FINISHED TRANSCRIPT

ITU AI FOR GOOD GLOBAL SUMMIT
BREAKTHROUGH GROUPS ON INVESTMENT AND ECONOMIC ASPECTS

GENEVA, SWITZERLAND
9 JUNE 2017
14:30 CET
ROOM H
INVESTING FOR IMPACT WITH AI

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

*****

*****

>> ASH FONTANA: I'm just going to get started. People or may or may not trickle in. We may or may not have more. Who knows? So the title of this session is: Investing for Impact with AI. A distinction being it's not impact investing, and it's not just investing in stuff other than AI.

The background of this session is -- I think we all have a moral imperative to ensure a certain quality of life for everyone on the planet. Now, that's very hard when you have on the one hand populations increasing, on the other hand certain natural resources decreasing. So we're presented with quite a significant optimization problem. Now, thankfully, artificial intelligence technologies and slightly more accurate machine learning technologies are optimization technologies so there's an opportunity to use them to solve for this particular moral imperative.

So that's what we're going to talk about today. Practically that means we'll talk about opportunities to use this technology to solve for this moral imperative and we'll talk about risks and other considerations.

Now, we'll use examples in energy and food or feeding the planet.

The aim of this session though is not to talk about those opportunities in particular. It's to have some grounding so we can come up with general principles for investing so that gets to the goal of this session. The goal of this session is to come up with up to three guidelines for people who are looking to invest in these technologies for maximum impact. These guidelines have to be high-level enough such that people in fields as diverse as policy, impact investing, for-profit investing, practitioners, machine learning practitioners can use them but specific enough so that we can actually develop something that can be followed. The structure of the session is we'll have three relatively quick talks that will sort of try to get you thinking about this topic or get us all thinking about this topic.

Then we're going to have a series of questions broken down into sections. If we've got a lot of questions, we'll ask them as appropriate, and depending on the flow of the conversation, then we'll try to come up with some preliminary guidelines. Then we'll seek some feedback, and then we'll try to refine them and we'll hopefully lock them down for presentation later on. I'll quickly introduce the panelists. I won't do extensive introductions because you can read about people, if you want to look further.

I'll start here. Christian Steinruecken studied a whole bunch of areas, studied a whole bunch of things within computer science and then in machine learning, and indeed did his thesis on a language for quantum machine learning and is now working on a company called Invenia which optimizes the flow, helps optimize the flow of electricity on the power grid. He'll talk about that.

Sven Strohband has again a very long history in computer science, has worked for a lot of the -- for some of the car companies as well in that field but now is CTO of Khosla Ventures and CTO of another venture firm.

And David Hunter worked in both in the military but also in the finance industry, and then studied reinforcement learning technologies at Oxford, which is a branch of machine learning and now is working on a company called Optimal Agriculture that is building the world's first autonomous agents to operate a greenhouse to grow food, and he'll talk about that.

Ankur Vora is with the Gates Foundation, and leads strategy and operations in certain areas there, and previously studied economics and all sorts of interesting areas within that.

Rigas is helping us today by reporting what we're going to talk about and making sure we capture everything.

And François at the end has a really interesting background understanding complex biological systems which is obviously relevant to food production but also weather, and all these other things we'll talk about today.

To get things started, I'll start with Sven or Christian, whoever

wants to go first. Christian, you want to go? Christian will go first and give us a bit of a kickoff talk to get you thinking about power grids and how they work and how we might do something interesting there.

>> CHRISTIAN STEINRUECKEN: All right, so I'm going to start with a question. So we live in an age where we have a really -- an abundance of interesting challenges and these challenges are not just personal, they're global challenges, and my central question is going to be: How can we make machines most useful to humans?

So my name is Christian Steinruecken. I'm at the University of Cambridge and the CTO of a company called Invenia Labs, and this company essentially is one of the examples of how we can apply artificial intelligence technology to the environment, and I'm going to tell you more about that in a little while.

So you all remember these, right? So these are development goals, and one of the big realizations is that of course, it's very difficult to make headway on these, but AI can actually help us with at east some of them and so my hope is that I can sort of show a few of these examples perhaps where AI can already make a difference today with today's level of technology. Electricity is one such example and electricity is a special case I guess because electrical grids that we're probably all familiar with are something that we all rely and depend on. They're of crucial value to human life, and they're really critical in infrastructure so if something goes wrong in a grid if there's a blackout for example, a power outage, bad things can happen. It means human life are lost. We critically depend on that sort of infrastructure.

So naturally it's interesting to think about how -- we don't think much about it. We sort of take for granted there's electricity all the time and a lot of what we do is powered by it, so the question is: How are they actually run? There are many challenges to actually operating a grid, and some of them are that the generation and the demand must balance at all times, so they must balance and if they don't then bad things can happen. Some of the generators that we use, for example, coal power stations or nuclear power stations need advance notice in order to generate the power. They can't turn on in a second. They need time to prepare. There are some generators like solar farms and wind farms that generate whenever the sun is shining or the wind is blowing and they're not necessarily ready when we need them so there's a planning problem and that planning problem is hugely important to get right.

With good planning we can have a reasonably well operating grid. With bad planning we can have all sorts of economic waste, increased pollution such as CO2 emissions but also particulate emissions and in bad cases even blackouts. Clearly we want to avoid that it is just very difficult to plan well. There's so many unforeseen things. It's not just about the physical complexity of the system that we

can simulate and extend with physical stimulation, it's about the unknowns that come from humans. We don't know to what extent humans are going to consume power or we don't know when someone drives into a transmission line or something goes wrong in the grid but that happens all the time and AI can help with that sort of thing.

It's an example of how automated decisions combined with machine learning can actually make a difference in that electrical grids can be run more efficiently. It is possible to act on many kinds of problems before they even occur, and we can reduce economic waste to the tune of billions of dollars in electricity grid and lower CO2 emissions and pollution all of that at the same time if we get it right. This is an example of how AI technology today can make a difference and it doesn't really cost anyone anything. Simply by making better decisions we can have a good impact on the environment, good impact on the planet, and it benefits society as a whole.

So that is something that I think is very important to bear in mind. It can be a win-win to deploy this sort of technology. It doesn't necessarily have to be only a financial win. It can be also an overall win for the planet and that's I think what I want to stress.

So AI can be used to solve big problems, global size problems, and that's an opportunity that is really up to us to take up. So let's do it. Thank you.

[Applause]

>> ASH FONTANA: Awesome. Thanks, Christian. I think the thing to take away from that is it's these opportunities exist to solve fundamental resource allocation problems purely with software without any capital investment and without taking anything out of the system in sort of a non-winner -- in a situation when there doesn't have to be a loser. So Sven now will talk a little bit about a company that's using robotics to help produce more food. And again, these are just talks to introduce you to some of the areas and give us some food for, well, some examples to use later on.

>> SVEN STROHBAND: You can say "food for thought."

>> ASH FONTANA: Sure.

>> SVEN STROHBAND: Test, test? Perfect. Okay, so I think a lot of this panel is going to be about resource management and how to distribute resources. I wanted to take a slightly different tact. Back in the day I used to build self-driving cars so I'm kind of partial to autonomous systems that move and so in our portfolio there's one company that I thought might actually be interesting to discuss here. This is not about resource distribution. This is actually getting more out of your resource without causing environmental harm.

So here's a picture of a tractor, and what this tractor drags over a field doesn't look particularly spectacular. It's some white

shroud, and in this white shroud there are 16 pods and these pods look at the ground and they identify each and every plant that they see and then they act on each and every plant that they see. I'll tell you a little bit of what that does. One of the biggest inputs you have when you're a farmer is your seeds, and also your chemicals that you will use on the field.

One big portion of the chemicals are herbicides, so whatever you grow, you will unfortunately have some parts of the field that is occupied by weeds and the common way in the Western world to get rid of it is to use so-called broadcast chemicals, so the way this works is you genetically engineer a food crop. This food crop is then resistant against a broad-based herbicide, and then you just spray the entire field with this particular herbicide which is going to kill everything hopefully but not your crop. And so the most famous example of that is a product called round-up from Monsanto, for example. So all the intelligence here is in the genetic engineering, and the associated herbicide.

But in AI, we can start moving this intelligence into a computer, and we can make the computer make the decisions of what is a weed and what is a crop. And so in this picture that you can see, you can see basically the vision system targeting the weeds, and leaving the crops alone, and it actually tells you how certain the computer is at any given point in time.

And if you look very closely to the right hand side, you can see that some of the plants have been sprayed with something, and some of the plants have been left alone. And I'm not a farmer, so I can't really tell you exactly how the weeds are supposed to look like, but the things that got sprayed were the weeds, and the things that got left alone were the plants.

And so what does that do? So it has actually a bunch of effects, but the biggest issue that it solves is it makes the input cost for the farmer a lot lower. As you might imagine if you have to spray an entire field versus only the plants that are actually required, that makes a big difference in the total amount of herbicide that you would actually spray. It actually turns out, if you spray plants directly, you don't even have to use herbicide. You can actually use fertilizer because fertilizer in high dosage to a plant is actually toxic.

That means you might be able to get rid of fertilizer altogether and the other thing that kind of breaks this dependence on the genetically engineered crop and herbicide combination, so this means you can use all kinds of seeds all of a sudden, and you can use all kinds of herbicides including high dosage fertilizer, for example.

And I wanted to show you, I have no idea if this is going to work. Can somebody click on this video? Because I can't with this. Oh, perfect, thank you.

So I just wanted to show you this in action how this actually

looks like. Can you see the spray that just kills the weeds? This is actually really low-speed. This is half a mile per hour, if I remember correctly. The next one you're going to see is I think 2 miles an hour, and we go 4 to 6 miles an hour over the field, so, yeah, this one is 2 miles an hour and it's from about a year ago now so this is a little bit how this actually looks in the field so it's kind of speedy going over this and this is just one part so multiplex that times 16, it's basically however in the tractor can do.

Why am I showing you this? I'm showing you this because this already exists and it works and farmers use it so when I got invited here, one of the thoughts, we have about 50 companies in our portfolio that have lots to do with AI and we have about 220 portfolio companies and I was trying to pick a company where I thought this base technology is not only useful for a farmer in the U.S. growing soybeans, for example but this could also help in resource maximization in the developing world, because it breaks the sort of cycle of having to have genetically engineered crops and herbicides and if weeds become resistant, for example, to the herbicide, for us, it means one week of retraining. If you have to genetically engineer a new crop that's not a one week affair but that will take considerably longer than that so we're also a lot more adaptive, a lot more adaptive to crops you might want to grow in this region because identifying them again is a retraining process. The physical system is exactly the same.

So I just wanted to give an example of AI technology that is nicely productized already, but it isn't used in the scenarios we talked about today so I'm happy to expand upon that later.

>> ASH FONTANA: Cool. Thank you very much, Sven.

[Applause]

I think -- .

I think certainly in terms of progression in the conversation something to take away from that is how we can use these technologies not just to allocate actual physical resources better, but to reduce negative externalities from the use of those resources so many things to take from it but that's certainly one thing that we'd like to introduce into the conversation later on.

Now, David's going to take it up a level, and talk a little bit more about trying to develop a model for thinking about investing in these technologies. So I'll let him take it from here.

>> DAVID HUNTER: Hi, everyone. Does that work? I'm David Hunter, the founder and CEO of Optimal, and I'm going to talk about mental models. What is a mental model? It's a structured framework for thinking about or making sense of some particular real-world situation. And they're used a lot by really great strategic thinkers and we use mental models in our day-to-day decision making. The example is the 80/20 rule.

It sounds kind of obvious to use these models but so few of us use them systematically and deliberately for decision-making and really the use of these models can make all the difference.

So given the goal of this talk, I wanted to think about mental models for helping us to think about how to invest in AI technologies in order to maximize the impact for mankind. This actually is a problem my co-founder and I faced when we started Optimal. We wanted to apply from some of the machine learning techniques we'd been researching to maximize or improve human nutrition, and the problem we faced was how do we allocate our time and our resources in order to maximize our impact on human nutrition?

And one mental models which we found quite useful was to think about the higher-order impact of technologies. What do I mean by that? I mean the impact created by the second and third-order, higher-order effects of technology you're creating.

And specifically, to seek technologies with increasingly positive higher-order effects in order to maximize impact.

So a good example of this is the work we're doing within high tech greenhouses. We're optimizing and automating hydroponic greenhouses as an example of large scope in Holland and the first-order impact of our AI technology is reduce the energy consumption within existing greenhouses. This sounds great but farming is a lot bigger than indoor farming, so why didn't we start with infield farming? For that you have to consider the higher-order impacts. One higher-order impact is you can reduce the cost of producing food in this way and you increase the amount of greenhouses that we deploy and use.

This is really interesting because hydroponic greenhouses use 80 the 95% less water than field farming with far less nutrients. You can grow with zero or minimum pesticides. You can grow in all climates and deploy these things out sides cities and towns and pretty much eliminate most of the waste within the food supply chain.

Another higher-order impact is if we have more widespread use of highly controllable indoor farming environments and we can precisely control the growing conditions of individual crops and actually optimize nutritional content or phytochemical content of the food that we produce.

So considering higher-order effects are hugely I'm important. Many complex real world systems the higher-order effects dwarf the first order effects. They often have opposite desirability. This is particularly important in AI and machine learning particularly in situations where you're applying machine learning for optimization. It's easy to get seduced by the first order effects which tend to be some kind of reduced resource consumption.

It's very easy not to think about the higher-order effects. But

if you don't think about the higher-order effects, you're stuck with that sort of thinking you end up drawing the same conclusions as everyone else and you end up with kind of not as much impact as you could have created.

So, yeah, in summary, a good mental models we've used is to seek technologies with increasingly positive higher-order effects in order to maximize impact. But, yeah, I should caveat this by saying one word of warning: All models are wrong and you need more than one model. So what I'd like to do is a little bit of a mind meld.

So we've got very smart people in this room, very diverse backgrounds. We've all got our own ways of thinking about things, in particular ways of thinking about how to invest in AI in order to create maximum impact so what I propose is you get a little bit bored during the session, or afterwards, email me your particular ways of thinking about things, your particular mental models, and I will promise to collate these and I'll send them back to anyone that contributed. So please contribute. This could be the highest ROI action you do today. Thank you.

>> ASH FONTANA: Awesome.

[Applause]

Thanks, Dave. So that's a good way to kick off this discussion and what we're going to aim to do today is quite frankly make this a little bit more opinionated and direct than a lot of other discussions. We're getting down to dollars and how dollars should be deployed. So we're going to run through a few different categories of questions. The first category of questions is going to be about opportunities. Where are the opportunity spaces to put money? The second category of questions, and I'm going to sort of kick that off with Ankur, the second category of questions is around the types of problems: What are the characteristics of problems most suited for a lot of capital to be put towards them today. And I'm going to kick that one off with Sven. I haven't told everyone I'm doing this.

The third category is about risks and possible constraints we can put around technologies that can be used for these purposes, and I think Christian and Dave, I'm going to kick that, and François actually, I'm going to kick that around with them.

And then the final category, and certainly at certain points we'll open it up to the audience, as well, because as you can see this panel is not particularly diverse, where we'll talk about policies, and those policy discussions, thinking about ethical frameworks, thinking about how to interact with governments, they will lead into the construction of the guidelines, which will be the concrete output of this session that will carry through to a report written after the session.

So I'll kick things off with Ankur. So Ankur, I wanted to just start with you and talk about the most significant societal problems that are yet to receive significant investment in this field, and in

your position. In the Gates Foundation you have a lot of capital, tens of billions of dollars to deploy in a responsible way, and focus on some of the world's biggest problems, and you've already solved a lot of those problems in health care and education.

But particularly in this field, where is an opportunity to use artificial intelligence technology, and we won't get definitional about what that means, to solve some big societal problem?

>> ASH FONTANA: Right. Thanks, Ash. So the -- I'll stick to the construct we've got which is about (?) -

[Captioner cannot hear speaker]

There are a set of problems we're trying to tackle. I'll focus on SDGs 3 and 4, so the mortality goals and the [(?) goals and for both those goals we've made tremendous progress. We've halved the number of children who die under the age of 5. The number is (?) dying under the age of 5 each year and that just seems ridiculous and it's ridiculous because nobody dies of these things in the developed world. They die of things like a quarter die from malaria. A quarter die of diarrhea, the solution to die Ray yeah is salt water. It's on every street corner and --

[Captioner cannot hear speaker]

[Captioner standing by for audio]

>> ASH FONTANA: -- to achieve one of those goals. Perhaps opening up a fairly controversial area, it is underinvested as in we have underinvested in technology that can optimize human labor and health care organisations. Perhaps that's underinvested because there's no corporate incentive in most parts of the world, ex-U.S., to innovate in labor, in optimizing labor in health care.

This sort of leads to my next question which is: Why are corporates investing sufficiently and where are corporates investing insufficiently either in this field or in other fields? As someone who may have capital to allocate to these technologies you don't want to double up. You don't want to invest where a huge corporation is already investing a lot of money.

So what areas do you stay away from and you think have sufficient investment? And perhaps what areas don't?

>> ANKUR VORA: If I just stick with the same example I'd probably go slightly different which is where are the areas we could partner with the private sector a lot more to extract more value for the causes that we care about? There are probably a few things so if I stick with the same example (?) companies are investors investing in algorithms so I don't have to invest in sort of getting the predictive algorithms smarter because people are using it for various purposes trying to figure out who to market a product to so that's being done. We have to figure out how to copy that into places we care about so that's great. We can do that.

Companies are investing collecting data for commercial purposes so that's great. One of the things we want to require is (?)

mothers and their children, I'm going to require -- we're going to require patient-level information, student-level information, and actually it's a really hard thing to do but -- companies are collecting a portion of that information. They're getting demographic information, the household income information, the -- information which is useful for the -- so one thing we could do is say, how do we make sure that the (?) companies can get the commercial value? Because -- but how do you make sure that they do provide access of that data -- sorry, this is better -- access of that data to the UN system, to the foundational system to people who care about these problems so that's one place.

The other place is that companies are great at collecting information, but they don't have an incentive of collecting a set of information that we'd want. We want to know student by student learning outcome information. That's something they're just not incentivized enough to do so this is a place where somebody like the Gates Foundation loves to play a role in that the risk return does not make sense for a private sector company to invest in it. How can we subsidize it? How can we change the economics for a company to say: I will invest in it as long as you make me whole in the (?) way. So that's the way of the Department.

>> ASH FONTANA: I'm glad you introduced that to the discussion as well. Certainly as we sat thinking about guidelines we've got to think about guidelines not just for what systems to build, but how to feed the data into these systems or where to get that data from, and for that, we need to look to where these data sets are being developed by both private organisations and also public organisations.

So just to sort of round out a way to try to understand the opportunity space, where is private investment going today? So perhaps Sven or even the entrepreneurs in the room can sort of discuss through their conversations with other investors where they seem to think private investors are really excited today. And the point of bringing this up is to really understand again when not to double up rather than just investing in the same spaces everyone else is investing in so I don't know who wants to take that one but where private investment seems to be going today.

>> SVEN STROHBAND: Test, test? Perfect. I don't think it's quite fair to just pick a particular area. I think the demarcation line is actually a little bit different. If I just talk about the areas, we are actually quite interested in health care and we invest a lot in AI in health care. But we don't actually invest as a commercial DC into the problems that the Gates Foundation would have a particular interest in and we actually work together with the Gates Foundation on some things, but so there are these factors that might overlap, like health care for example. We do care about, but the thing that is different is that the application area is just

very, very separate and it has a couple of consequences so to go back for example to this farming thing, we're happy to invest in the technology that can be productized and widely deployed in the large market because that will result in the large valuable company and so hence there's enough capital flowing into this and there's a lot of capital in AI these days so that we will just do but the adaptation of this particular technology to a problem in the developing world is something that is tricky.

A, the economic incentive is different so investing in it is tricky. The other thing is very often the CEO or the founders don't all know that ecosystem, so they don't even know exactly how to adapt it, and very specific to AI, we don't have any data sets, so we don't actually know even if we didn't have the skilled problem and there would be a large economic opportunity there's still a data set problem, so for example, there we don't even know which crops and we don't have millions of images of the crops that people care about and the weeds as they're prevalent in that particular region because it turns out weeds are actually Regional, they're not globally the same and we're happy to solve this problem like in Australia and then in the U.S. and so on because they're large developed markets and it works but the good news here is that it doesn't actually take that much more to make a technology like this useful in the developing world. What it does take is somebody who actually knows the ecosystem really well and our companies could partner with somebody to do that and it takes because it's AI, all the hardware is the same. It takes a different data set to retrain on and we are actually quite good at retraining on novel data sets. We know exactly what we need to do. It's just we don't have a data set like that. So given those two things, we could actually start deploying things that look like that in the developing world. So similarly in health care, almost all of our health care investments are geared towards the U.S. and maybe the European market.

But a lot of the techniques really carry over and the product itself might actually carry over, too, if the training set were actually specific to a training set in the developing world but in the private sector there's very little incentive to gather that data set and there's very little understanding, and I'm certainly on the forefront of not understanding, what exactly is the correct adaptation for that particular market. So those are in my mind the two big things that prevent AI to be more dominant there. Not necessarily a particular Sector.

>> ASH FONTANA: So that leads to my next question, which is around: Let's start trying to develop a framework for understanding where machine learning is applicable. So what are the characteristics of the type of problem that you should look for to see if machine learning will actually solve that problem? And so Sven, you've mentioned to me offline, you know, certain things that

you think about when making an investment in a company as a private sector investor, that you look for to make sure that the problem is likely to be solved.

And Christian or anyone else, feel free to jump in here when sort of trying to think -- trying to help people develop a framework for investing in these technologies.

>> CHRISTIAN STEINRUECKEN: Yeah, so maybe I'll say something here. So I think one of the mentions that's important is machine learning can be deployed essentially where humans are often deployed when it comes to making decisions and especially when it's a problem of the scale where it exceeds the capability of the human, and anything that is driven by data or driven by mathematics or fast decisions, machine learning is likely to be able to help, but one of the bizarre things I guess is that the private sector investments cannot solve all problems and that's simply because the incentives aren't quite right for all types of problems.

So for example reducing pollution without having some sort of economic driver behind it is going to be difficult for private sector investment. So there is a necessity for Governmental international bodies to take action too and maybe leverage the existing ecosystems for startups or industry to incentivize positive effects of that nature.

I think it's -- I think of AI as an enabling technology really that is not unlike the electrification of cities that has enabled a whole range of different technologies, and AI is a bit like that. The question is also: How can we make that technology accessible to countries all over the world, to people of different levels of education? How can we make the benefits of that technology widely applicable and widely accessible?

>> ASH FONTANA: Okay. Any other sort of ideas from the panel around what makes a good problem that's solvable by this? David perhaps you've got something to say here in that you've been on the journey to find the right application for the technology you're working on.

>> DAVID HUNTER: Yeah, I think a couple of points. One, just to go back to Sven's point, which is about, you know, high scalable deployable products in large markets, part of this is about resource allocation. Allocating money that's already there and is trying to be allocated in different types of AI companies and figuring out how to do that so I think a good question there to ask yourself is: Which companies are developing technologies for rich customers in developed world where that technology could be replicated at low marginal cost in the developing world?

So that's an interesting way to think about it and I think Sven said the same. To these companies for example what we're doing is we're understanding crop biology algorithmically so we're trying to augment human level decision-making process around farming and

growing crops and it's easy to spin up an AI farmer and deploy them in a developing world country and although we may not be selling that to a rich customer, there are a lot of charities that want to invest in this kind of thing and allocate their money to these kinds of projects, so that's a good way of thinking about things.

>> ASH FONTANA: So, François, do you think that's a realistic tractable problem space? That is, trying to model complex biological systems to the point where you can have an autonomous agent control certain elements in that system? How far along are we there? And is that even tractable?

>> FRANÇOIS TADDÉI: I think one of the key characteristics of biological systems is they evolve. They mutate and they start exploring other parts of the space than we are used to. And typically what AI is very good at is analyzing data from the past and inferring from it what may happen in the future in categorizing if everything is like in the past then the prediction is going to be good enough but if suddenly you have something that is radically different it's going to be very difficult for AI to deal with and so we're not there yet. Maybe we come because if you look at it, nearly (?) AI software is based on biological iteration somewhere. (?) so maybe we reach one day the level where AI is biologically quite enough that it can mimic our biological system in clever ways but we have 10 to the 4 bacteria on the planet and 10 times more viruses so those are huge numbers and being able to predict the (?) they're exploring will be very, very tough because we can analyze evolution of the past but we're very bad at predicting tomorrow. We try engineering this and that but we constantly have surprises when we do this, in simple microcosm within laboratories so if you start doing this in the wild, you want to reengineer an ecosystem, you might be somewhat successful but you might be surprised (?). So that's one point.

Another thing that biologics can tell us is male and female (?) very differently throughout the developing world and so I'm ready to invest maybe the most important thing I have here which is time so if any woman wants to jump on the stage and share, I'll share may time with her so she can tell you what investing for impact would mean because typically women tend to invest very differently their own time, which is the ultimate resource that is (?) for every one of us and also when they have money they start investing money very differently for the family. All sorts of movements have shown this.

So (?) minorities, women aren't even a minority, they're a majority but if we start inviting people in the minority of decision makers to start investing themselves, we'd probably see very different way to invest and very different type of impact. So talking about economy and investing there's a (?) in the economy that says one of the best Public Policy investments is investing in the very youngest one. Typically women understand this very easily

but Public Policy rarely do. Do. We may have (?) free child care from birth is very rare so why don't we do this? Even if we show this is the best possible investment? And can we have AI for these questions?

So why might they be crying is a question (?) have, even males tend so have those questions especially at night, and what should we do? But there start to be software that are taking this -- you can record the voice of babies and if you have enough (?) to record you can start making a typology. Why do they cry? Is it the same cry as a hunger cry or is it a pain cry, is it a whatever? So far those AIs are very humble results but I guess the more we do, the better we become.

And I think the other type of investing for impact would be AI is somewhat good at optimizing but what if (?) optimized? AI (?) for this so it's for us to decide and again we have to invest in being able to decide what we want to invest our time or energy or AI into.

So I think investing in the (?) of human and artificial intelligence is probably one of the most interesting investments because we can play on the (?) --

[Captioner is having much trouble hearing this speaker]

-- unique strength that AI is very far from having. (?) not just human intelligence but starts to be the collective (?). How do we do this and how do we (?) can we invest in defining what are the best ways forward for the (?) for human and artificial intelligence for individual and collective agents at large?

And I think those questions coming from the (?) intelligence on this planet kept rising.

[Captioner is having much trouble hearing speaker]

If you're able to predict, you invest more wisely, but are we (?) to predict? If you want to predict temperature, molecules can do it at the nano scale but if you want to predict temperature at the plan tear scale, 100 years from now, you start needing many more forms of intelligence so we need more and more of this but I think the strength of artificial intelligence so far is the ability to reflect on what this is all about which AI so far is not doing. Maybe one day. We're not there yet. This consciousness about the fact that of course we're (?) but we're lucky enough to be alive and we have the responsibility that comes with consciousness. Individual and collective consciousness and what we want to do with it so I'll end up maybe citing a poet T.S.Elliot. Maybe you know this poet which says, where is the life we have lost in living? Where is the wisdom we have lost in knowledge, where is the knowledge we have lost in information? Where is the information we're losing in data? And part of this AI discussion is, how can we reverse T.S.Elliot poem and go from data to information to knowledge to life, and see if there are many (?). So far we don't have AI

for wisdom and if we were to have it where would it come from?

We all have some piece of wisdom. None of us has all wisdom, but we have distributed wisdom. Can we progressively start sharing this? Can we build (?) technologies and (?) technologies that will help us take the best of everyone and share it and can we learn from each other? If we go back to (?), was saying coming to him these three ways of becoming wiser. The first one is learn from your mistake. The second one is to think deeply enough and the third one is to imitate the wisest people around you, (?) the wisest people around you and how do you identify? Which is one of the big problems because we tend to copy all sorts of things on the Web these days.

Before the Web we tended to do this, so how do we do this? And can we build collective intelligence and artificial intelligence 2458 help us even investigate those questions by sharing what are every one of our mistakes? Are we ready to share our mistakes? Because AI is going to make mistakes but maybe AI can learn from its mistake or build the AI that can do this. We have to train the people to do the same. We have to think deeply enough about these issues and we have to build the technology to share at scale in order to promote this higher level consciousness and wisdom.

>> ASH FONTANA: Yeah, that contribution certainly took the conversation to another level. It's a very broad ranging sort of commentary there that introduced some pretty interesting elements to this conversation around collective intelligence, and really in goal-setting really. So perhaps, David, you had something to say on that and then I'm just going to open it up for an opportunity to -- for people to explore areas they've been thinking about and see what the panel's thoughts are on the tractability of those. But David you had something first? Speak directly into the microphone.

>> DAVID HUNTER: Yeah, I had a very quick point. That was some great points raised. I think one of the most important ones to me was about really taking the time and how do you close this knowledge gap between the people who really understand AI technology and the people who understand the main problems and this kind of thing? It's -- there's a lot of misconceptions around what AI can do and specifically machine learning, and what machine learning cannot do.

And I think some way or think about some way to close that gap is really important. Just as a very simple example, we can now train a machine learning computer vision algorithm to recognize and diagnose cancer within x-rays, and other medical scans to a super-human level. So that's pretty amazing.

On the other hand, we -- it's just on the edge of technical impossibility to get a robot to fold a t-shirt for us. So just think about those two things and what would naturally you would think of as being simpler than the other, if I hadn't told you that.

So I think it's really important to try and think about how to close the knowledge gap between what is colloquially understood as

the capabilities of AI machine learning and what actually is possible.

>> ASH FONTANA: Cool, and just to wrap up this Section on sort of opportunities and trying to define problems or find problems that are worth investing in or likely where an investment is likely to yield some sort of positive outcome, if anyone in the audience is thinking about applying this technology to a particular area, feel free to bring that up at this point, and perhaps the panel will have some thoughts on whether or not there is a likely solution to that.

So if anyone does have anything that comes to mind, feel free to raise your hand and bring it up right now. And then after that, we're going to move on to risks and other considerations of applying these technologies.

Yeah? Go for it here, front row.

>> So I'm (?) and I run a think tank in India, so I'm much more on the finance side, and I came to look at the confluence of AI and finance, and therefore this section. So what I didn't find a lot over the last two days were specific examples of where AI is helping the social problem so the Blue River technology was one of the few ones that were expressed, so my big challenge is, if I could generalize I would say global north is where more of the money is, global south is where more of the issues are so my question actually is on these technologies. Does AI have the capacity from how to take an issue and the social entrepreneur trying to work on that issue globally? It's an optimization of a website or something.

But like the solution have to reach the poorest corner of the world and maybe they'll adapt it with a frugal methodology. But I think the whole AI for Good, good means you must -- I focus more on SDG-1 so that's like elimination of extreme poverty, so a second aspect was just referring to the AI for prosperity panel. How do we even identify the 5 million children? Is there a way that AI could do that? One we were just saying then everybody will have to have an identification number.

So in India right now we're launching an identification card, are there benefits of that and are there any AI solutions for that? Those two questions of can we kind of get everybody all working on one particular problem together? And the second, how do we identify people in the world with problems?

>> ASH FONTANA: Yeah, so maybe we'll start with the second I'm not sure if Christian was going to comment on the second one which is sort of sociographic modeling using Government data, what opportunities are there to do that? Christian, did you have something to say on that, or the first issue?

>> CHRISTIAN STEINRUECKEN: Yeah. I think that generally, one of the big problems with AI is that it's not immediately applicable without first getting it ready so you need to have the right substrate for it to apply and that means that that is one of the

reasons why for example, applying it to x-rays where we already have digital imagery is for current technology more or less straightforward but to apply it for something that is purely social where we don't have data readily available is going to be more difficult so there's going to be some real human effort or even Governmental effort needed to make it ready to provide the sufficient infrastructure to make it possible to deploy these sorts of technologies.

So I'm not sure it's fundamentally an intelligence problem in the first instance. There are other problems that are very important that need to be solved first.

>> ASH FONTANA: Yes. Sven and then Ankur?

>> SVEN STROHBAND: I have kind of a yardstick that I use if I do a new investment in an AI company and maybe this yardstick can be useful here. And the yardstick is you have to have a clean articulated data set that you have access to, so in your case, it would be if you're trying to find all these people you really would need to know them. This is your data set. And you would need to have a nicely labeled data set, who is interested in what.

The second thing is typically I like to have lots and lots of humans in a large chain that make decisions, and then you have at the end of it, an optimization function that you actually understand. So you know exactly the behavior you want to drive, the outcome that you want to have, how many false positives you're willing to tolerate, and you can articulate the economic value of this optimization function.

If all these three things, then I'm basically interested in looking at it, because it has a high probability to actually work. If one of those three things isn't there, it is hard to solve, because you might have to do some groundwork first to actually solve that problem.

One of the things that I've become a fan of to make AI applicable in situations where you might care about is if it is highly productized in a different context and only requires a minimum amount of retraining, because at that point, it might be feasible to get a novel data set, and then maybe it's an NGO that finances the collection of that particular data set, maybe it's the UN, whoever finances that collection, but everything else has already been solved: The product is completely productized, it's produced in very large quantities, it's therefore quite cheap and has at least a shot to be used in a scenario where the financial constraints are more rigid.

So maybe that's helpful, so I'm a little bit more on the very, very applied side of this, because I think scale matters in almost all of these things.

>> ASH FONTANA: Yeah, Ankur, did you want to add to that?

>> ANKUR VORA: I was going to go to your question number 2, which

is a great question, and I think everybody struggles with it, at least everybody -- people I talk to. The two ways people have thought about it, one is the one you brought up, which is India's unique ID is that sort of the basis of doing more things? So EMR in the U.S. or EMR in other parts of the world, people voluntarily sign up to get an ID. You do the unique ID and have a tag against it so there's a system around that.

The other one which I found interesting when I heard about it is that you could, if you just aggregate a little bit up, so for SDG-1 if you care about poverty, World Bank does the service called SMS and they try to figure out who's poor, who's not poor, et cetera, et cetera. With AI now we have geospatial maps. We just take satellite pictures and then you train your maps, your satellite maps, or the machines, and the predictive power of these maps is significantly high. In fact, people claim that it does better than the survey does in terms of trying to figure out which household is poor or not and it also does better in terms of census numbers.

So when we try to figure out when we're doing a polio campaign and we want to know where children are, there is a census route you go and you figure out where the people are. And the other route is you use AI and the geospatial maps. Actually that does better than the census data that we have. So people have found ways around it, but at some point if you do want to do the identification of which child is at risk, you're going to have to get to the unique ID.

>> ASH FONTANA: Yeah, so data collection being very, very important, fundamental. We'll take one more sort of request for feedback from the back. Yeah?

>> Hi, I'm from ITU Statistics. Correct me if I'm wrong, but what I seem to be getting from the discussion is that one of the main obstacles for investing for impact with AI is basically information, and the countries that could use the impact really are data deserts, on many levels, from official statistics to the data that is fed into AI systems.

And it seems to me that one potential solution is a platform that connects local entrepreneurs and local people on the ground, because sometimes it's not just about impact investment, really. Good old investment is really what you need in many of the least developed countries. Sometimes it's local people and how you connect it with people who have these grand ideas and how to use AI to make money and solve problems basically.

So do you think that sort of works? To what extent is that feasible? How can that be done?

>> ASH FONTANA: Yeah, that's a fantastic idea because there are private companies out there, like there's a company called telemobile, there's a whole bunch of companies that lend money, write small microinsurance policies whatever else, these are privately funded for-profit companies, that are collecting really

interesting demographic, sociographic and psychographic et cetera data on populations in countries where the Government isn't doing such a good job at collecting that data.

And again, just thinking a little bit, 10 minutes forward, about the guidelines that could come out of this session, it could be to create a fund, for example -- this is to your suggestion, just sort of trying to coalesce into a guideline -- could be to create a fund that gives grants to private companies, if they collect certain pieces of data, and make that data available through an API to others in that country, whether they're Government organisations or other private organisations or whatever else, and just giving that grant could start sort of a data flywheel where you get some data points, you can combine them with other data points, et cetera, et cetera, that you eventually build up a data set which could be as good as a large scale Government data acquisition project. I personally think that's a fantastic idea knowing how far ahead some private companies are but making it clear what data points are needed and then providing some sort of financial incentive to collect them could really spur data collection because a lot of these private companies will just see that as another source of non-diluted funding for their company and obviously it has positive effects in terms of what can be built to identify those poverty or otherwise.

I'm not sure if anyone else had any ideas around that or any excellent on that really good idea.

Cool, we'll take one more from over here.

There we go, it's on.

>> I'm from a China based company and invest in projects that have both value and high impact for society and one of our representatives is a bike riding company. We were lead investor. We have got already 100 times return based on their current marketing capital.

So our question is that we have already did lots of investigations of China based AI startups. Most of them are deep learning based but based on our review we have these conclusions that it's very unlikely if not impossible for any deep learning based startups to give us 100 times more return, so that is our opinion. I'd like to hear more about your opinion based on the U.S. market both from the investor side and from the startup side.

>> ASH FONTANA: I don't quite understand the question but Christian might have a response.

>> CHRISTIAN STEINRUECKEN: Yeah, so I think one important thing to bear in mind about deep learning is that it has some fundamental restrictions. One is, for example, that once you train a deep neural network it is very difficult to introspect what is actually happening inside that neural network. It makes it very difficult to build things on top of the network that the neural network produces

so the question about the value creation, how can you leverage that sort of technology? There's different kinds of technology to make it easy to introspect and build something more durable and neural network is oftentimes an end product. It's the last stage you use in the stack of machine learning technology.

In terms of the question of how can you make sure that you'll have a good return, I think an analogy would be just because a company uses electricity doesn't make it a good company and just because a company uses AI doesn't make it necessarily a good company, too. I think there's no magic bullet like that, right? So...

>> SVEN STROHBAND: I think there was a time when you could invest in a company just because it says "AI" on the cover. I think this time is now done, so this is at least my philosophy. We invest in a solution that might have a very big AI component as part of the solution, but it needs to solve an actual problem. And that problem better be valuable to lots of people, and then if you do that, you can have a good return.

So for example we do this in computer security. We do there in agriculture, as one of the examples. We do this in medicine. But you need to solve the real problem, so just saying, well, this is a deep-learning company doesn't actually tell me anything about what this company actually does and which problem it actually solves.

There was a way I will admit of investments that were just general AI companies, but I don't see that really going on as time moves on.

>> ASH FONTANA: Sure.

(?).

>> -- companies in China and they were developing next generation of (?) camera system for face recognition and smart surveillance and they're competing with incumbents in the surveillance system providers and the problem they face is that they cannot collect the data faster than they do but that's -- but that problem is satisfied given how much money is poured into it. The bigger problem is the second one that is not quite servicing other domains which is about the service team, and the incumbent already built up a very sophisticated National-wise service network such as each client, they have this dedicated team for all kinds of service, which for startup companies really hard to view that kind of dedicated (?). So that's a two part problem. First is solve the real problem, but they need more data.

Second, even if they have more data they need customer service so I'm not sure if you encounter similar kind of two problems in the United States.

>> ASH FONTANA: Yeah, sure. I mean, I think on the commercial side, you'll see that again and again in terms of startups trying to compete with incumbents.

I'd like to move the discussion on for about 10 minutes talking about risks and other considerations we should have in terms of constraining the application of these technologies, the allocation of natural resources, or to solve big societal problems and then we'll try and coalesce on some guidelines to take forward and be concrete output from the session. Yeah?

>> My question belongs here, but talking about principles for investment, I'd like to kind of take from what François said. Are we going to talk about the need for investment from those people on whom the impact is to be felt? Is there any way of aggregating actual needs of the population that we are trying to influence? Because we all seem to be talking top-down. How about using AI to aggregate needs, to aggregate those things which are of extreme importance to the populations we want to affect?

Is there anything being done on this? Or at the very least, I think one of the principles for this kind of impact investment should be taking care of actual needs of actual people, aggregating them in whichever way possible.

>> ASH FONTANA: Yeah so I think that absolutely fits into the next phase of this discussion, which is what do we need to keep in mind as we make these investments? So one is measurement, and what is the actual impact of these investments on the people that it affects? Another is risks and constraints.

So we can start if anyone has any ideas to kick us off, otherwise, I can start, but on either of those two points, either measurement or risks, we should be thinking about. Sven?

>> SVEN STROHBAND: So I only have one part about this, which is: If we're using AI to distribute resources, there will very naturally be some parties that will get more of them, and some parties that will get less of them. And if we trust AI to make the resource allocation problem, then at the very least, assuming we're doing it, the AI does it more efficient than the humans used to do it, then one of the constraints that ranks very high for me is that the system needs to be inspectable, and in particular, the optimization function needs to be public.

So if I want the AI to favor Group A over Group B and that's my objective function, I can certainly train an AI to do that, but that should be a decision that is made on a political level in a society. That should not be some sort of dark black box that is completely uninspectable.

So I think any AI that makes resource allocation questions that affect millions and millions of people needs to be inspectable and we need to be able to reason about the objective function that we give this AI.

>> ASH FONTANA: Yeah. Christian?

>> CHRISTIAN STEINRUECKEN: I think the question you raised is probably one of the most important ones, and it is very, very

difficult to solve, so AI as a technology has fantastic promise, but it does by nature come with barriers. It's up to humans to decide how to apply the technology, and that's a really big and difficult issue, and it needs to be solved, and I think to answer your question from earlier, there isn't currently a good solution to this, and to echo Sven's comment here, it's really important to make sure that we build a tools necessary to control the technology.

One of them is transparency, that we have a way to look inside the box of what we're building, and to understand what it means and what impacts it will have on the environment. The other is we need to give the right incentives. We need to make sure that as on a political level but also internationally we give the right incentives to build technology that will actually reach the communities we care about and help people globally. If possible, we should favor technologies that have a good impact for our future, for our climate, for our environment that reduce O2 emissions that reduce pollution, that help with social justice. These sorts of things, these are difficult problems to tackle. AI is a tool a transformative technology that can help us but it's not solving it for us at this stage.

>> ASH FONTANA: Yeah, and I would even go a little bit further and make a stronger statement which is I don't believe there's any machine learning practitioner today that believes it's possible to build unbiased machine learning algorithms. They are inherently biased. You have to set weights, you have set functions at the beginning of running a training process or otherwise, which necessarily introduces a bias.

Now, that bias can be in the form of a heuristic about how something may work or it can be in the form of a goal you're trying to achieve. So because of that, I think that measurement is obviously very important, transparency, observability, decomposability, and not just important but are crucial, because if we don't do that, we have absolutely no hope of understanding whether or not why certain goals have been achieved for the people that this technology impacts.

Yeah, we'll go to François at the end and then back to David.

>> FRANÇOIS TADDÉI: I'll try to answer your question, there is no perfect solution as was stated before but there is a few hints that are I think interesting. And it has to do with crowdsourcing and citizen science.

So I'll give several examples. One of them is in order to understand for instance the complex nature of galaxies, and there is so many pictures of galaxies these days. There's so many astronomers don't have to do it. They launched a seasoned science project and invited people to identify galaxies, and they had better machines and many more, so they use AI to use the -- that analyzes, produced by the citizens to categorize and speed up, and then what the

machines could not do the citizens would do. What the citizens could not do the scientists would be do. You had this free change of cooperation between expert intelligence, collective human intelligence, and artificial intelligence so that's a nice example of the revolution. It exists for galaxies, and the information about galaxies where enough people contribute but it's also true I'll give another extreme science perspective.

So in the Congo forest where people don't have access so far to even literacy, they came with some technologies. They are typically Smartphone like, and designed for people to be able to map their forest and the threat to their forest in terms of who is polluting it, with cutting trees illegally, with hunting animals, with mining illegally and so on. And they contributed to map collectively their forest and help NGOs that could launch international campaign to defend the rights of the people living there versus all sorts of poachers and so on. This is yet another example.

So there is ways of today inviting people to collect data about what they care about. So say pollution, air pollution, some pollution or whatever, it's quite easy for citizens to contribute to do this. To give an example, in China, even the Environmental Agency wanted to reduce pollution.

So what they did is they invited citizens to contribute to identify pollution source, and then they gave something like one month warning to companies say: You have one month to come up with a solution for this. If you don't do it, we'll make this public. And then no one will want to buy your products anymore. So you can create all sorts of things that people care about, and try to collect enough data, and create the public fear and the public debate sort of the (?) of today where people can come up with ideas and start discussing it globally. This is just the beginning of these type of things and I think it will be evermore and with the cell phone that you can hack by building an app, you can start building maps of all sorts of things you care about and build a political movement so that your issues will be addressed.

>> ASH FONTANA: David, did you want to add? Sorry, go.

>> The population I'm thinking of, they do not have cell phones and they're not on the net. I'm not thinking of people who walk around with cell phones. I'm talking about the people at the very basic level. I'm sure you'll understand what I'm talking about, coming from my own country, India, that is, and there's a lot of the population which does not have this kind of access, so how do we connect them? How do we connect to what they think?

>> ASH FONTANA: Yeah. Go for it, Ankur.

>> ANKUR VORA: I'll do it very quickly which is I think one way we can do it, I agree with everybody that it's actually quite hard because it's hard to reach these people. It's hard to collect the right information. But if you expand the definition of what

information you need from the ground, I can give you things where people don't tell you but you can pick up the information from other sources.

So for example, in the presentation at 1:00, Eric was going to talk about cholera and we actually have models that can predict based on soil information, based on what's happening over there on where a cholera epidemic is going to happen. I can put diagnostic on malaria and I can figure out where a malaria outbreak is happening.

I can put a camera in a health systems supply warehouse and figure out what is out of stock or not. So I can collect all of that information without asking people, and then I can figure out what the right thing to do is so if we expand the definition of collecting information from the ground, then there are ways to do it.

>> ASH FONTANA: David, did you still want to make that comment?

So given we have 10 minutes left, I'm going to try and coalesce all of this into some guidelines to output from this session.

So rather than going full crowdsource mode on this, I'm going to suggest two, and then ask everyone for some feedback.

So these guidelines should have the word "should" in them, so there should be some sort of imperative in them. They should relate to one of the UN or Xprize organisations in terms of it should be a directive for them to do something, and again, they should be sort of broad enough that they can be applied at the policy level but narrow enough that they can be applied for the sort of people we're talking about that are perhaps allocating capital and solving some of these problems.

So, one is a fairly straight forward one, which is that the UN should create a regulatory effort to mandate transparency, observability and decomposability of any models that allocate natural resources. So I'm keeping it -- I'm suggesting that we keep it to physical resources.

Obviously observability and decomposability are concepts that would need to be fleshed out, but I think some practitioners would be very helpful in doing that. So that's one that I would suggest.

Another that I would suggest to the lady from ITU's point is that the UN, or I'm not sure if the UN is the right organisation for this, but should incentivize private organisations to collect population-level data points, and now we can point -- one idea I had here is we can point to data points or surveys currently run by UN or UN-related bodies like the World Economic Forum or UNESCO or something else and there should be an incentive programme for private organisations to collect those same data points and the economic we can put on that incentive programme is prance linked to the current budgets for those surveys. That could be a way to do it. I'm not sure.

I'm obviously open to feedback there but they're the two I would

just put out there, for feedback from anyone. And we can -- we're allowed up to three so we can either delete one of them and add two, or so on.

So any feedback? Any extras? Any others? Rigas?

>> RIGAS HADZILACOS: Just to give -- having heard a lot of very cool ideas also from the crowd, just to put some more out there for feedback, I would like maybe trying to create some citizen science projects and enhancement through international organisations to collect a lot of data on issues that we care about and the people care about. Obviously something that could be added is for those places where we cannot have citizen science to try to see if we can have some students doing citizen science on satellite imagery to collect the data in areas that we could.

And another one that came in terms of could we push for, as an international community, push for allowing minority voices in decision-making to also be included in the decision of where investment goes in this area? I thought that was something that would help.

And the last one is back to the comment from the lady in the first row, can we develop an AI tool that can be a matchmaking process between the local, Regional, global challenges, and the solution providers, the scientists and the researchers?

>> ASH FONTANA: Yeah. Thank you very much for making note of those suggestions along the way and bringing them up at this point.

So feedback. I'll take some more from the audience there. Yeah? Sure. Weave got to get through the guidelines but we'll take the question if it's heading in that direction.

>> I feel obliged to answer and I came a bit late so I apologize if this was already covered but I was just curious to know two things. One is, do you foresee also AI technologies to help for the impact investors to source fields and also to scout for the (?) especially in the early stages it's very much human focused so would the AI systems help you count humans throughout the planet so it helps you identify the high potential ones you want to invest in and also think about investing in the future in 10, 20, 30 years down the line, what does impact investing look like?

I find it a bit ironic that impact investing has been sitting on the front line of innovation disruption, it's been operating quite the same way for the past decade so operationally speaking will it be materialized the same way? Will money be a prevalent part of the equation? Or is there anything else that's on your minds? Thank you.

>> ASH FONTANA: I think given that we have to come up with guidelines in 5 minutes, we don't have the capacity to open that avenue for discussion. I'd certainly be happy to talk about it afterwards.

I didn't give any background on what I've done and why I'm

sitting here, but I started the business side of AngelList, which is the biggest fundraising platform in the world, and we certainly had some ideas around this, so I'd be more than happy to talk about it afterwards.

So what should we go with? What guidelines do people have a preference for? Feedback on them? Starting with David.

>> DAVID HUNTER: I think a few times people were talking about matching problems with solutions. So really in sort of the developed world, a capitalist society that's the role of the tech entrepreneur, understand, get to the cutting edge of some sort of technology, really understand some problem in great detail and figure out how you can use the technology to solve the problem. In the developing world, where there's not necessarily rich paying customers and large markets you haven't got that role fulfilled.

I guess traditionally that role is fulfilled by charities and this kind of thing so I really like, it would be good to flesh out maybe a little bit that guiding principle around connecting problems which is the point you raised, I think with technologies in developing world and I think you sort of mentioned about possibly, about creating a fund or something so maybe we could sort of have it defined about what it might be.

>> ASH FONTANA: Or even a repository where organisations can post datasets and try to explain the problem or what they're trying to extract from the dataset and let others work on that. I was on the Board on a company called Kaggle which did this for some public good purposes to solve some social problems like diagnosing certain medical conditions and also private purposes as well so insurance companies would put up a dataset and say we're trying to predict this. All the Kaggle members try and build a model, certainly building something like that, which is François's suggestion, as well. François, Christian, then Sven.

>> FRANÇOIS TADDÉI: You mentioned we had a platform that would maybe group the various suggestions we heard because for me they're complementary. I don't know what we call this, fourth SDG or whatever so you want to know what issues people care about, want to know data, and we want to know solutions and we want to know who is good at what, and you want to know who can contact for, and if we had people that could put together all these various dimensions from, you know, the problems to the solution, going through the data and the delta of the data in the sense of we implemented the solution, know we implemented this solution, the problem was reduced by this much and we can measure the efficiency of the impact, and then if we want to invest we say well, these things reduce this problem by this much by investing for so much, and it is in that area under those circumstances. Would that be scalable? Would that be applicable elsewhere? If we have open data, open solution, an open network of problem-solvers and if we create incentive for

people to join, I'm sure some entrepreneurs and companies would be interested but also students.

The students would be learning a lot from such a database and they could be learning to be problem solvers themselves and they could start contributing to their community and maybe build an NGO or a company wherever they are to solve for X, whatever X is.

>> ASH FONTANA: I think that's a fantastic idea. I would just clarify one thing, or at least make a suggestion that would clarify the incentives on the platform, which is to post a problem, you also have to post some corresponding data. You have to contribute that as a poster of a problem. Because then you avoid the situation where you have lots of problems but no actual work able to be done. Just to your points about students and student participation, for what it's worth, Kaggle, the biggest such platform for commercial purposes, has a huge involvement by the education communities, universities, high school students, a lot of people go on there to learn how to do machine learning because that's the only place they can find open datasets to work on. Were there anymore suggestions here? Otherwise we might be able to coalesce on guidelines.

>> CHRISTIAN STEINRUECKEN: It's also important to note different people are motivated by different things and not everyone has necessarily monetary incentives. Plenty of people are happy to apply their skills in NGO or for other purposes so I think it's important to guidelines that will incentivize to harness these sorts of contributions so the key thing AI requires typically is first of all the knowledge in the base technology but secondly data so if we can have guidelines that incentivize countries or society to make data available, high quality data ideally that would be fantastic.

>> ASH FONTANA: So what about adding an additional one, which is -- I'm not sure what the appropriate organisation is -- should support the development of a community where Governments can submit problem statement, or non-profits can submit problem statements and associated datasets, and a community of machine learning practitioners can work on those problem solutions and datasets, something like that.

I don't want to draw a direct analogy to a private company but Sven, do you have something to add on that one?

>> SVEN STROHBAND: I think it's a great idea. It's not entirely obvious to me what it has to do with investment necessarily, so in general, I'm in favor of it, but a little bit more completely on the investment side I think there is an opportunity. AI is very, very heavily invested field from the private sector in companies and also folks that look like the straight-up venture investor. And it would be a pity if these products are not applicable in the developing world, because they haven't ever been adapted, and to adapt them, there's two things. I need to adapt them to the correct problem, and I need to get a dataset that makes it specific to that particular

problem. But hopefully 90% of the product can remain the same.

And so one idea that's kind of rumbling in my mind is: Should we have a fund that's basically, its only purpose is to take the best of breed technology that has been funded by other folks for similar usage, let's say in U.S. hospitals, for example, and basically try to just adapt it to that other use case, so that the investment in adaptation is actually smaller, because a lot of the stuff has already been built, but it is very, very specific then.

And then this fund would go for two things. One is what I would call -- and the company would call it product management, so it actually fits the intended audience. The other one is specific data gathering that is required to retrain the AI for that specific purpose.

>> ASH FONTANA: So how would you put that in a sentence that would be a guideline? X should Y for --

>> SVEN STROHBAND: Okay. There should be a search fund for develop AI technologies that adapts these technologies to problems that exist similarly in the developing world.

>> ASH FONTANA: Yeah, great.

>> SVEN STROHBAND: And the two tasks are, one is a product management task and somebody has to do it that understands the ground reality, so certainly not me but somebody who actually understands that, and the second thing is data collection costs money. And it requires money to do that. And to me, that could be a fruitful investment, because it leverages, well, all the other money that quite naturally is pouring into AI right now anyway.

>> ASH FONTANA: Cool. So I think we can arbitrarily just decide on three right now, because we've got to end.

So I think we'll go with the one about creating a dataset platform for where people can submit problems and associated datasets and the community of machine learning practitioners can work on those and to Sven's point about what that has to do with investment just as an aside, I think any solutions that get traction on that will attract investment. We've seen this on other platforms, as in any solutions that are achieving a certain level of predictive accuracy will attract further investment in and around those problems so I think we'll go with that one.

I think we'll go with the one about decomposability and observability because it's fairly obvious and crucial and I think the third maybe is this idea of a search fund.

I can certainly imagine a whole lot of image structure that could be built around the search fund. It could hire people good at the functions, product management functions and data acquisition functions. There does need to be a more explicit bridge between organisations that are willing to share part of their intellectual property or application in the developing world and people with knowledge of those problems in the developing world.

I'm not using terms that everyone's going to agree with but I'm just trying to wrap it up. Rigas, do you want to clarify anything?

>> RIGAS HADZILACOS: Just because I will be reporting back to the Plenary, and the one about data traceability and openness has been raised by every breakout group up till now on different topics. I'm wondering, I think that the one that you formulated before on the other fund that could give incentives to private investors with data, I would suggest that I open up by saying, we also came to the same conclusion as everybody else that data needs to be traceable and open.

>> ASH FONTANA: Models yeah.

>> RIGAS HADZILACOS: But for three more concrete --

>> ASH FONTANA: Let's leave that one out.

>> RIGAS HADZILACOS: I will mention it in the opening but it will be on the --

>> ASH FONTANA: So the third one then instead of that would be to your suggestion, which would be to create a fund that incentivizes private organisations to collect data that is currently collected through UN and associated body surveys, which is just leveraging the work that all these organisations are currently doing.

Cool, all right, thank you for pulling us up and helping us avoid being too generic.

All right, thanks, everyone. Thanks for your contributions.

[Applause]

[End of session]

*****

*****