AI for Good Global Summit session: Breakthrough Groups on Investment and Economic Aspects
June 9, 2017
7:00 AM CT

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

***

***

>> CHRIS FABIAN: Hello, everybody.

I'm going to let the panel introduce themselves. I will start off. Welcoming you to our discussion on KPIs, three letters, which I was excited to look up this morning. -- when I was asked to moderate this panel. KPI is not an API, nor is it anything else that a human has been used to doing the time before machines, but it's way to help us indicate our success in certain ventures.

We brought together a group of experts to talk about how in our respective fields we're looking at understanding the impact of the emerging technological space. That's defined for this panel as -- well, we'll let everyone give their definition of AI. And so our panel, who is interspersed among all of you comes from various expertises. Some from hard academia, some from the space of rigorous social science and others of us from the field of philosophy and literature, which doesn't qualify me in any way to be here.

And we wanted to ask the panelists three specific questions each. Or to frame the indicators in those -- in a few buckets. I think we all agree, it's important after the three days of discussions to be able to have some way to measure the success of this community. Did we all just come to Geneva, as fun of a city as it is and as much opportunity as it provides us for evening entertainment? Did we come out here to hang out and talk or do things? And if we came to do things, how do we measure that doing?

Why are we here? Not as human beings, be but people in the space of AI for good and how do we measure that? But without further ado, we will ask the panelists to SBLUS themselves quickly and I'll go in and ask one first question and we can have an open

discussion with all of you. Thumbs up if that's okay, thumbs down if you hate it. Great, you can all stay.

Say who you are and why you're here. Come and join us in the front somewhere. It's a very casual room.

>> Hi, I'm Sara Jacobs. I run a nonprofit.

>> I'm Joe Konstan.

>> I'm Kyle Nel. I run innovations.

>> He was a fake panelist.

>> Hi, everyone --

>> CHRIS FABIAN: The first and only audience member.

>> Hi I'm Dave Larry. And I also passed teaching and literature and full-time into investments, I'm currently managing partner and a general partner in a couple of investment funds in Canada.

>> SEAN MCGREGOR: I'm Sean McGregor, for another seven days, I'm a Ph.D. student at Oregon State University, at which point I will defend, I'm not planning my dropout beforehand. My project area that I think brings me to be a panelist and reporter on in the session is in my research optimizing wildfire decision making.

So a fire starts in a forest, and you want to decide whether you let it burn or suppress it, if you let it burn, you have less fire risk in the future. You have more ecological diversity. Also, if you let it burn, a lot of stuff will burn. So you need to figure out how to balance those over that time frame. And the way I am interpreting that on the context of KPIs, I suppose, you need to be able to explain why a decision is the best one. And you need to be able to do that to multiple stakeholders who may each have their own KPI when it comes to fire. If they own a home in the region, they may only care about their home burning down, rather than whether the forest is producing timber.

And so, it's I think a lot of the questions we're trying to answer in the course of this conference has similar properties of diversity, of value and what they care about when solving these big challenges and sustainability.

>> CHRIS FABIAN: I'm going to stop you there. But that's a great intro. Gretchen. And you want to slide in here and get in? I'm also super mean as a moderator. It's not because I don't love you, it's because I love you so much. Move in.

>> Hi, everyone, I'm Gretchen, I'm a senior producer at Lowe's and my background is in research and translating insights into action.

>> NIKITA AGGARWAL: Hi, I'm Nikita Aggarwal and I'm a lawyer and academic, and I'm interested in the design of liability around AI.

>> CHRIS FABIAN: I'm Chris Fabian. So I wanted to frame the conversation first and give everybody a chance to respond, and we'll do the opposite order to respond to that framing of KPIs, and

then, I wanted to ask pairs of people a question if that seems right. But first of all, show of hands in the audience, who is, like, super familiar with KPIs? Who has really never heard of a KPI before?

Who is, like, midly familiar and in the room because they want to know more? Who is coming from the engineering side of things? Who is coming from the policy side of things? Who is coming from finance? Academia? Okay. You know your audience. What do you do on the left side? Blue shirt, in the front, what do you do? HCR -- oh, sorry, who comes from the world of development? Who comes from the UN world? That's you. And Sara, what do you do? Which one? Okay. In the back? Hybrid. Okay. We covered the left side. They didn't raise their hands. So I wanted to figure out who was there.

As we look at indicators, we've come to the end of a few days on artificial intelligence. And how do you indicate whether or not we're doing something right, there's a bunch of things. The strategic development goals. How many of them are there? 19? Right? 17? How many indicators are there? 110? 169 indicators. Targets. Target's not an indicator. How many of those? 230 indicators. Okay. That's good. Great. So we can only measure 20% of those target indicators. Yep. There's only 20% of those are tier 1 as Sara informed me. That's my dynamic information processing right there.

Here's the funny thing, I had somebody at a very important workshop the other day. I only look at SDG7. Actually, 7.4. Are you really talking like that already? We just went through this process of defining this whole framework and you're already on, like, number decimal point other number. You lost the point of what the UN does. Which is, we didn't join to do accounting, I hope, or archaeology I really hope.

And we need to look at it differently in how we understand the world around us. As we're looking at the indicators, I want to propose three possible areas. One is sets of indicators about how well we're doing the work we're supposed to do. All of that stuff.

Can we use artificial intelligence and machine learning to understand the change that's happening in the world better? We do something to an education system, something else happens. Did that come from what we did? Or what somebody else did? Did that come from the intelligence or something else? That's one. What are the indicators that we could look at for that? Is this technology actually making a change?

Second is, indicators about the technology itself, the thing itself, is this technology good? Can we interrogate it? Can we look into the tech? This is that middle level stuff. And the third types of indicators, the indicators of the input level. What went into making this technology? The data, is the data biased? Or is it fair? Is it equitable across the spread of humanity?

Is the team that created it a bunch of dudes in San Francisco or a diverse team of people that can solve a problem? I

wanted to go around the panel to see if people had responses or thoughts on that. And I'll give everybody five minutes to come in on that. Why don't we start with Nikita.

>> NIKITA AGGARWAL: So I like how you phased the approach to KPIs in this space. And it's sort of -- there's an input perspective and there's an output perspective. I think something to be conscious of, the output perspective is new. We've already been trying to measure social impact in many fields already. We need to learn from that. And there are a lot of -- there's no one answer, but there are a lot of practices and approaches that are being tested. The other thing to be conscious of, we shouldn't try to over quantify this discussion.

And there's a lot of intuitive appeal to evidence-based policy making. But it's not so easily done because you have a political variable that you can't measure, right? So it's a really important overall goal to sort of achieve targets. But how you measure it might not be so precise as we would like to think if it was a more purely scientific pursuit. But the idea of having three phases is key. And to think about the values in what you put in, like the data, the transparency of the data. And then, as well as the outcomes is, I think, a really good way to structure the discussion.

>> CHRIS FABIAN: Can I ask you to talk about overquantify and what that means in this context.

>> NIKITA AGGARWAL: For me, it is to think about the quality of the indicators as well as the sort of numbers. If we're talking about education -- in the previous panel, someone mentioned the returns based approach to measuring education, at least historically. And what we're not talking about is simply looking at how much salaries someone's earning out of school. It's also looking at broader quality of life indicators. And even the jobs they get, like, how much stress are they under? And how happy are they? And how much well-being do they derive? Please come and join us.

>> CHRIS FABIAN: Can she add to it now?

>> Yeah, I come from a social/political science background. And often times, we see in government, I worked in the State Department for a while. If you pick what to measure, if you pick the wrong thing, you're totally changing the incentive structure around how everyone does their work. And so, sometimes, it's actually better to not pick something to measure than to pick the wrong thing to measure because it totally distorts people's incentives and decision making.

>> CHRIS FABIAN: Awesome. And we'll come back to you for more. Gretchen.

>> And I might add to that. I would agree with what you said. But I think as we do, as we select certain things to measure, we have to be conscious of the fact that there are going to be unexpected outcomes. And that's probably when it comes to AI, the

biggest thing that scares people and why we're having a whole conference about how AI can be used for good, the fact that the speed with which it has the potential to change things could be so huge, and we may not understand the unintended consequences that could come from it.

So as we're, you know, putting a stake in the ground and saying this is what we plan to measure, or if we're not, being sure that we're looking at the consequences outside of those specific things that we are measuring so that we are understanding the full scale and breadth of the impact. But then, at the same time, recognizing that the -- that there's a benefit to not necessarily leveraging the full power of an AI, necessarily, right away. And the need to test, see what the reaction is, what the output is, what the consequences are. And then, scale thoughtfully as we understand that.

>> CHRIS FABIAN: Can you talk about the different speeds? For those of us who may not have been very much involved in that type of user engagement or research, do they operate at different speeds? Are there two different ways of working in those two halves of the brain.

>> I guess speedwise, quant can be instantaneous. But I think the reason both are important because quant can tell you factually is it working against something we're trying to do? Do we like it? Do we like the output of this matters, as well? You know, a plane is faster than a bicycle. But if I prefer to, I don't know, take a train instead and see the scenery, then get there in a minute, then there's benefit there, too.

And as we do roll out AI or, you know, mechanism solutions for AI, we have to consider do people like what's happening as a result of it? And there may be communities or industries or groups of people who they prefer and they're better off in their minds without that solution.

And so that's something else to consider.

>> CHRIS FABIAN: Awesome. Cool. Sean?

>> SEAN MCGREGOR: If I could attempt to talk about two different kinds of machine learning systems that you could conceivably want to measure. One would be to give an example, a system that takes satellite imagery and labels areas as having poverty or not having poverty, and then you could decide whether you deploy resources to those areas to mitigate food security issues and related.

Another one would be a system that directly makes a decision on the basis of that information that the system chooses where the resources are deployed to. In the first one, you don't necessarily need to come up with new performance indicators. It's all ability whether your system is appropriately labeling poverty or nonpoverty from a satellite imagery.

And that's pretty easy. But when you start making

decisions with the system is when the performance indicators get really muddy. And it's uncertain on whether you're benefitting in the world, or not.

And you want to add, too?

>> I just want to push back on that slightly.

>> SEAN MCGREGOR: Okay.

>> If it were 100% correct, we might agree, poverty or not poverty. But if we're labeling and we're 95% correct, we do need to think hard about how do we measure things like the bias? Are we disproportionately finding poverty in certain places RAR than others that may then lead through the system to the idea that certain types of poor communities don't get the kind of support they should when other ones do. Whether that's by race or urbanization or other things.

And so, I'm not sure there's not a hard problem even for the simple AI.

>> SEAN MCGREGOR: And I agree to an extent. I guess, I would just make the argument that that's a hard problem. And when you go into decisions and changing the world, you go from like hard or to ten or a thousand times hard. You have limited capacity to look at the world that would've resulted had you not made the decision. You can verify satellite data on the ground or you could have a person between the satellite data being labeled and the deployment in the real world.

You don't have that if the decision is being made.

>> CHRIS FABIAN: I'm going to come back and ask you guys about data sets in a minute. Awesome. Thanks both of you. David? Should be good to go.

>> That's yours.

>> CHRIS FABIAN: If you want to see people's faces, do that in the first two rows. Come join us here. Just putting it out there.

>> Thanks. First, just to the three, the opening statement in terms of where we're going with this and the focus on measurements. I think it was well-stated in terms of having the three areas we look at and output type of model. I'm not sure -- I guess, a couple of comments. One, I would be -- I'm not sure there's a clear boundary between those. And I think that's something we want to keep in mind as we're looking at it.

And I would just suggest that maybe some of the things, like, things like bias, et cetera, might find their way in between those edges and how we judge it. Second observation would be that one of the things that I think I've heard over the last couple of days, every other -- just about every second presentation is the word exponential, and the thing that goes like this.

>> CHRIS FABIAN: Who loves exponential? Who is sick of hearing exponential. All right.

>> No, more than twice, it's exponential.

>> CHRIS FABIAN: I did bias the results there. Guys, come

sit in the front two rows and you can see everybody. We're doing this differently.

>> Checking us out, first, before they make the decision. But just something to keep in mind. We'll be with that pace, what we decide to use -- if we decide we're going to measure something a certain way, we're going to have to think a little differently in terms of, the measurement, if we start looking at some of these things, which let's say it takes five years to solve two of these. And that five years, two things are going to change. One, the ability we have with these, particularly artificial intelligence within itself.

And secondly, the pace at which the change can occur. With artificial intelligence and its capabilities, perhaps, we arrange an environment or create an environment where the measurement is actually created within the environment. And may actually become an algorithmic progression itself, and that will be a fundamental change and we have to find a new acronym from KPI.

>> I would like to say we should never change the name from KPI. Exponentially. So I -- you know, like everybody, I have no linear path to, like, this position, but I'm a behavioral scientist by trade who got in on the qualitative side who got into neuroscience and built and applied a neuroscience company. And working within a large fortune 50 company making cool stuff, autonomous robots. So we've been doing AI lots of different times. I think the thing that's hard is we as people want to understand why things work.

AI gets us to the point that all of the things that make up AI. Even now, I don't understand why it's spitting out the results that it's spitting out. But I know it's affecting a behavior change that I want but I don't understand why. So there's this chasm between, I have a need to understand. And then, when I go before the board or important people, I am compelled to explain why it's doing the thing it's doing. And I don't know why.

But my rational brain is trying to come up with reasons why. And I can come up with a good construct, maybe a lawyer style. But it doesn't mean it's actually right.

So either have to get really focused on I only really -- my primary concern is affecting a behavior change or outcome. Or we're focused on understanding why that is. But I don't think you can do both at the same speed. And I've seen that happen.

I used to work at Walmart, yes, Walmart, before all of my nonprofit friends murdered me afterwards. I wanted to work within the belly of the beast to see how things would work. I took all of these crazy things I was building, like machine learning at early stage, like deep -- deep neural network stuff. And we applied it to our behavioral database, which is receipt data. And we compared it to the largest thing, which was search engine results from Google. Back in the day, it was about 60 years ago, you could regress data

against theirs, it was awesome. And we were finding stuff like mittens, searches in Saskatchewan correlated highly with orange sales in California. And we didn't understand why. It got to that point. And I had no idea why and I set up parameters, like, sales are going to go up, get ready, and they didn't -- we had no mechanism for that. And we had to get comfortable with acting without understanding. It's already here, but it's going to get more profound and bigger. Those are the questions we need to ask ourselves.

Rather, what is our thought -- what are we really trying to get to?

>> CHRIS FABIAN: Fantastic. You want to come in on that for a second? Yeah. Interject away.

>> NIKITA AGGARWAL: I feel like I have to respond because, you made a dig. I have to defend. I'm not sure whether we're necessarily optimizing for scrutability. How would you -- decipher it. That's right. Decipher the black box. But actually, that might not be our focus. It might be making sure if we go back to this simple dichotomy of input and output, saying you have a representative sample in your data that it's fair, that it's been obtained in a transparent manner and so on and so on, and the system has been built in a way that conforms with certain standards that we have all agreed on and we're measuring whether that system works to produce the outputs we want to target and then you reassess.

But we don't necessarily want, you know, a blow-by-blow description of how that system reached that result. Because that's in my view not really useful and also quite difficult.

>> Let me spoil the great dilemma. Because I've been to both Florida and to Saskatchewan. And orange season is in the winter so is mitten season. And I bring that up because much of my work in the field. Much of my work in the field of recommender systems. And in the early work in the field, one of the points that I and some others made regularly is all of the fancy algorithms, if your goal is a good recommender system, they're useless. Go to a supermarket and draw in marker on the cart, buy bread and bananas.

And in the US, no matter what you do, people will buy the bread and bananas. The beautiful thing, they buy it even if you don't write it on the cart. Because they always buy bread and bananas and you learn that, uh-oh, the system is perfect and it's useless. And we built a lot of useless and, yet, perfect systems. And that's what I'm going to push back at your three-way tricotmy and argue it should be a dichotomy. The idea of measuring output is important, and the idea of measuring inputs is important, but the technology should be nothing more than one more of the inputs.

And when we privilege it as something we measure separately. Oh, the AI is working, I would argue, the AI can never be working. Unless the output is working. And that anybody who allows their engineers and their research scientists to say, you know, hey,

you know, in essence, they're coming back and saying, we've got a great classifier, it's just a shame that it didn't fit in this problem. That's sort of like the surgeon coming back at the end and saying, the surgery was a brilliant success, it's a shame the patient died.

>> CHRIS FABIAN: We're going to do one more panelist. Or you want to let him pop in there? Do it. Go for it. And who are you?

>> Whether an area is poor or rich based on satellite images. And then, there could be another system whose goal is to make decisions that lower the effective classification rate of poor areas. So the second system, if you follow me, could decide or could observe that, say, planting flowers on the roof -- of tricking the other system into thinking an area is not poor.

Externally, the whole system appears to work well. System one, the classification system finds less poor areas. System two makes good decisions, plant flowers, so less poverty. But on the ground, nothing changes.

>> CHRIS FABIAN: We'll come back for the next round over questions. We'll reflect on this idea. I hate KPI so much. Let's call it algorithmic equity. But Sara?

>> Now I regret going last because everyone said my interesting points. It is actually an assumption that a more diverse team can create a more equitable algorithm and there has been studies that show that, in fact, it's true, but it's still an assumption we have going in. We have them throughout. So I found that as long as you're very clear about them, it's much easier to, then, fix your model as you go back.

And the other thing is, just, when you look at this big round matrix of goals that the group tries to move towards. We're talking about how AI can work towards them. I think it's important to think about how sometimes AI's not the best thing to work towards them. And to have a way of thinking about when it does make sense to use this tool and when it doesn't.

And, a corollary of that is also to not let the perfect be the enemy of the good, which is that most policymakers and decision makers as they're working on making decisions around these things are operating with little information right now. Even if you can improve the information a little bit, you may be moving the ball forward even if we don't yet have a perfect model for everything.

>> CHRIS FABIAN: Cool. I want to ask one more round of questions. And then open it up to our audience/copanelists. Thank you for joining us today, guys. The question I want to ask now is about this idea of algorithmic equity. If we are able -- and I'll preface this. UNICEF works in a space of equity. We believe we have models. We do have models that show economically, it's more valuable to invest in the bottom quintile of humanity. If you put the school in the hardest to reach place where the poor Education people are,

that school is worth 1.6 times as much as if you built it anywhere else in terms of the amount that it can change the economic factors and the social factors around it. And that's based on old and kind of basic math. UNICEF calls it equity, which was confusing to me coming from the start-up world. I figured it out.

It's like equity, equity. And I wonder if there's a space for thinking about algorithmic equity and what that would mean. And how this community, and after we ask our esteemed panelists and our other panelists will ask, but is there a space to come around algorithmic equity? That might distribute content to a social platform, media platform. Can we ask that the algorithms equitable in the way that we make building schools equitable? What if you said, take the bottom quintile and get it to as many nodes as possible. Can we make search work that way or some of the data collection actually force an interaction with that bottom quintile? And so I wonder if, as we're thinking about this -- these indicators for the eye of Soron. The SDGs, as we're looking at how we fix the world, if we need to have a position on the integrity and equity of the algorithms that we're building and working on. And if that can be something we can define.

I ask everybody to reflect on what a world would look like where there is that equity BIMT into what we're doing.

>> Yeah. First of all, I think there's a lot to be learned from other spaces in this regard. The development space, initially, was working a lot towards that easy to reach communities. And about 15 years ago, there became a push around this do no harm framework, which is actually a framework now that every development program kind of runs itself through to make sure that there's no unintended consequences that actually would make things worse.

And so, I know we've been talking about do no harm a lot and it's an abstract concept. But in theory you could build something that filters everything that we're doing in the space, as well, to look for those consequences before you start implementing or as you're implementing.

So I think that's one piece of it. And I also think, it goes back to the three buckets or two buckets or whatever that we've been talking about, right? Which is that first you need equity in the input section and it's pretty clear to me, although, this is an assumption, that you can't get algorithmic equity until you have equitable inputs. Then you need to make sure that those inputs are being used in an equitable way in the system and then you have to make sure, also, that the outputs and the outcomes are equitable, which may or may not follow from the other two pieces.

>> I mostly agree. I think it is possible to develop algorithms that are explicitly designed to take in inequitable inputs and improve the equity. Maybe not establish perfect equity. But none of this matters until you have a social consensus about what that

equity is. You were talking about who we distribute information to. What about who we distribute information from? What if Google were to decide tomorrow that instead of giving you the pages that if things are best but are often most popular, it would give you the least read pages on the topic you're looking for.

Whether they're in your language, whether they're relevant to you, whether they're any good. Google would be out of business a lot faster than we'll be done with this session. And there's a reason that that's not what people want. Even if it's a more equitable way to expose people's content to readers, they have the right to exposed to the equity they want.

We also have to think about intergenerational equity. And that's the thing that scares me the most about the rush to AI. I fully believe that if you go out and start creating tools that do a fantastic job diagnosing tumors, that there's a set of people who will be helped.

I even believe that with the right government action that set of people could include to a large extent poorer people and people in areas that today are less served by technology. But I also believe that it will be very hard to do that. Without disempowering the set of experts in radiology that are the ones who need to diagnose the new things that there's no data to learn from. And make it so that our next generation actually has systems that can protect them from the types of cancer that haven't evolved yet, but will evolve over the next 20 years.

And to me, this is the same problem that you get with self-driving cars, I actually believe could work because they could communicate with each other. And I think, people are so bad at driving, there's a chance. I worry that we're breeding pilots who are not ready to take over when the autopilot is incapable of handling it. And we're seeing evidence of that. And I don't want us to start breeding teachers and doctors and others who have been, basically, so de-skilled by having their hands taken off on the day-to-day part of the job that when they're finally needed they're not able to help my kids and my kids' kids to come up with a convenient solution that might help me in my last years.

>> CHRIS FABIAN: Everybody's hitting the two to three minute mark. I have red and yellow cards, usually. You are a beautiful set of humans. Kyle?

>> KYLE NEL: I think we're overcomplicating it. I think every movement since the dawn of time has had an easy to understand narrative or story behind it. Good and bad. It's so complicated, it's hard to assign metrics of success or anything to it because we all have different views on the success criteria are. And a lot of times, they're in conflict with one another.

I think until we have a universal story or, at least, a couple of stories for different parts of where we're all trying to

go, I don't think we'll ever have consensus on what this equitable world view should be.

And in that vacuum, a lot of crazy stuff is going to happen until then. So I think really what we need to do is -- for the next conference, there should be a world narrative conference. Like, what is our group human story? That sounds super corny, but that's how we do stuff at Lowe's. I hire science fiction writers, we give them our trend data, send them out on different time horizons and they come back with stories. We turn those short stories into comic books and hand them out to the executive team and the board and they tell us what to go build. That's how we build all the crazy stuff. There's no reason we couldn't apply that same thing here.

And then, work backwards. Now we've agreed upon, this is the future we're all marching towards, then we'll start to deconstruct what the systems need to be in order to make that happen.

>> CHRIS FABIAN: I'm not somebody that loves science fiction or anything, Star Wars shirt. Don't those narratives exist for humanity? I like this idea. Don't -- aren't those narratives something that -- it's pretty low-hanging fruit. There are a lot of books and multiple books in series that talk about the very dark or slightly less dark future of humanity. Is that sort of what you're suggesting we pull some of that existing? Or you say there's a different one?

>> I think we need a different one that's not based on either an extremely dark zombie apocalypse or terminator 2. It's somewhere in the middle. Star Trek the Next Generation, unless we agree on that. Unless we agree or a big portion of the population agrees that's what we're going to try to do, we're never going to get to the place we want to go. I feel like where everyone's trying to build systems that aren't laddering up to anything of -- like a cohesive value or a defined future that we're trying to achieve.

>> CHRIS FABIAN: The Water Knife is not either super dark or unrealistic.

>> KYLE NEL: Who has read that? If you think about the great movements that -- so, think about manifest destiny, bad, good thing, that whole thing. But think about how crazy that was. 13 plucky colonies deciding all of a sudden, we're going to go coast to coast. It was a belief that it was going to happen and, boom, it happened over time. There are lots of different things where there's an understanding of where we're trying to go. And with all of the political uncertainty. I'm trying not to go philosophical, but without the things, there's not a cohesive vision of where we're trying to go.

Feels like we're trying to manage bad stuff from not happening.

>> CHRIS FABIAN: More of science fiction talk later. Dave?

>> I agree about the next generation, too.

>> CHRIS FABIAN: I got a T-shirt in my bag.

>> What a great panel to be on. I don't expect by the time we get around here that we're going to have one panelist who has the same understanding of what this -- an equitable algorithmic equity is. You know, and I'll use our own company as an example. You know, our theme, our mission is knowledge for everyone by everyone. Now, to me, that's about as equitable as it can get.

Our goal is to provide knowledge for the whole world, to let people create knowledge for the world.

And this is not -- we don't need advertising. This is not what this is about. Just -- we built a company from that point of view. And that's what we're setting out to do. And it is directly related to the beautiful colors up on the wall. I think -- I don't know where we're going to find an answer for that.

But what we're really saying, we've got a lot of stuff we're going to pay attention to and we're going to struggle to find a balance for it. I think understanding that as we move along and write the story is going to be really important piece of the plot line that we keep in mind all the way through as a thread.

>> CHRIS FABIAN: Can I ask you to respond to Joe's question? If Google was to dump out the least viewed page, right? Is there a ground above? Is there a space above that where you could have curated applicable content that's maybe not the most popular by populous numbers? Or popular in the populist community? Is that type of curation something that can be helpful?

>> I think your point was that if whatever the system is not putting out -- the people who are consuming it want, they're going to go out of business. Rather, it's Google doing that. That's an extreme example, but they make a good point. So I think, I mean, given, again, the stuff that we know about AI and other systems, it's possible for you and I to do a Google search and get very, very different results. Because I have a real interest in science fiction and you don't. You know, so I think -- I just don't know how deep that's going to get and how much of me it's going to get into that system and then, that's going to have an impact on what comes out and then does algorithmic equity being an equity of one. And then, it's equal by individual.

>> CHRIS FABIAN: All Kardashian kilobytes.

>> I'm passing this off.

>> CHRIS FABIAN: Got enough mics there.

>> So one of the wonderful things about computer science and artificial intelligence in particular, is there's so much interest from industry that the field has been glow quickly because the funding has grown quickly. The problem, then, though, is that a particular set of research agendas are developed almost to the exclusion of ones that you would more traditionally find in academia.

For instance, there's a research track that I don't think

is very well represented in the summit called fairness, accountability and transparency. We'll say it's a stand-in here for current question. But, you know, in our march of advancing the field in areas that are related to advertising or similar things that have commercial needs, we may have forgotten that it's important to fund these cross-cutting ethical concerns and being able to test for them and to measure them is going to be critical. In particular, because computer machinery systems optimize to things. And if you don't include ethical concerns, they won't care about it. And you'll have to figure out what it's doing after the fact.

I have a dual role of the panel is stand up at the end of it and say what our findings are. Just in the interest of me having something to say, I'll be more formal and say that we should fund research efforts and testing or we're testing is performance indicator here. Of fairness, accountability and transparency of algorithms. And I use that FAT line of research because there's an established line of research that would be fun to research.

>> CHRIS FABIAN: We're going to let you report what we decide to let you report, Sean. Thank you so much, Gretchen?

>> I can say I watched you type it, so I know you didn't come with it in your pocket. No, I mean -- I don't know I have anything that wild to add. I think it's when we talk about equity and I'm probably not using the same definition of equity that UNICEF does or that anyone here does because with all values, we all have a different perception of what they mean and what, you know, that -- and we value different values differently if we place different emphasis on them.

So, yeah, I think that only complicates the solution.

>> CHRIS FABIAN: Let me ask you about the work you do and how you draw out. And this might seem like a basic question, I think it's important as we're thinking about design of systems. How do you draw out the voices? They don't provide data that's like in any poll. May not want to provide data. How do you draw out those user insights? Is that something that could be an open area? That's something that worries me a lot. It's easy to build all of these great, cool things for all of us. And there's a whole group of people just getting -- and the thing with Uber in Florida, public transportation through the drivers, leaves anybody without a smartphone. If you're a poor person without a phone, you can no longer take the bus.

How do you ingest those stories? Is that something you think or worry about?

>> You do think about it and worry about it and you're conscious there are forgotten voices. And you do your diligence to the extent you can to mine the social data or any of the passive information that's being ingested. But then, I think, too, as you are moving into an AI world where there are sensors and tracking and things like that, you have to be conscious of who is open to that

and who is not.

I don't know that I have a very sensible answer.

>> CHRIS FABIAN: Awesome, thanks. Nikita?

>> NIKITA AGGARWAL: It's such a loaded and, like, complicated topic, right? Like, how do you define equity? Whose values? There are so many themes we've been drawing up. Maybe I'll highlight two concerns I have.

One is that, you know, people who are taking data have some responsibility in how they represent it. I don't subscribe to the view that, you know, that there's a neutral platform that's just taking in data and putting it out. There are some values being attributed to those who are managing our data. And those need to be somewhat reflective of the society that you serve, right?

But then, that comes to the second point that there suspect that much universalism in this value set. It is by and large relative. So we also maybe need to be comfortable with some differences and how we look at equity.

And it's not about -- it's not about universality. It's about maybe equality opportunity. It's about fair access. But I don't have a necessary -- like a single answer. And maybe the last point is -- going back to what Sara said, actually, is about not letting perfect be the enemy of the good. I think that in many ways, an aggregate level, we're getting a lot of -- we stand to get a lot of benefit from the deployment of AI and better use of data.

There are people who don't have smartphones that can't use Uber. That's not a good thing. But if you are at a societal level getting progress. And we shouldn't, sort of, let that be halted in an attempt to overengineer the solution.

>> CHRIS FABIAN: Cool. Great. With that final statement, which I may argue with, but I'll do that in a minute, I would like to open the floor up, first of all -- I'd like to thank -- give our panelists a round of applause. That was awesome.

And I know it's super weird for a setting, but yeah, feels good. Feels good to be weird. So we wanted to take a little bit of time now, maybe 15 minutes or so to have some questions from the audience. And so, because I'm a terrible moderator. If you ask your question for more than 30 seconds, I'm going to raise my black flag of plague. If it's a comment, that's cool, as long as it's not more than 30 seconds. How many people have a question/comment, but more questions? Three. Great. We'll do the first three. One, two, three and then I will ask panelists to just pick which ones they want to respond to. Take it away.

>> Hi. I'm working for about ten years on the measurement. At the EU level and the international level. And as you know, I'm involved in the development of AGI. So I have no idea it's going to take. Artificial general intelligence.

>> CHRIS FABIAN: Thank you.

>> I was thinking and it's a question. Not a comment. But maybe I will not manage in 30 seconds.

>> CHRIS FABIAN: I'm watching you.

>> Okay. Thinking about what we can do with AI today -- so not, like, in the far future. There's all these problems, tricking the measurement system as soon as we measure whatever, causation, unexpected consequences, the perceptions of people versus the behavior and the data they're able to get together. I think AI could help us a lot to try to solve all of these problems to establish kind of a system, which would measure the things much more profoundly.

It could monitor continuously the number of data, which we can -- we're not thinking we should monitor, so on.

Second, it could suggest us new data. If you program it cleverly, it could help us to actually identify new data. It could help us with the perception.

>> CHRIS FABIAN: The plague flag is coming soon.

>> Okay. Perceptions, it's difficult to ask people often. So they might be good proxy to identify the proxies. And this should be done where we --

>> CHRIS FABIAN: There's your flag.

>> The values.

>> CHRIS FABIAN: We'll take that as a comment/question.

>> Somebody has to pay for all of this. Who pays and how much? And I think in the real world, this needs to be landed, and yes, it's a complex sector and it's a complex philosophy and agenda. And what came from that in the years and development in the space is a new branch called Pharmaco economics. And if you want to derive value, question, how do you do it? You can't necessarily talk about heuristics, you have to land it in a way in that somebody who is writing a check needs to know how much to write it for.

>> CHRIS FABIAN: Beautiful. Exactly 30 seconds. One question in the back. Question/comment?

>> Thank you. It's a comment. Actually, I joined this session because I was told to come and take notes. But then, I have trouble realizing what you're saying. I think we need to narrow that discussion. But we are talking about what AI could do to help the world evolving this framework.

Let me --

>> CHRIS FABIAN: Coming up on 15 seconds.

>> The UN strategic planning network, all the agencies and all of that have been trying to map what they do. These activities.

>> CHRIS FABIAN: Boom. That is a wonderful comment. Can you sum it up in a sentence?

>> I want to say that AI is useful because we need to benchmark already what the situation is make sense out of data and help to see if we're doing the right thing. Satellite images are very important. Could somebody plant flowers --

>> CHRIS FABIAN: It's a lot of comments.

>> We need AI to make sense. Too complicated what you're saying. Anticipating the problems B uh the actual situation is that there's a lot of data out there and we would need to make sense of that.

>> CHRIS FABIAN: Beautiful. I like that.

>> We need to know what the situation is first.

>> CHRIS FABIAN: I think that ties in very closely to the first question. I think there's a thread there, which I believe many of the people on the panel agree with. Thank you for your comment/question. We -- so one of the things I heard to this point, there's stuff we can do with AI today.

Whatever we call it. So -- and I think this AI term is very dangerous because it can get us lost in a sort of cerebral future search for what will be. Can forest fires -- can identifying forest fires be used to identify something else? Can you use the data sets we have now to make immediate action? And so, I would totally agree with these comments that in that big pie, there's stuff that needs to be done right now. I'm going to make comment one and question three into one and ask the panelists. If, you know, how do we -- if we're looking at indicators that can be shown for immediate success in the next six months, one year, what could those possibly be? And the second question, really, the question of the economics around an emerging AI eco system and what that might look like. Should I go back around the room and have everyone take one or two of those they want? Sound fair?

There's one thumbs up and everybody else is upset.

>> On that economics point, I think looking at the how the governance structure on the internet is important. That's the space where private sector companies who stand to make money from the internet actually are in favor of standardization and therefore are paying a large amount of the structure and standard work that needs to happen around it.

So I think it could be an interesting model. On the what we can use for this now. Of the indicators that have been created for the SDGs, they're divided into three tiers. Tier 1, we have data for, that we can collect. Tier 2 is one that we have proxy indicators that we think are correlated. And tier 3, we have no idea how to measure this. I think that AI could be really helpful in the tier 2 and 1 and 2 space specifically and especially making the tier 2 space, making the correlates make more sense. And I think there's a lot of space to be done in that right now with the data we already have.

But also, those are political decisions that, like, AI can't solve for, right? So only 60 countries that have signed on to the SDGs have actually voluntarily signed on to be part of the measurement plan for the SDGs, that's not something AI can solve for.

>> CHRIS FABIAN: Cool. I didn't want even give you the black flag. It's for them, too. Not just you guys.

>> I want to express support for the idea that AI is a distracting and dangerous term. Mostly, what we're talking about, I'm going to call technology. Remember, the early example of AI was the thermos mug. It would make your hot stuff stay hot and cold stuff stay cold. And knew the difference between what was hot and cold so it could do that.

And, frankly, most of what we're talking about, all of this, none of them have to do with reasoning. Simply learning from data is not about reasoning and understanding. That doesn't make it not valuable, but it does mean that until we talk about things that actually create new concepts all on their own, what we're talking about is technology that we have to deploy to problems that we have to set.

And having said that, the issue of where AI can help us right now is the technology is there to gather lots more stuff. We still need humans to be creative about picking what things do we care about and measure? We can measure prosperity in new ways? Measure learning and education in new ways.

The issue of the economics of this, I think, yes, we need it, it's going to evolve. I don't think it will work if somebody tries to preplan it and impose it. But none of this will succeed until we come up with a new way of thinking about the relationship between technology, data, ownership, and accountability and responsibility. And that's both a new legal framework and a new economic framework.

>> CHRIS FABIAN: Awesome.

>> KYLE NEL: That was good. What do you got? I think the internet framework makes a lot of sense. Closest to where we are with all of this stuff. But I think we have to try stuff. There's so much talking about all of these things. We have to start building some stuff. Some of it's going to work, some of it's not. Businesses are going to try it, governments are going to try it. But only in the doing are we going to learn and got to learn real quick. It's coming one way or the other.

>> CHRIS FABIAN: A quick plug. On our data science team. If anyone wants to get doing with some not AI technology, talk to those guys, they're doing interesting stuff.

>> I think that three-tier, it's a bit of a gift, I suppose. I don't know if it was put in place with AI in mind, I would hope not and I'm sure it wasn't. But the opportunity, I think, is to, yeah, speed tool. We can take more stuff and churn it faster and find different stuff faster. Maybe we can fail faster if we're going to get stuff done now, which is one of the things in the start-up world is a mantra now.

What about the third tier? Maybe that's the opportunity. Maybe that's sort of why we're here. And maybe that has more to do

with economies, whether it's complex programming. Because there are people who really, really know about how complicated and how sophisticated you can do the programming. And people who know an awful lot about the bottom quartile who are desperately completely disenfranchised. Wouldn't have a clue, what's in that bottle? Let alone what's in AI. Maybe in there is the opportunity for this dialogue, which is nice, which is broad, not just one area of technology. Or maybe that's the opportunity to come and really try to figure out. I don't think we're going to know -- I think you're going to report what we have to do for that level. I don't think we know. I'll say we don't know. But I think we have the opportunity to start to try. And I think that might be something we want to start on and encourage in our report back is that you've identified some strength and where we don't. And maybe this is a place to start.

>> Yeah. Thanks. I just want to say, you know, we talk a lot about finding new ways to collect data in, especially, the development space, the government space, there's a lot of data that is just sitting there that no one knows what to do with. I think there's a big component of it that's, like, actually just using the data we have in a faster way. And to the point of failing fast and learning and doing, I think that's right. I think there also needs to be some sort of framework or values that we decide upon that in failing we are not going to do things that could potentially hurt this amount of people in this way. We get in a dangerous space, it's okay, we'll try it without doing some of the thinking first.

>> CHRIS FABIAN: So to try and translate that into something --

>> SEAN MCGREGOR: Go into the findings, develop a framework to be applied to artificial intelligence systems, deployed in the real world. I'm looking for a way to --

>> CHRIS FABIAN: I would say, you know, this sounds -- a lot of the frameworks exist and exist in traditional development space and technology and exist in IRBs. So I actually think it's pointing -- I think the recommendation -- it's adopt or adapt, maybe. We've got the principles of digital development, which UNICEF and 40 other UN agencies have signed off on, including ITU. Take those adapt them, adopt them.

I think it's really -- don't reinvent that but be like stamp, stamp, stamp. Look at that room is full of thumbs is that cool? Or did you want to say something else? What we can do with AI today or who pays for it.

>> SEAN MCGREGOR: I was going to add, we'll hang out for a couple of minutes after the panel.

>> CHRIS FABIAN: Not me. I add zero value to this. Cool. Gretchen, you want to add anything?

>> Yeah. I think I like the idea of, you know, putting things out there and, then, by doing something if we take action,

it will help us all understand implications and I'm sure it will provide clarity on how we want this to go. You know, the values that we have will become more solidified when we see the impact.

You can't understand your reaction to something that hasn't happened or you don't even know will or could happen. So by taking action, then we'll secure everyone's belief set on that. And it may come together, it may not. May divide us. We won't know.

>> A possible outcome of this is to say, you know, that we're recommending we not wait until we have perfect indicators before encouraging experimentation subject to some pre-ethical thought about the amount of damage that will be caused.

>> CHRIS FABIAN: I don't think anybody should go into a discussion, especially, with the world we're coming from thinking you're going to cause physical damage or emotional damage to humans. I think you can have some research damage or reputational damage or that kind of thing for us, that's fine. But I also think there are models of developing a future, instead of future optional value and future value models where we can do small investments and play around with stuff and test it.

And this is, like, what's the future value of the test? If it fails, it won't hurt anybody. And as something becomes more core business and this is the way the UN does the work. Then it's going to have real impact on humanity. There, we want to be more careful. I actually think it's probably there's stages of work. Prototyping is one thing, putting something in a governance system is another. We'd want to differentiate, where is the playing around? And where is the research and thinking plus research and this is now part of a sovereign government.

>> We should recognize, it's all happening anyway. We -- people are being harmed by these systems today. You know, you can pick your favorite example. But I promise you, Uber didn't have a very detailed analysis on how much better or worse the life of the people who ended up driving for them would be. They had a business model that seemed like it was successful and deployed it and trying it.

Some people are better off for it, some people are worse off. And we don't have a framework in this world or at least in most of the countries of this world to stop some of this experimentation from happening. What we do have is the ability to encourage some of the experimentation and the directions we think might actually advance these longer term sustainable development goals.

>> CHRIS FABIAN: That did you want make Uber any more or less of an aspirational company to work at.

>> NIKITA AGGARWAL: I think to add to what Joe said, I think it's important to recognize that frameworks aren't static, they're dynamic. We're not talking about coming up with a set of rules and principles that should always work and that, you know, if they don't

work, we're a failure, you can reassess and you can change. This is a dynamic process. And in the spirit of that and in the spirit of what's already been said, we shouldn't hesitate to run small pilot schemes and contained experiments because you will learn from doing or learn from empirical evidence.

And to respond to the question about, you know, who pays and obviously that's critical. The pharmaceutical industry probably can give us a lot of lessons here in how we approach experimentation. But that's a key question in thinking about liability, especially.

We have existing product liability rules and the question will become how easily transferable they are. I think it'll be more direct in a world of specific artificial intelligence and less so as they become more and more autonomous. And then you have questions of who is at fault.

You know, also, we need, like, AI insurance and that will be the deep pockets that kind of fix it all up.

>> There's an example and it predates the internet.

>> CHRIS FABIAN: What was there before?

>> Coal. But the notion of a common carrier, which is a long established idea that if you adhere to certain standards of neutrality, you're shielded from liability. So the telephone company, as long as it didn't pick what was said, it wasn't responsible for what was said. It would be an interesting thing to try to apply that to algorithms here. That said if you are testably or provably only delivering an unbiased result of the data and configuration that comes in, then you might be shieldable from the liability that of somebody applying that by putting in data and controls and then using the output.

And that might be a way to encourage people and companies to take more effort and innovation on building these neutral and transparent algorithms so that other people can spend their time experimenting on their application.

>> SEAN MCGREGOR: I would be a little concerned over the degree of modelling and training since you can still produce bad outcome even with --

>> The user would be liable, but not the --

>> CHRIS FABIAN: Interesting idea. Yeah. And the other thing you can take from pharmaceutical industry is how they do clinical trials, which is you register a clinical trial and at the end of it, say what happened. And that's an important thing. It means that suddenly we're not all doing our stuff in our own labs and nobody's talking to each other. At some level, you want to play, you've got to put a start and end and marker in it.

>> And you opt in. Whether it's the individuals who are participating in the clinical trial or the organizations who are building these models and creating them and using them internally for their own purposes. There's some acceptance of the outcome that

you can test with this small group of individuals, communities, people that are willing. And then, those who see the outcome can decide whether or not they want to implement.

>> I think the pharmaceutical example is a good one, it's also important to remember that it relies on strong institutions that have legitimacy and authority to enforce. I don't know I would necessarily trust right now a drug developed in China, for instance, that the FDA hadn't approved yet. So I think that as we're looking at a world that institutions, state institutions that we rely on for this. There are models for how you work on public goods, how you minimize bad externalities from companies, but we're living in a world where those tools are going to be different and going to be less strong. How do we design a governance model that applies in that world, as well?

>> CHRIS FABIAN: Awesome. Any other comments/questions? We love questions. No. Just -- if you think we started at 2:00, which would be incorrect because nobody arrived until 2:30, then we're over time. If you think we started at 2:30, then we're dead on time. I'd like to really do a few things. So first of all, I want to take a minute to sum up some of the stuff that we heard and please feel free to boo or yea or whatever or do nothing. But just so that we can have a good caption and Sean can take something back to the repertoire stage. He's got a thumbs up for that.

Somebody shout the most important thing you heard. The thing that came out of this that would stick with you. Audience? Earn your bread. Nothing important came out of the last hour of your life.

>> Go do something.

>> CHRIS FABIAN: One of the things interesting, the ability to correlate some of the work here happening in other fields. I would say, look for indicators or processes for developing indicators that may have come out of other fields, for example, pharmaceuticals and Pharmaco economics. What's another one?

>> Look to other fields for inspiration.

>> CHRIS FABIAN: Particularly Pharmacoeconomics. I was pointing to the guy behind you, but that's good. Start doing. I think the second one I would highlight, there's an urgency. An ability to wait for the whole perfect structure or an ability to act quite quickly on some very easy -- I'm not going to say it. Easy to approach problems that may not be exciting to researchers because they're not super complex, but we can make a big difference.

>> Look for indicators but don't strive to optimize indicators, that's important.

>> CHRIS FABIAN: How can we say that in another way?

>> Metrics are not targets.

>> CHRIS FABIAN: Metrics are not targets.

>> Otherwise the system will self-optimize towards the metrics, but not towards good.

>> CHRIS FABIAN: Perfect. Yeah. Good. I like that. Beautiful.

>> You can't evaluate against the measure of the system is designed to optimize for.

>> CHRIS FABIAN: I like it.

>> And if others agree, at the top of that list is -- I call it findings rather than -- we got rid of actionize. We got food for thought. Anything can be a finding. The first finding is the performance of the AI is not what matters, the performance on the application and the separation between the internal evaluation that's used to help you tune a system and evaluate how to build one has to be separated from the external evaluation of outcomes that look at the impact of that system in its context and usage.

>> CHRIS FABIAN: That was clearly expressed. That's four. Sean's got his -- you got one, too? Yeah, it's very much what he said. Very much. It's a tool. Cool. Well, I will take that comment as our -- we can figure that out. I don't want to keep people a ton of time past there. And we will make sure the -- guys, audience, I want to thank you for spending such an amazing hour and a half in this room. Panelists, thank you for being wonderful brains and all of you, thank you for coming to this event. These three days mark the beginning of a lot more discussions. Please give a big round of applause to all of yourselves.

(Applause)

And get out of here.