. We don't know basically what's going to happen, but one thing we definitely do know that there will be major disruptions in the workforce. And this is certainly going to -- it's certainly not going to be evenly felt throughout the globe. Right? Most of the studies actually which I am familiar with are focused on the developing world, studies where the job issues of segments is examined in the developed world. There are less studies probably in the developing or transitioning economies. But the way we felt and the way I personally feel is that the automation will certainly not be felt evenly, but a lot of disruptions will be happening in the developing and transition economies as well.

Now, from then on, we look at the issues of social stability, how

this is going to result or what is going to be the effect and impact on social stability, and with such issues like increased migration and crime. Right? People will lose jobs, and people will tend to either migrate or do something else, and this something else might be translated into the increase of the crime rates. Right? And this is certainly sort of the essence of this track, definitely trust and belief in AI, and that certainly will be under mind.

Now, some of this -- okay. One of the solutions, actually, and certainly this is not the solution, one of a solution will be enhance understanding of the impact of developing countries on economies to the automation. And what we are -- what we felt is that we would need to identify actions for countries to take the mitigation potential -- potential negative impacts of it. And one of the important aspects of it would be to foster political support to implement actions and build trust and belief in artificial intelligence.

The approach which we developed and which we would like to present here -- and I will be available for discussing it throughout the -- during the breakout sessions -- would be some of the points which we want to develop during this project, and we would like to actually identify the points of contact for AI in the countries, and this would be designated by the Ministry of Foreign Affairs of the targeted regions and targeted countries. We would like to conduct initial general assessment of impact on artificial intelligence. Would like to follow up in that country's specific assessments and do pilot --

on some of the pilot countries. We would like to conduct workshops with national authorities and concerned entities to validate assessments on building local ownerships. We want to co-develop roadmaps of concrete actions to mitigate the impact and conduct awareness-raising workshops for policymakers.

It's a comprehensive study. It's a comprehensive assessment. And it's a comprehensive action-oriented activity which will be conducted in numerous countries. This in principle will create an infrastructure where we would understand the impact of AI automation to the social stability, migration and crimes, and create an infrastructure where all these approaches could be worked together and best practices could be developed, which would be shared throughout the regions.

Implementing partners of this project, we will be leading the project as UNICRI. We have many years of experience conducting such activities in different fields, and I will explain some of them a little bit later on in the slides and in my conversations during the breakout sessions. We have two partners at the moment, World Economic Forum it, who agreed to be partner with this session, and also the Center for the future intelligence. We are obviously open to have other partners as well during the implementation and during the development of this project as well.

A little bit about UNICRI. UNICRI just created a Center on artificial intelligence and robotics. We started this in 2015, continued with -- it started with creation of the program on AI and

robotics and translated into creation of a center which is established in The Hague.

Now, a little more about -- a little bit more about our activities and the infrastructure and resources which we have in place. We have network of partner countries, which we created already ten years ago, and working through it on some other issues. This network was created for the chemical, biological, radiological, and nuclear risk mitigation. And we have 55-plus partner countries, and we have similar methodology where we started to define the national focal points, national teams, and national action plans on the issue of chemical, biological, radiological, and nuclear risk mitigation.

Here we have the eight regional Secretariats which cover these 55-plus countries, and we have our individuals aware of situations within the countries and working with national authorities within the countries. And we are certainly looking at using this and utilizing this infrastructure to conduct and carry out the project which I presented.

I take this opportunity to also announce one important meeting which we are going to conduct together with Interpol. This is the first global meeting on opportunities and risk on artificial intelligence and robotics for law enforcement, and this will take place in Singapore 11th and 12th of July, and some of the findings and some of the interactions which we are going to have with Interpol Member States would also be very useful for the project which I presented here.

I thank you very much. This is my short presentation. But we will leave more time for discussion, for questions and your suggestions for the project.

(Applause)

>> STEPHEN CAVE: Thank you very much. Do put up your hands if you have questions. I will also take some from the app.

So first off, how different is this initiative from what the World Economic Forum is already doing? I saw that they are one of your partners.

>> IRAKLI BERIDZE: Right. So thank you very much, Stephen, for this question. World Economic Forum is one of our partners. The World Economic Forum would like to actually contribute with their data, with their knowledge and understanding on issues related to the automation, issues related to jobs and which segments will be pressured, but this project will mostly focus on the effects of automation on social stability and issues like migration and crime, so how this is going to affect these bigger issues which are ahead.

>> STEPHEN CAVE: Thank you. We have another question here. What are the specific sectors in developing and transition economies that you see being vulnerable to AI-driven labor market disruption? Is there any early evidence of this already taking place?

>> IRAKLI BERIDZE: Well, I mean, basically, this would be one of the outcomes of the project, to understand which sectors will be affected, which segments will be pressured, and certainly I have not seen a comprehensive study on these issues yet; therefore, it is very

difficult to discuss which sectors will be affected. We probably can feel some early signs of it, that while technology will be entering the markets where cheap labor advantage, for example, having advantage might be pressured, and this would affect the consequences which was described in the project. So this is something we would like to study through the project as well.

>> STEPHEN CAVE: Thank you.

Any questions from the floor? Yes, please.

>> I think it's really important what you are doing. My question is to innovate, you also need to consider innovating on the design of how we are doing things, and what I heard is a very traditional design, which starts with assessment and finishes with building awareness. So what other options, have you considered other options for innovating on the how?

>> IRAKLI BERIDZE: Yes, thank you very much. I think it's a very good question. Certainly the approach is quite traditional, building awareness and doing assessments. If anybody can suggest some better options and better methodology, we would certainly consider it. But to my understanding, assessment and awareness is something very, very important. We know we need to understand what is at hand, and we need to understand -- and the countries actually or the regions at hand should really have the awareness of it.

Right now, for example, I travel a lot throughout developing world, transition economy countries, and I don't see much discourse related to the AI-driven automation. I have not seen any discussions

either in the parliaments or in other settings in many, many countries throughout my travels which would reflect the discussions. Therefore, that leads me to believe that the awareness raising is very, very important there as well. Certainly assessment.

So if we have some other tools and other approaches, we would be happy to incorporate into the project.

>> STEPHEN CAVE: Yes, thank you, and which brings us back to the point that all of these projects are in the development phase, and so your input and suggestions are very welcome.

Now, we had a question over there.

>> Yes, two questions indeed. First, yours was a very nice presentation, but it is an official presentation from a UN body. Don't you think that any official rhetorics and analysis is always a little bit biased because you start from the assumption that actually poverty is the main factor for instability and crime. It is sometimes true, but you very well know that we could have a list of half the countries of the world disproving that. But you cannot say that because you are UN official and you have to prove that there is the relation between the good for the world and the human rights and UN action. And so all your statement is in a way already written in the marble of the UN thinking. This is my first question.

And the second one, it's a bit outside the scope, but still you are the relevant person. I am a Swiss citizen, and I believe I live in a democratic country. Then I realize that all the field of social reality, the one which is more opaque and on which there is least

data is what happens in court, in the court of justice, if you are not a journalist attending the trial there, there is no way you can make any statistical study of any kind of offense because it is considered -- I don't know, in fact, I never could find out why it is so secret, but if you are not in the room, data are lost for eternity. What do you think of that, can it be improved, and can we one day make statistics on what kind of offenses, how often, on which type, et cetera, so that we understand better this social and judiciary --

>> STEPHEN CAVE: Thank you very much, Irakli.

>> IRAKLI BERIDZE: Right, so some reflections there.

We can put the UN aside and talk freely and discuss the issue. For example, when we are talking about the -- yeah, if you hear me. So yeah, we can put the UN aside and talk freely about it and discuss it, what we have at hand and what is there an issue. The project, what we presented, Rand the understanding what we have, and there is a very clear correlation that if big or large numbers of the population will have no income and will live in poverty will have serious problems and will have negative consequences to their livelihood; right? And in numerous cases, this is translating into increase of migration, increase of crime rates, increase of negative effects of it. So basically, this project is trying to examine and will try to examine that type of negative effects of AI which is going to undermine or which may undermine the trust in artificial intelligence. So basically, that was the essence and understanding

of it.

Certainly, throughout the implementation of the project, we can examine different types of reality, different types of data, and if we find that this is not the case, certainly we are going to come out and say it. And I don't think that United Nations would have a problem of actually saying that certain issues don't have correlation which we believed previously. So I am quite open.

On your second remark, I did not quite understand what was the question, so if you might want to rephrase it.

>> STEPHEN CAVE: Can I suggest you talk about it in the break because we have some more questions more directly relevant to the project that are coming in.

First one from the app and then one over in the corner. From the app -- perhaps this is also a little bit of a tangent but important perhaps for background. Could you say something about why east Asian countries are not represented in the network of partners of UNICRI and perhaps whether this would impact on this project in any way?

>> IRAKLI BERIDZE: Right. Certainly it will not. The example which I gave was the infrastructure which we built for the CBRN risk mitigation and nonproliferation project which we conducted -- are continuously conducting throughout ten years in support or with support of the European Union, and this is not an inclusive area where we are going to conduct that particular project. So that's where we have the infrastructure at the moment built. This is a continuing, growing infrastructure, and other regions could be

incorporatedeasily in the project As well.

>> STEPHEN CAVE: Thank you. A question over there as well.

>> (?) global network an ethics. I just came back from Nigeria teaching ethics to students, and I talked about AI and ethics. And then students came to me and said they are approached every week -- not every month -- every week by the (Inaudible). There is a poverty relation. You are a student. You don't have money for your student fees. You are approached by a sidewalk criminal. What to do? You accept the offer. I did very concrete coaching for students how to remain with your integrity and how to refuse that and what would be the alternative. So my question, how to deal with that in order to strengthen the ethical internal integrity of students, which is very concretely related to your topic? How can this be strengthened? And I am sure -- most of the students say I don't want to become a cyber criminal, but if I can do it for one month and I get $10,000 or $5,000, that's a lot for me. I can pay my whole student fee with one month cyber criminality. My student fee for one year. So this is very concrete example.

>> STEPHEN CAVE: Thank you very much, yes, how can we resist temptation? I think that's probably a question for all of us. Irakli, do you have --

>> IRAKLI BERIDZE: Right, it is certainly an issue, and it is certainly a very, very big issue. There is no straightforward answer to it. There is no straightforward solution to it. I think it's quite a complex approach which requires to build networks of

infrastructures in the country, to build the -- wealth in the countries, to increase the economical output in the countries, to have better answers to the issues of criminality and better understanding of the new technologies and how this is going to reflecting.

For example, the meeting which I announced which we are going to conduct together with Interpol in July in Singapore is going to actually tackle some of these issues that how to deal with emerging technologies, how to deal with AI-driven technologies when this is going to be used by criminals and what are the infrastructures which countries need to build, what are the best practices or tools could be developed which would be shared with the countries, which they will be answering to that type of problems and issues.

>> STEPHEN CAVE: Thank you very much. We've run out of time for questions, so please join me in thanking Irakli.

(Applause)

>> IRAKLI BERIDZE: Thank you.

>> STEPHEN CAVE: Now, I hope that's given you a sense of the scope of the challenges involved in earning the trust of the stakeholder communities who are going to be affected by this technology, but also a sense of what's possible, and I encourage you all to think about how you might contribute to these projects, and also, I hope you will be inspired to think of projects of your own.

Do come along at 4:30 to talk about the projects of most interest

to you, and do come back in half an hour to hear about building trust across boundaries. And please join me in thanking all our panelists for their excellent project suggestions again.

(Applause)

And now time for coffee.

***

***