RAW FILE

ITU

AI FOR GOOD GLOBAL SUMMIT

GENEVA, SWITZERLAND

DAY 3, 17 MAY 2018

14:00 CET

POPOV ROOM

PROJECTS IN ACTION: TOWARDS AI AND DATA COMMONS - PART TWO

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 afternoon. Welcome to the second part of AI and Data Commons, and this is the last combination of everything you have done so far.

Happy to see that so many of you are here and participating. So we have an awesome panel, Dr. Chaesub Lee, Director of ITU; Francesca Rossi, researcher of IBM that you have all met; the CEO of a can be called BigchainDB, a centralized data and blockchain. And Stuart Russell, who is a researcher at University of California Berkeley.

It's my pleasure to start this session called Towards AI and Data Commons, which we heard this morning on the first panel the importance of data commons, and you brought a number of good questions about what should be shared or not shared and how this can play. In fact, data commons could be a portion or a piece of a larger system of framework. We heard about TrustFactory, the need for that. We talked about many different opportunities of projects, such as a Project Zero that Stuart mentioned before or the ideal city or the city platform. So when you think about all of these, it seems that there is a need for a connecting framework where different stakeholders can work together to make this AI for Good a reality.

So when we started this conference, the goal of this conference was to identify practical applications of AI with the potential to

accelerate the SDGs. And this is a thing we all did in the past two days to identify core applications among different sectors that could somehow demonstrate the roadmap towards the SDGs. But also we said as part of the goal of this conference are to formulate strategies to ensure that trusted, safe, and inclusive development and dissemination of AI technologies are possible and are equitably accessible to everyone. So the benefit of AI should be accessible to everyone. So this is a little bit of context of this session on making sure that however we build/design/deploy AI and we think about it and be part of it, how can we make sure that everyone in the world benefits from it, and the benefits can be defined in many ways.

We have two more panopelists joining us, Wendell Wallach that you have seen on Day 1. He is going to obe remote and joining us. And Vin sen sew Aquaro from United Nations is going to be joining us.

We would like to start with Wendell that give us first day discussion of onward, outward, and inward of AI for Good to give us a sense of from a project perspective what he thinks we should be doing towards a commonality of AI for all. Can you hear us, Wendell?

>> WENDELL WALLACH: Great. Can you all see me? Unfortunately, I am not seeing you at the moment, so how are we doing? Am I being broadcast?

>> We can hear you and we can see you.

>> WENDELL WALLACH: And can you see me?

>> Yeah. Sight and sound, yes.

>> WENDELL WALLACH: For some reason -- okay. Good. So my

apologies that I had to skip out on the conference a little early and I am only able to talk with you remotely. I am actually going to focus primarily on agile and comprehensive governance for AI, and I am going to begin with a few comments building on what Amir just said on why a dij Cal commons and governance are central to AI -- digital commons and governance are central to AI for Good.
In a sense, we can thinking about AI for Good about these discrete projects you all have been focusing on throughout these three days. My big concern is not that we don't put a lot of discrete projects in place, but more that the projects that we do put in place mask some other ways in which AI can become harmful, which really serve the self-interest of small stakeholders or at least small interest within society.

So at its best, AI for Good is a total movement where we are explicitly focused on enshearing that the adoption of AI is for the betterment of humanity and the adoption of AI is what the overall trajectory of artificial intelligence will be.

So when we talk about these projects such as a digital commons or agile governance, these are great words, but as you will be hearing more and more throughout these sessions, there are a great deal of both technical and practical and political considerations that will have to be met before they are realizable. But in both cases, what we are talking about is putting in place frameworks for cooperation, frameworks for people working together and hopefully working together in the positive interest of all of humanity.

So that's the basis for this conversation. Now, again, our working together is only going to be frusful to the extent that we are also mitigating the potentially harmful side of AI. So if, for example, AI decimates large numbers of jobs and that happens in a very quick period, then we are going to need some kind of economic, governmental, political economic response to that challenge. And whether that's the creation of new forms of work or whether that's a universal basic income, those are the conversations that are going on.

So let me move forward with two elements. One is to just give you a general framework to think what agile and comprehensive governance might be. And secondly, I want to finish up with giving you a very brief description of the BTI 4AI project. We use the number 4. I am using the word "governance" and not the word "government." Government is often thought of in terms of hard laws and regulations. Governance is a term now being applied to comprehensively looking at the vast array of mechanisms, and those mechanisms can be anything from a technological solution to imbuing technologies with values to corporate oversight to a plethora of other mechanisms sometimes referred to as soft law and soft governance. So soft law and soft governance include industry practices and procedures, standards, laboratory practices and procedures, insurance policy, a vast array of different activities. And the wonderful thing about these activities are they can be very agile. They can come up quickly if needed and they can be dissolved

if new technologies or other factors make them no longer useful. The weakness of them is that they aren't enforceable. So there will always have to be some partnership with governments for forms of enforcement of soft law when it is truly important that that soft law get enforced.

But the basic law here is that government is the last solution, not the first solution. That's really what we are directing ourselves toward. In that regard, we are proposing a new framework, probably instituted in new bodies, forms of civil society, that will engage in lightly coordinating the activities between the various stakeholders, being a source of a true multistakeholder platform, monitoring comprehensively the development of the field, and underscoring where there are gaps that are not already being addressed by other bodies. And if there are gaps, then looking through this plethora of different mechanisms for the best way to address those gaps.

Now, when my colleague Gary mar chant and I first came up with this model, we were thinking of it largely in terms it of the governance of emerging technologies and why it fits, why the existing governmental models really didn't work for these fast-appearing technological possibilities. But the more and more we got into this framework, people asked us to talk with them about the application of this way of thinking to all kinds of fields, so we are viewing this now as a new framework for which we need some pilot projects to see if we can flesh it out in the same way as you will be talking

in the rest of the session about how you might flesh out a digital commons.

In that regard, our thoughts were originally reflected on how you might put in a project in artificial intelligence for robots or synthetic biology that could be utilized by a nation in working through its strategy. But the more and more I thought about this, it became clear that we should start with an international body with complementary national bodies, and in many cases those complementary national bodies are already being created by countries to serve their own particular purposes. So it may just be a coordination between a plethora of different kinds of national bodies.

But this gave birth to what we call the Building Global Infrastructure for the Agile and Comprehensive Governance of AI Project, or BGI4AI. And as a first step in building global infrastructure, it's exactly what I am doing with you here. Getting people to think more creatively about governance and hopefully having some of these ideas adopted as various other bodies or various other nations put their governance frameworks in place. But recently over the last year have been working on convening an AI global conference for the governance of these emerging technologies. And the hopes are that we will have a governance Congress convened within the next year or so, and that what will be born out of this will be a new NGO that can begin to fulfill this more comprehensive oversight of artificial intelligence.

Now, as with a digital commons, there are so many technical

problems that need to be solved in that regard, so many political and practical problems. And furthermore, there will be the questions of how the digital platform can complement this agile governance and how the agile governance might, for example, ensure that the digital platform is truly beneficial, is truly AI for Good. So thank you very much.

(Applause)

>> Thank you very much. Thank you, Wendell. This morning we realize more about this idea of trust and importance of trust. Governance is definitely level of performance, and finding guidance and rules that can be implemented from the design steps to the deployment steps that captures the governance rules will be important.

Let's turn to Vincenzo from the United Nations and hear from him how he sees the importance of the assimilation of AI globally.

Vincenzo?

>> VINCENZO AQUARO: Thank you very much, Amir for first inviting me for this very important panel and sessions. Can you hear me, just to be sure? Hello?

>> Yes, we can hear you.

>> VINCENZO AQUARO: Okay, perfect, perfect.

So let me first give you a quick overview about other initiatives that the United Nations have been doing, starting from the United Nations High-Level Committee on Programs as well as Management that have noted that the cross-cutting nature and the transformative

potential of artificial intelligence is a driver of accelerated and structural change. The Committee agreed that the multidimensional nature of (?) such as artificial intelligence require an integrated, cross-ectorasectoral, and collaborative approach that mobilize and engage the entire United Nations systems.

Also at the Internet Governance Forum, artificial intelligence is considered as a key topical issue, as it is defined as part of its multi-strategy of its current mandate. I recall to all of you that Internet Governance Forum is a multistakeholder forum that works in this direction. And DESA is also convening every year the science and technology education forum, and a recording session has been included of the emerging issues on AI.

So why I am saying this? Because this Summit is already one of the most important, prestigious international fora on this subject matter. For the second edition, a multistakeholder community of experts from international organization, academia, private sector, and civil societies have been intensely discussed about the specific action-oriented artificial intelligence project, able to create impact and help achieve some specific sustain Sustainable Development Goals. Breakthrough teams have been presented during the Summit to demonstrate that the potential of impact of AI with some specific SDGs.

During these days, experts have strongly highlighted how AI could be one of the most powerful tools and effective enabler for achieving the SDGs. There were raised also concern, such as potential impact

of -- ramification for privacy, cybersecurity issues, and some other ethical concerns. But experts have also highlighted that it is not possible to determine effective solutions working in a silo because of the social implications of AI are really so complex. So we need to work all together in a multistakeholder, a multidisciplinary community to maximize the benefits of AI while working to address its potential challenges.

As outcome of the Summit, all the approved initiatives would be funneled in specific concrete action that will be implemented in the coming months. So for the first time, the outcomes of an international UN Summit are concrete action for building capacity on AI for specific SDGs rather than reporting on them. This is very important. But these concrete artificial intelligence summit outcomes can be also an unprecedented opportunity for AI community to show to the world in concrete ways the way AI and data commons could address some of these humanity's most pressing large-scale challenges and generate economic growth and prosperity as well as fight poverty and improve the quality of human life everywhere.

Going back to 2030 agenda, the 2030 Agenda is already spotlighted on the important principle of leaving no one behind. Beyond the set of measurable goals and targets, the 2030 Agenda is a universal agreement and a call for action for pursuing peace, well-being of people, environments, planet protection, prosperity, and partnership between the countries and within the countries. The famous five P. Those principles must be always applicable and

relevant for all countries, whether high, middle, or low-income countries, superpowerers, or small state. This is the universality of the 2030 Agenda. And this long-term universal principle should also be applied by the whole artificial intelligence community if we wanted to support the creation and the promotion of AI for the common good. For this reason, the principle of leaving no one behind should also or always be applied for artificial intelligence. As mentioned before, artificial intelligence is a tool and a powerful enabler to achieving Sustainable Development Goals. That's why AI should become app indispensable and universal -- an indispensable and universal resource for all of humanity to be equity plea distributed within and among Member States and to be available for everyone or every government, no matter its level of development or capacity.

In this regard, a novative multistakeholder participatory framework for AI commons should be identified, developed, and shaped by institutions and policies at the local, national, and global level. So one of the biggest challenges for the international agencies, governments, and the global AI community is to create a common framework to regulate the proper use of AI while not stifling innovation.

The UN general assembly also recognized that the pace and scope of rapid technological change can have a far-reaching implication, both positive and negative, for the achievement of sustainable development, requiring international multistakeholder cooperation

in order to benefit from opportunities and address challenges. There are thousands of applications already in place all over the world that could easily be deployed and made available for all. Many Member States in their digital transformation are already starting in close collaboration with private sector, academia, and civil society, are the results of their national R&D on AI can be used for the common good.

But what's presented at the summit in terms of concrete solutions is just the tip of the iceberg for all possible AI applications and solutions in supporting the SDGs. But what was really unique and potentially infinitely more valuable than the solution presented per se was these multistakeholder participatory frameworks composed by ITU plus 32 UN sister agencies working together with XPRIZE Foundation, with the Association for Computing Machinery, and with this huge and committed international community of experts and practitioners from governments, academia, private sector, civil society that have attended the meeting working together as one and making AI for Good the leading United Nations platform for dialogue of artificial intelligence.

So more than a (?) I was seeing it as a real ecosystem in which different communities and key players could experiment and propose and replicate new, inclusive, participative, artificial intelligence business models and solutions for the common good. In this regard, the whole ecosystem could and should work also as powerful combination platform for achieving the SDGs and to advocate

and solicit Member States and the other key stakeholders to a collective effort to create the condition to make artificial intelligence more inclusive, more accessible, and more affordable to everyone. It is really time for leaving no one behind on artificial intelligence, and we should all work in this direction.

So thank you so much for your attention.

(Applause)

>> AMIR BANIFATEMI: Thank you very much, Vincenzo. This is a powerful message and a hopeful and positive one.

Going along leaving no one behind, remember that the goal of there conference was to identify strategies for development and dissemination of technologies in an equitable way. Let's try to go towards the finishing of what could be a platform or a common platform for AI for Good? We are all talking about AI for Good. What would that mean?

If we look at the opportunities today, of course we need to use more and more AI to solve problems that are currently not addressed. We all did it together in the past two days, and especially yesterday, and we identified new projects and new opportunities. It is also important for the community of AI practitioners to identify and see and be knowledgeable about core problems and not just a few of them. It is also important for the community of practitioners to have a way to collaborate with problem owners, and problem owners are the ones that are identifying core problems that are closest to the issues and are screaming to solve them. And sometimes those problems are

not big enough to have commercial viability behind them, but they are certainly big problems. And the problem owners, the way we define it and we talk about it, need to have a way to collaborate with people who have problem-solving capabilities.

It's also important to have what we try to define as safe environment, an opportunity to have feasibility checked. It's also interesting to know that an idea can create a pilot and a pilot can grow. And for a pilot to grow, it needs to scale; it needs to scale to services. So how can we enable all of that? How can we find a way to make this happen? As you think about all the projects and initiatives that were explained yesterday, you probably recognize some of the foundations here. So if you try to talk about AI for Good and put it in quadrants, on one side you have arc I practitioners with researchers, developers, technologists, innovators, engineers. On the other side, you have problem owners, and the problem owners could be every one of you, could be a city, could be a mayor, could be a government, could be the midwife or the woman with young children. Anyone who understands a problem exists and wants to solve it. And in some cases, AI can be instrumental in this solution.

We, of course, talk about data commons, and we are not going to go over that, but data commons are very important in this picture. And obviously, we need services and technologies to make this happen, to have this collaboration happen, to have algorithms being built and deployed and developed and have access to basic cloud and computing services, including algorithms.

So if this AI for Good has to be in the middle of all this four pillars of four groups, how can we present it? We need a collaboration framework. How can we bring people together to collaborate? We talked about multistakeholder collaboration, we talked about alliance. But in practicality, we need to develop solutions that takes a problem and gets a resolution. This could be in the areas of health -- I am sorry if this is very small -- health, education, and different verticals. So participation, collaboration, and creating incentives to resolve those can be framed as a circle of collaboration.

So basically, if we want to think about a framework of AI for Good that we can call AI commons, will be a way to connect all of this. So what does that mean? It means that not only do we try to solve problems by putting practitioners together, but also we need to talk beyond service providers that provide computing and class services when you talk about scalable platforms. And the thought of scalable platforms was given to me by Stuart, and Stewart made me and us aware that in fact we need to go beyond pilot. We need to scale those pilots to global services. And he mentioned briefly this morning. I would like to invite Stuart to probably emphasize that and amplify that.

>> STUART RUSSELL: Sure. Thank you, Amir.

So it became very clear in the satellite discussion that the expense and engineering expertise and resources involved in converting a pilot project that had demonstrated success in solving

a particular problem into a global service that was then going to be available to the problem owners wherever they might be and all the countries in the world who might face that problem, that that transition was the main stumbling block. We have lots and lots of ideas. We have lots of ways that we can use AI for all of the SDGs, but we continually fail to make this transition. And then projects result in a publication or a report, a nice talk with some smiley faces, and then they get filed away on a shelf, and then the next year someone else has sometimes even the same idea and does another project that in three years' time, after a lot of effort, will get filed away on a shelf. And yet the real problems on the ground continue.

And we have to break this cycle. And I think one way to break the cycle is to have this shared infrastructure because when you think about projects that go from the raw data supply, which might be the cell phone company, it might be satellites, it might be financial transactions, it might be any number of things -- so going from the raw data supply, you need storage, you need computing facilities facilities, then you need a little bit of AI. So AI is a little bit of the solution. Then you need ways of disseminating it to the customer, so to speak, the person who wants to find out what is the weather today going to be this afternoon in Geneva, and do I need my umbrella or whatever it might be.

So almost all of that is common to almost all of the applications of AI, and therefore, if we can develop that infrastructure, we'll

be able to both amortize the cost across many, many projects; we will also increase the value of the data itself, which will then encourage further investments in creating new data sources. There's no point in launching a whole bunch of high-resolution satellites, you know, if we need 8 million people to look at the data that's coming down. We need to have AI to interpret that data. And if it's not there, then there's no value in launching the satellite. So you create a virtuous circle because people are more willing to build applications because they can turn them into global services quickly. People are more willing to create data sources because they can find more customers for those data sources.

Now, there are still some complications, and I want to just talk a little bit about my own experience with United Nations, and the main project I have done with the UN is with the Comprehensive Nuclear-Test-Ban Treaty organization, and we did some things right. So it started, I think, in March of 2009, where CTBTO invited some machine learning researchers to Vienna to learn about the problem. I have heard this from a number of people. Miguel from Global Pulse goes further. He says you need to send your Graduate students to be in the field in Africa or South America or whatever it is. They need to actually be there not just learning about the problem but living the problem. That might be a bit extreme, but I think there's certainly truth in that idea.

So AI people cannot just view this as, oh, someone is going to give me data, I am going to run an algorithm on it, then I am going

to send them the results. You have to spend days or weeks with the problem owners learning about the problem. And that was a very, very important process. Ever since then we have had fantastic collaborations with the technical people at CTBTO. Elena Tomuta is one of them, and she was here presenting in the Trust track yesterday.

So the other stakeholders besides the technical people in Vienna who run the technical Secretariat, are the Member States and the part of the United Nations that deals with contracting. So the University of California and the United Nations both see themselves as sovereign organizations with absolute power. At least the inging divisions of these -- the contracting divisions of these organizations do. So it took two years to write the first contract. We had already solved the problem long before we got any funding through a contract because it took two years to write the contract. The second contract took another two years to write, even though it ended up exactly the same as the first contract, it took them two years to agree to have the same contract that they had already agreed on before, so we got the second tranche of funding after four years, my students had long since graduated and all the rest. That was one sort of failure. I think part of what we want to do with the commons is have fairly standardized kind of agreements, and don't underestimate how difficult this is. I think almost every research project that involves academic researchers, corporations, you know, public data sources, public organizations will go through the same kinds of difficulties unless we standardize to some extent what the agreements look like.

And then, of course, there were the Member States, and it turns out that not all the Member States are actually wanted better monitoring of the nuclear test ban Treaty. So it took from roughly late 2010 where basically we had a solution to the problem until January of this year for the switch to be flipped and our software to be actually running, and this is again something I should have anticipated, but getting all the stakeholders onboard ahead of time would be really helpful.

So I almost feel as if every one of these projects is going to take place in the data commons almost needs a kind of shepherd who has experience in bringing these kinds of projects to completion and making sure that you don't run into all these roadblocks and these different ways of failing that an AI researcher simply is not equipped to anticipate.

We need to have the stack and infrastructure and resources to deliver solutions from pilot to global service. I think that will make a huge difference. But it isn't just a technical problem. It's a sort of a governance problem, a sociologic problem, a coming-together problem that requires a different kind of experience, certainly, than I have

>> AMIR BANIFATEMI: We saw the model which was four quadrants. I am going to invite Trent to probably dive into that framework a little bit better and deeper. Do you want to stay there? I can switch the slides for you.

>> TRENT McCONAGHY: Actually, before we get to that slide,

Stuart's experience reminded me I spent the last almost 20 years building AI software to help design computer chips, so that phone on your desk or in your pocket, inside of that, there's about 15 chips, and every single one of those needs to be designed by an electrical engineering sitting inside Qualcomm, Sony, whatever. When I first started doing this back in the late 90s, I am like great, if we build it they will come. It was pretty naive to think so. It turns out a fast sales cycle is six months. Selling to some of these big enterprises is a matter of two or three or four years. So over the years, we learned to be patient as engineers. We also learned that the value of shepherds within the company was critical. So very much echoing what Stuart learned in his experiences.

From that experience in AI over the past two decades, I spent the last few years in the area of blockchain and ecosystems. The last couple, I have started to see a lot of intersection and seeing some pain points that are coming from the AI world and how blockchain technology can be interesting.

One of the biggest pain points is data and data silos. And many of the issues and problems that many of people have been talking about in the last couple of days here. Sort of from that, myself and my colleagues have been developing a framework towards this, and it's not just us. Out there in the world of blockchains and ecosystems, there's a lot of interesting technology that's emerging that can help to solve or reduce friction in the political problems and so on, some things that are permissionless, some things that are permissioned,

and so on.

So with that and taking as an input the discussions of the last few days and so on, Amir and I and others have been iterating -- and Stuart -- iterating on sort of a high-level conceptual stack of what this might look like. How these parts may interact together to start fleshing this out towards this idea of data commons.

This is a conceptual stack of how this looks. And I'll state from the beginning, so I am an engineer, and this is actually something that is not only buildable but is getting built, actually because it's a critical need.

So what this is, at the very heart in the middle, it says AI commons, this utility network, and this is a scalable platform Stuart had been discussing. It's a decentralized network that no single entity owns or controls it, and it has built-in data governance and a whole lot of other features I will talk about in the next slide. Within that at the higher level of the ecosystem, you've got a few different actors. You have the problem owners, whether for satellite data or tracking cattle like we saw earlier today or the logging or misbehavior of logging in rain forests. So these different problem owners integrating with a problem's frontend or one of many front-ends. And these problems can get stored in these commons, something that no single owner or entity controls.

Then you have sandboxes, frontends and so on, that AI practitioners and problem solvers can work with and build against. So there's various frontends. This is not only existing data science

tools like baizeDB like was being demonstrated Earl area, but also Web-based applications and so on.

Finally, we have the suppliers for data and for computer, and these can be things like traditional cloud providers. It can also be people just making their hard drives available from within the enterprise with appropriate privacy controls. And this is actually not only for a data commons, where the data is free, but also for marketplaces. And it's really crucial that you actually have free data side by side with priced data because one makes the other stronger. If you have the free data, then the price data might migrate elsewhere to their own data silos, and then vice versa. It's important that these sit side by side.

In these AI commons, once again, the core block at the center, within it, you've got data sets there, and this is not only the raw training data, it's also cleaned-up data, the models, the problems, and it's also the sem plats for various sandboxes and so on. No, ma'am only the data, you have -- not only the data, you have also the storage and compute. And the smidgen of AI algorithms that you need too, like Stuart mentioned.

This is a high-level framework that is not only buildable but getting built from various folks in the blockchain ecosystem with a strong eye towards the real-world problems of the SDGs and so on.

Onto the next slide, please, Amir. So it give you guys a feel of some of the features -- and this is meant to echo some of the discussions of the last few days -- think of these as inputs. So

from a high level, decentralized. So multistakeholder, but not relying on any one party to say yes, I will fund this, but rather, many, many people helping to run this ecosystem together. No single entity owning and controlling it. Free and priced. So reconciling the data commons with data marketplaces. A variety of data, so raw training data with all shapes and size $, structured, unstructured, and so on; cleaned data models. Built-in curation. Right? You don't want people just putting in garbage, or if you do, you want the high-quality data sets to bubble up to the top. You want to incentivize when people clean the data so there can actually be curation on this. And finally, incentives to share the data. In the world of blockchains, one of the most underappreciated things of blockchains, they are actually incentive machines. Bitcoin, they are actually getting people to do stuff to add to security in the bitcoin network, to run compute for Ethereum, and you can leverage it as a way to incentivize people to share their data. This is a superpower of decentralized technology.

That's the highest level. Some other features that are really key. Privacy is a big one, for storage and for compute. We have things like GDPR, the general data protection rules, which have a very strong bias towards protecting personally identifiable information. There's ways to reconcile that now, and you can basically -- a system like this really needs to support local storage, where it's behind a firewall, but encrypted, to making that available on a centralized cloud, or even a fully decentralized

cloud, things like Swarm. Also privacy compute. It can be local or on the cloud or decentralized cloud. Leveraging emerging technologies out of many places, including Berkeley, you know, multiparty computation, this sort of thing.

So there's a lot of new technologies that have emerged, even the last two or three years, that are finally efficient and scalable for actually preserving privacy. They were conceived of 20 years ago, 30 years ago, zero knowledge proof, that sort of thing, and finally they are efficient. So there's talk of trade-offs between security and privacy and convenience. Those trade-offs are not as harsh as they used to be, and that's really great to understand.

Going forward, more features of this platform and it's getting built, data governance. So basically, at the core you can have permissions, just like you might have in Google Drive and otherwise, where you can have fully permissionless, anyone can see all the data, light you might see in the bitcoin network. And you can go the opposite where only one person or a handful can see a particular data set or do a certain compute. And you can have small groups forming around sandboxes and so on. So if you have permissioning, it allows you to have global or local -- allows you to have sandboxes, that sort of thing.

Another key thing, there is this idea of knowing who did what, and the key technology behind this is it digital signatures. This is a feature that just comes out of the box when you use blockchain technology. But it's really useful for understanding the neck

thing, which is really crucial towards trust, this idea of trust in the data. Well, in order to trust the data, you -- a starting point is you need to know who put the data in there in a secure fashion, and you need to know the history of that data, when it went in as the raw data, how did the compute happen, where was it stored again as the cleaned data, then how was the model built, and so on. You can actually record these steps in a decentralized permissionless fashion securely, and that's really the provenance.

Finally there's been the talk of various labels, you know, is this data set good for space? Is it binary valued versus a regression problem? Or ontologies and taxonomies, and if you have labels, it allows you to bypass actually having a strict taxonomy, and instead you can have these emergent ontologies. That's really useful because there are some data sets that are useful across the board, things like tutor sentiment. That's good for space data sometimes and otherwise.

Finally, the sort of fourth high-level feature is related to legals and interoperability. It's really crucial that we have this sort of network of networks effect, just like the original Internet itself. So there's a lot of great platforms out there for doing decentralized AI compute, things like singularity net and otherwise, as well as for decentralized data, things like wit net and otherwise. So you really want to be able to connect these together, and that's the work we are doing, for example, for Ocean, for example. A network of utility networks all playing together and interoperating, and this

can come out of the box.

Also IP itself. There are a long history of IP standards. It turns out there is a common standard, something called LCC rights reference model, and this has been brought into the blockchain decentralized world with something called qual IP. It reconciles decades of history of standards with JSON and Semantic Web and all this. Finally, there's discussion several times today about legal documents. So it's really useful to have these template legal documents for IP licensing and otherwise. This can come out of the box as well. So basically, the whole idea here is to make it as frictionless as possible for people to be able to make their data available, to share it, to do compute, and so on in this decentralized permissionless fashion. NPS that was a lot of content. I will stop and breathe and just maybe summarize by saying there's a lot of really big problems out there, and the SDGs are a great way to summarize those challenges, and approaching those challenges at a higher level, we can talk about using AI, blockchain, but drilling into that some more, adding more flesh to the bones, a common platform for the data commons, for the AI commons, is actually something that not only can be built, it is being built, and that's basically what I was talking about today. Thank you.

>> AMIR BANIFATEMI: Thank you very much.

(Applause)

Great job of highlighting the small diagram to this one.

So we talked about AI commons and sandbox in the middle. I would

like to invite Francesca to talk more about the sandbox and give you a different flair of what it means.

>> FRANCESCA ROSSI: Many things have been said already about AI commons and platform, so that's my take on it, my way of summarizing it is that this should be a way for everybody to be involved in creating AI, and possibly not just creating any AI, but creating in my view AI that can be trustworthy and beneficial.

Why do I think this is a good idea? It is because AI will impact everybody. So for this reason, I think that everybody should be involved in having some voice being held in what AI should be and have ways to facilitate those voices to be expressed.

So that's why I think it's really very, very important that everybody should have the opportunity to build some form of AI. Some of them would be very simple, would be just contribute to have an AI system to function well. Some will build it from scratch because they will have more expertise in doing that. But everybody should be involved in that. So that's why I think the ingredients to create AI systems should be made available and simple to use in a kind of tool kit where one can make different pieces, mix and match them, and in a way that should be as compositional as possible.

Of course, description has been said, it matches with what I say here, which is the framework should be very practical, should be very useful for everybody, with real, concrete collaboration mechanisms between the different players. And also between the different

contributors that can come from different sectors, different stakeholders, different cultures. So in order to make sure that the decision of how certain AI system is going to be designed and developed and deployed is not only the decision of one, but a collective decision-making process.

And again, it's good that it's called an AI commons and not just a data commons because -- also because of other reasons. It's because AI is not, in my view, just machine learning, just data-driven algorithms, but AI contains many other techniques and many other subdisciplines, like planning, scheduling, search, knowledge or presentation, reasoning that in many sectors are needed together with the data-driven algorithms and be combined together with the data-driven algorithms, so there must be a way also to not only take the data, curate the data, provide them in order to build data-driven machine learning approaches, but to combine them with these other techniques.

Who should be involved in this? The picture that you saw before involved AI predictioners and problem holders. I also would add other stakeholders. I would add those that maybe do not produce AI but they innovate in AI, so the researchers, the research part of AI. I would also put as some of the stakeholders that should be involved the social scientists because if you are going to develop AI for some problem owners that want the problem to be solved for some community, you want somebody to tell you what that deployment of that AI system in that community is going to mean for that

community. How is that community going to change in its functioning, in its interaction among people? So social scientists are very important to be in the picture because you want to understand the consequences of that system that you are going to design, develop, and then deploy into some society.

And the data subjects, using a term of GDPR, should be involved as well. So we, the people who are going to use that AI system that is going to be deployed that are going probably to provide our own data to that system to be able to function in the best way, we should be also in the picture since the design moment. As well as the Policymakers. As we said before, one of the ingredients to help people to build AI is to help people understand what policies are there, what legal frameworks are there around that AI system.

And again, I just say that all the time that the only approach to make this and many other initiatives around AI successful at this point, it has to be multicultural, multigenerational, multidisciplinary, and multistakeholder. So this is my four multi-something, multi-X that I always said.

So finally, I would like to go back to something that I talked also this morning and we had a whole track yesterday to this idea that I would like these AI commons to be really instrumental in making everybody build not just building AI, but building trustworthy AI, building responsible, beneficial AI. And for that I mean that they should be, you know, pieces of the tools, the toolkit that should help people developing the AI system to understand how to detect bias,

how to avoid introducing unintended bias, how to check whether there is bias or not in the model and maybe how to integrate it, and how to audit or certify that something is biased or not in a certain context in which you want to deliver.

There should be also tools to allow them to understand whether the system is going to behave with the right values that are needed and relevant for the community that is going to be impacted by this AI system and whether there is a need for explainability modules that allow the system to be able to explain why it is making certain decisions, not others, or why it's recommending humans to make certain decisions and not others. There should be ways to help these developers to understand what it means, what kind of explanation you should give, what it means in that context to give, as the GDPR would say, a meaningful explanation, and how to Mr. It.

So to help designers and developers -- how to build it.

So to help designers and developers of these systems to understand what they mean in the context and then to resolve them. So I hope that this initiative, the AI commons initiative, can collaborate and partner with other initiatives that can explore this platform and can also contribute to the platform by injecting, you know, ways and brainstorming and guidelines and best practices on how to develop and so how to frame the way these AI commons will work. So I am thinking, for example, about the partnership on AI, the IEEE ethically aligned design, the TrustFactory.ai that we talked this morning, and others that may have something to say to contribute to

frame and shape these AI commons.

So I really welcome this initiative. I hope we will be very successful.

(Applause)

>> AMIR BANIFATEMI: Well, you have heard from different speakers. We hope that this explains a little bit the foundation. We talked about Project Zero. This could be also a Project Zero that could be led by all of you or including all of you to make it happen.

I would like to turn to Chaesub Lee, who is the Director at UN to have his feedback on the necessity of capacity building and how this can contribute to not only SDG but digital mandate of making AI more universally accessible.

>> CHAESUB LEE: Yeah, as host of this event, I am very pleased to see these presentations as a project in action. All participants indicate that they are looking to take action. As Vincenzo mentioned, ITU is working collaboratively with agencies, how we can work together to realize your conclusions.

Also, give some of my views. We working very hard, holistic pictures. It's very nice. But we do not forget about this AI for bad community. They are working very hard, I think. They do not need consensus. We take that into account. I take words from someone talking about pilot project that will be great idea. We have to move forward but also take into account, Professor Russell, you mentioned about the transition that will be one of the great missions of this summit, how we can ensure this transition is a good success,

that will be a great way forward. So having said that, it's a great conclusion. I am very happy I can deliver this conclusion to our UN agencies. Thank you very much.

(Applause)

>> AMIR BANIFATEMI: Thank you very much for all. This concludes this session. We hope that you all join this new initiative called AI Commons, which is the framework of all of our work. It will be shaped by the common and multistakeholder approach, and looking forward to sharing work on that with you. Thank you very much.

(Applause)

***

***