Raw file.
May 16, 2018.
7:00 a.m. CST.
1400 CET.
ITU.
AI for good global summit.
AI breakthrough teams.
Room C1.

(captioning test).

AI breakthrough Teams.

The Heart of the Summit.

C: AI for Health Promotion, Prevention and Education.

>> Good afternoon, everyone. Welcome back for those of you that you were here earlier, and welcome to the ones who were not here in the morning.

This afternoon's session is devoted on promotion, prevention and

education, but also it will have an interesting component, on ethics and policy, with a presentation and hopefully discussion. Before I get to that real stuff, I want to introduce myself. My name is Effy Vayena, professor of bioethics at the Swiss federal Institute of Technology in Zurich. This morning you have two of those, I am in the one in Zurich.

My own work is very much on the issues around ethics, governance and policy of data and their uses, and if I were to briefly tell you what I do, I think I'm trying to do two things. One is to try to change a little bit the perception that when we think about ethics and policy, that is something that is there to stop development and technology and that is a hard sell. But I'm trying.

The second thing, because I think if we are to stop the right things from being done then we would have a problem with kinds of ethics. The other thing I'm trying to do and I think meetings like this are very exciting and inspiring, is to try and think that innovation is critical, not only in technology, but what we also need to accompany and enable this innovation in technology is innovation in ethics and policy and governance. These are two themes that enthuse me very much. I try to work on some very interesting, looking through some of the examples I'm trying to do that very thing here.

We will keep the order of the programme, we have one panelist who is coming later, I think. But let's start with Lukasz. I will

ask the panelist to introduce themselves, so we don't waste time.

>> LUKASZ KIDZINSKI: Hello. Great. Let me introduce myself, I'm a computer scientist, a techie from Stanford. There are many good things about being a computer scientist. I can earn a lot of money. I can, people think that I'm somehow smart because I know mathematics, and people like me because I can fix their browser, their printer or anything they like. But there are also bad sides of being a computer scientist, which might be surprising to some of you.

Whenever we encounter a problem that is beyond our mathematical expertise, we are sometimes lost and don't know what to do. We are frustrated. Many of those problems come from health care. I come from a family with some cancer history. Whenever you encounter a person like your close relatives are terminally ill, you realize that your mathematical framework is completely useless and you are getting frustrated that you haven't chosen the medical career, and that you cannot help your relatives.

But actually, so I imagine that many of you can relate to this story to some extent. What is interesting for me in AI is that the situation has changed dramatically, thanks to AI. Now as a computer scientist, I can sit in my house at my desk and work on algorithms for solving medical problems at scale, without medical degree, collaborating with doctors but using my mathematical and computer science skills. Or I can join big corporation and make money, that is our desire two paths, I chose the first one. I chose Stanford.

I'm working in the lab where we deal with problems with movement, in general, and with human mobility.

Let me show some projects that I work on. The first project is related to neuroscience. When a doctor has a patient like the girl here with problems with walking, the first thing that the doctor has to do is send the child to a so-called gate laboratory where they are taking precise measurements of movement of this child, and the problem with this setting is that the laboratory is equipped with super high tech hardware and the visit in the gait lab costs $3,000. Wouldn't it be great if we could take a mobile phone, record the child with a mobile phone and get the same clinical variables that we get from this expensive lab and enable measurements at scale in developing countries, enable monitoring progress of child after surgery and prediagnostics and many other things. Yes, it would be great, let me answer that.

That is the technology that we are developing at Stanford. We are trying to take the videos of children, and extract information from these videos using neural networks, here open posts, algorithm called open post and we use the signals extracted from videos to predict certain clinical variables. It works great and drops down the price from $3,000 to basically zero dollars.

Another example around human mobility is classical deep learning example, where we are working on radiology. We developed a algorithm for extracting information about arthritis from a x-ray scan as here.

That is quite basic application of deep learning, five years ago it would be unheard of actually, that that is possible but now we can outperform human doctors. Stanford's top clinicians by big margin, actually, I got the results this week and we are quite far on that, compared to doctors with 20 years of expertise.

The third example that I want to show is like out of box and the moon shot, we are trying to clone a person. The idea is that when you are trying to do a surgery on a person, you would ideally want to have ten clones of this person, apply different surgeries on ten clones, to see how the person will react to certain surgeries, and then you could apply the surgery to actual patient.

It sounds a little bit unethical, I can imagine. So let me show you what, how our clones look like. Fortunately the clones are just computer simulations. We try to fit musculoskeletal model to patients to understand how they control their body. With that, we will be able to use AI to drive the skeletons, and now if we have a good model on how people walk and how people move, we can simulate surgeries in this model, and predict how the person will walk after the surgery or with prosthetic leg which should appear here and didn't, but so that's fine.

These are three projects that I'm working on. An interesting thing about the projects is that especially the last one, is that I did it around one year and a half at Stanford. I wouldn't be able to do it by myself. It was basically a lot of work to make something

like this happen. The way we approach that was that I constructed the last problem as a challenge, and I publish it on the platform crowd AI where we attracted 500 people from all over the world to deal with these problems.

They solved the problem for us. This motivated me to move forward with this project. I started a group at Stanford called health AI, a working group where we attract, we have 250 members trying to solve healthcare problems at scale using their expertise from Stanford.

Now the mission of the group is in line with what you have heard here from other speakers, where we try to build scalable solutions for healthcare. We want to democratize healthcare and make free healthcare for everyone. But that is just the big mission that we all have. How would we actually do that?

We started building a platform where we are putting the models that I've showed you before and make it available for everyone in the world. The x-ray project I showed where we outperform humans in x-rays annotations is available for everyone right now. You can enter the website, put your x-ray of your knee and get a prediction if you have osteoarthritis or not. We are putting more of these models. The problem is that we are techies. We can provide solutions. But we are looking for problems. We are looking for people who actually can take our solutions and apply them in practice, or can bring us problems that we can solve with our technologies and then we can make it available for everyone in the world.

I believe that by everyone in the world, I mean developing countries in particular, and I don't mean developing countries as like only because I'm a good person, so I want to go to developing countries. It is that in developing countries, these technologies can bring the most value, because we have doctors in the U.S., we have doctors in Europe, and the value that artificial intelligence adds there is marginal, while in developing countries, we can make huge difference right away.

Then reiterate on that, build better technologies, build data, leapfrog the classic clinical settings from developed world. That is basically the point. Now let me just finish with one assurance. I know that I may seem like a green millennial, or as like a Silicon Valley lunatic and that's correct actually. I'm glad that I can bring this Silicon Valley enthusiast here, and I hope that actually you will allow me to solve your problems and allow me to bring expertise from Stanford on board for the on-ground work in Africa, in other developing countries, because we have a lot of excitement, we have a lot of solutions. We just need problems.

The overall idea is to bring together engineers, doctors and entrepreneurs to discover, develop and deploy scalable solutions for healthcare, by entrepreneurs I mean everybody who is willing to work hands on in developing countries to implement the solutions in practice.

Thank you very much for your attention.

(applause).

>> EFFY VAYENA: Thank you, Lukasz. Thank you for keeping it to the 8 minutes exactly. I'd like to encourage the other panelists to keep up with this. Anne is our next speaker. We will take questions after everyone is finished.

>> ANNE TORILL NORDSLETTA: Thank you. This is a difficult task for me to speak after one from Silicon Valley. Wish me luck (chuckles).

My name is Anne Torill Nordsletta, Director for the health analytics department at the Norwegian center for e-health research. We are located in the northern part of Norway, above the arctic circle.

The Norwegian center for e-health research was established and given a new assignment in 2016. So we are to contribute with knowledge to a common national ICT solution for health and care services in Norway. Our center has been collaborating with the World Health Organization since 2002. The health analytics department at our center is working hard on building knowledge within health analytics.

That includes machine learning, deep learning, and also the use of natural language processing. We have researchers with deep knowledge in privacy preserving data mining, they are pulling out aggregated data from the general practitioner offices without taking off any sensitive information. So it's kind of aggregated data. But in Norway we have noticed that there is a basic need to share information to decision-makers and others about what does artificial intelligence really do, and how does it work within healthcare, what it is, and

there is also a need to give clearly definitions about what machine learning is, especially within healthcare.

What is needed, what are the risks? What are the possibilities and obstacles with health analytics? So artificial intelligence is currently very hyped and there is a lot of uncertainty about it. So it can be used in settings where it really isn't artificial intelligent or it isn't machine learning but they are kind of using the words even though.

Therefore, we at the center have recently been writing a report about machine learning within healthcare. It will be provided to important stakeholders. We are also working on a natural language processing project which I will be presenting here. The project is called nor in text but my project is artificial intelligence based on real world data. I think I need the .... (pause).

Sorry about that. I have to push the right button. There. So real world data are collected from a variety of sources relating to patient's health, as from services as environmental and social media data can be electronic health records and registries.

In Norway, the authorities are investing heavily in establishing a database for all the national health registries for research. But real world data also comes from a number of sources and that can be genomics, medical imaging, claims databases, patient monitoring devices, but also from many, many more sources than that I have mentioned here.

It is estimated, dependent on which literature you read, that 40 percentage or more of clinical data is unstructured in the electronic health records. This means much of it is not easily reusable for decision support, or within research.

So, in a previous collaborating study performed by the Norwegian center for e-health research, by the university of Trumsa and the university hospital of north Norway, patient histories as documented in clinical notes was used to predict anastomosis Le cage which is a common complication after colorectal cancer surgery.

This study used a word model with a use of machine learning where a large amount of data was included.

The prediction algorithm can be valuable for early detection and assist in preoperative planning, give early warning and decision support. The study had a sensitivity of a hundred percentage which is good, and a specificity of 72 percentage, so low specificity may result in expensive false alarms. The health analytics team with our research collaborators are preparing to use natural language processing tools on the same field text data we have used in this previous study.

So the most frequent document types to be included are nurses' notes, journal notes, outpatient notes, radiology reports, referrals and admission letters. So in this project we will extend the study for improving the previously reported specificity of predicting past operative complications in colorectal surgery.

The specificity, called the true negative false rate, was 72 percentage on the previous study. With the use of natural language processing, we will see if there is a possibility to improve the specificity. So specificity is the percentage of healthy people or patients who are not in risk having complications, and who are correctly identified as not having the condition. The prediction algorithm can be used for early detection, and identify risk patients.

And the improved specificity percentage will give less expensive false alarms. But the challenge with using natural language processing tools is that they are developed for English language, and are not directly applicable for Norwegian language. So therefore, we have partnered with professor at the Stockholm university, and he will help the team to adapt clinical text mining tools from Swedish to Norwegian.

The professor is actually publishing a book about clinical text mining, and it's coming out these days. He's been working on establishing Swedish health bank for clinical texts.

Real world data from other sources contribute to the study, that is my question. Perhaps if we managed to, for instance, add other clinical data, data from other clinics, and at home data. But there is a big possibility that we will face ethical and legal challenge because of this.

And especially, since the European Union's general data protection regulation, GDPR will come in force in Norway this month. So and

one concern with use of the machine learning as many other people here have been talking about or mentioned is the black box, that the model will often produce a result without an explanation, and that is kind of difficult if you want the clinicians to use the models, and they are not able to understand what the predictions are or how they have formed the predictions, so that is a big challenge.

So for the Norton text project we are working on where we are using natural language processing is that we have access to the data, but we have to follow strict ethical guidelines for accessing and handling and analyzing the data. That is very important. Thank you.

(applause).

>> EFFY VAYENA: Thank you. Our next presentation is P. Anandan.

>> P. ANANDAN: Thank you very much. Thank the organizers for giving me this opportunity. We have a new kind of being I'll explain a little bit, but why am I here, so this is a story about healthcare, I think could have been repeated years ago and now. There is a lot of advancements, lot of progress in terms of public health metrics. You just have to figure out which metric. But there is always a lot of things to do, in other words, we are still not where we want to be and we probably are not where we wanted to be even now and this is a continuing kind of problem.

I come from India, where there is about 1.3 billion people. Maybe about 800 million or so, I'm just making up the number, don't have access to regular mainstream healthcare doctors. They are from rural

and poor communities who don't have as much of an opportunity, and they can't afford to pay for healthcare. This tends to be a continuing problem. But fortunately, there is a system of, that's been in place for some time, run by the governments across the states in India. They are a hierarchical system of medical facilities.

They start with medical colleges and specialties, district hospital, but really the action starts when you come to these primary healthcare centers, and as it's listed there are about 25,000 across India that serve the rural communities across all the different states of India. But more importantly, there are about 155,000 what are called sub centers. These serve between 1,000 and 10,000 people, depending on where they are.

This is often the first touch point for people in villages. They go to the sub center, if they cannot make their way to the primary healthcare center which has more facilities and so forth.

A lot of the help is being given by end first mile or last mile healthcare workers, social healthcare activists. There are a million of them across India. These are not qualified medical professionals, they are not doctors. They haven't gone to medical school. In fact, in many health centers, there may be, in some parts of India, 65 percent of the doctor medical professional positions are vacant, because there is not enough opportunity, people don't want to move their families there. But other parts of India, you see them functioning well.

This is a continuing story. However, what we found is that there is a lot of effort to make these healthcare centers primary health center and sub centers function better. We know a number of NGOs and other organisations that have taken this as a charter. One of the things that they would like to do, in order to make up for the lack of human resource, medical expertise, is to bring technology to bear. To go back to some of the previous speakers in the morning and now what Lukasz was saying, I mention that, I don't think Lukasz is a green millennial, although he might be a millionaire, he is wearing green shoes though I think I should mention (chuckles).

Anyway, green millennial and others who have solutions and they are trying to reach the audience. Why am I saying this? We started, we have started in Mumbai as a not-for-profit AI research institute specifically with the idea of bringing AI technology for social good. Health is not the only domain. We are working in agriculture, education, public infrastructure, financial inclusion and other areas. It is funded by the two donors, who have the benefit of having the Prime Minister inaugurate us and the reason he did it, this is the key, is because there is not a lot of organisations that are trying to apply AI for social good working closely with the government and social sectors. Private companies dominate the area, but their technology does not reach the kind of people that we are trying to reach.

But we can apply the technology blindly but because as you might

know, AI learns in place. You have to build the right type of data in the right context in order to train the models. You can't just train it in one place and apply it another place. That means we have to put in the effort needed to build the models, build the solutions in place, based on perhaps generic technologies such as deep learning and other things that could have been built but they have to be retrained. We are dealing with a domain where people don't have technical expertise.

Our mission, the way we are going to do things is to work with social technical organisations, governments, to source problems, identify use cases, find data sets, develop solutions and deploy them, but the scaling can only be done by these external organisations.

One reason I mention this is because there has been a fair amount of conversation today about the need for the ecosystem. In fact what we believe is that the one way we can do this is not just do it ourselves, but by taking advantage of where we are or what we are going to do to turn ourselves into helpful innovation where anybody who wants to work in the sector, but wants to see their solutions executed in the domain, in the field, in the villages, can actually work with us or work through us. So we pretty much do everything in open. It is a not-for-profit. There is no issue about proprietary commercial value, things of that sort.

We have already started doing this. We earnestly begin our technical work this summer, at the end of the month. I don't have

a project to tell you that we are working on. But the initial projects that have come up in health domains are what you might imagine. Can we empower these front line healthcare workers, a nurse, a midwife, with technology to help solve the problems right away without having to wait to go to the doctor, without having to look for a doctor.

I've gone to some of the health centers, where there is a doctor who might show up and he will stay for four hours. And there are 150 patients who come by. There is not a lot a doctor can do in that time with that many patients. But can we reduce the load by giving them technology. Some of the talks that I heard today give me hope that we can actually, but a lot of the things -- I've been a technologist, I've worked in the industry. When technology reaches the audience, often it is not just technology but adapt ability to the audience that are involved.

Simple things like user interface assumptions or understanding about what people might know, what data may be available about all of that, tends to be not valid. A solution you develop doesn't translate. So there is quite a bit of redesign and adaptation that needs to be in place, which is something we will do. Of course, we invite other people to do.

The second one is at a mass scale. Diseases, pandemic prediction, every state government, state health official is desperately interested in doing that, can we help them do that. There are all kind of things. One thing about in India, even in the last decade,

the time that I've been in and out, the accessibility to technology through cell phones has incredibly increased. Fiber is available almost everywhere and with smart phones or not, something that we have to look for.

Bottom line is we can do this. But one of the big challenges is data. What we found is that when we started looking at these things, people are interested, there are many interesting problems people think they have data but that is not the right kind of data. They have not been collected with the purpose of doing analytics, doing machine learning. We have to scrounge around for the data, sharing data, issues with respect to privacy exists. But the other serious problem, ecosystem problem, we are dealing with a complex ecosystem, diverse organisations, the idea of incentives. How do we connect them and bring them together.

It's a challenge but one we believe can be met if we dedicate ourselves to do it. More importantly, I want to re-emphasize this, the fact, if you want to do this right now in this space and countries like India, you cannot do it with a commercial model in mind, because there is not much money to be made and the market is extremely fragmented. At least for the time being we have to de-risk it philanthropic initiatives are going to take the solutions to the end. As far as scaling is concerned, I'm sure in this case only the government can scale the efforts. So thank you.

(applause).

Cycle of how we plan to do our work, but in the interest of time I'll leave that out.

>> EFFY VAYENA: Thank you. We have very punctual speakers. Thank you very much. We are moving on to Matteo next.

>> MATTEO BERLUCCHI: Thank you. Lukasz, there is a third way, you don't have to choose between being a poor academic or rich computer scientist at Google, you can be a middle income entrepreneur like me. Which is okay (chuckles).

Matteo Berlucchi, I'm a dreamer, I'm Italian, and I have a passion for how technology can solve problems. As was said at the opening of this event, you should think about the problem first and then how to fix it, rather than do it the other way around.

I'm fortunate enough to have stumbled across this fantastic opportunity that I've been working on for the last three years, YourMD, here it is, start.

I'm a serial entrepreneur, kind of reasonably successful, I've done nine start-ups, and I'm recently joined the entrepreneur residence so I can share my battle stories with young entrepreneurs that are trying to build businesses.

YourMD has a simple vision and this is what I want to share with you. I hope you walk away from this presentation with one thing, is this, the point is if you look at the global situation with healthcare, and the reason why people go to the doctors, is fundamentally to get information. So the chief knowledge officer at the NHS said that

information is the single biggest problem in healthcare. We simply don't know what to do. The only people don't know what to do are physicians. So we have to go to them to find out what to do. What we are trying to do with your M.D. is address this problem.

The global shortage of doctors is actually worse than you think if you start looking at generalists. It's interesting that doctors are humans like us, they are people, and they like to make money. When you spend a lot of time qualifying as a doctor, what tends to happen is that these doctors tend to specialize, because they make more money, and to live in beautiful places because why not live in Los Angeles instead of the suburbs of an Indian city.

So the actual number, when you look at the number of doctors per 1,000 people, if you look at the number of generalists it's much lower. The average is 30 percent, 30 percent of doctors are generalists. But when you need to deliver primary care, which is a topic of this talk is promoting healthcare, the generalists have to do it.

You can see the numbers there, they are very very bad, even in the U.S., only 0.3 generalists per 1,000 people.

The role of primary care is fundamental because there are several studies that show that if you implement, invest in primary care, you have excellent results down the pipeline. So there is a study that was mentioned, in the economist, a couple weeks ago that I'm sure you read on universal healthcare mentioned in Costa Rica where

they invested in putting health workers and primary care workers in the field and they had very very good results.

What is the fundamental information problem? I'm going to walk you through the slide quickly. The role of a doctor is basically to keep you away from harm, be sure you go to home if you need to go home but the problem is if you see the person on the left when everyone in the world regardless of where they are, they have a question about their health, they have basically a long wait to see the generalist. In the UK it's a average of two weeks to see a doctor. In the U.S., in major cities up to 24 days.

Most people go to Google. We have all done it. Has anybody gone to Google to check your symptoms? Yes? (chuckles).

Now, Google is great. But they have a problem. The algorithm is not designed to address health questions, because they don't really know what is your situation. If we all type headache in Google now, we will get the same results. We could have 122 different reasons why we have a headache.

Google doesn't work, and to make the problem even worse, they give you a lot of links, and you don't know which links you need to click on, because you are not a doctor. So you are back to square one. And then the famous joke that people with headache end up diagnosing themselves with brain cancer because the fourth of fifth link is brain cancer and it has the same symptoms, and then you end up jumping to the wrong conclusion. We all end up having to go to

see the doctors. That is why you have a bottleneck.

The bottleneck is not that there are not enough doctors. The bottleneck is that we go too often to the doctors. We should go less to see the doctors, particularly for all those things that don't require a specialist. The things that you could self-care for. There is an interesting statistic from the NHS that says up to 80 percent of visits to the GP are for minor ailments. What is actually happening in reality is people are going to see the doctor, to get information. You go to see the doctor, doctor gives you the answer, you walk out and you are done.

But that is not very efficient. So what if you could take on the burden of the information side of primary care, which we call pre-primary care, which is that moment when you seek personal, safe and actionable health information through an automated system that doesn't have the issue of working hours, being tired, not having a doctor near you, and be able to tell people what they need to know, when they need to know it, in an instant, including going to a hospital.

If you are vomiting blood, it is not a matter of diagnosing what the problem is, it is a problem of getting to the hospital as quickly as possible. We have built this, we started three years ago, the service is available on every platform, where the only medically approved bot on Skype, we are on Internet of org on Facebook, we are on the web, IOS and Android, because we want accessibility. We want everybody to be able to use it. We don't want inequality.

Your M.D. is trying to give you the best possible information every time, so that if you can self care, you can self care, that reduces load on doctors, but if you need a test, you need to see a doctor, it tells you why you need to see the doctor. The key is also, one of the key things I think about health promotion is giving the actual information, because if you just use artificial intelligence to give people the likelihood of having a condition, you haven't finished the job, because if you tell me that I have migraine, I don't know what migraine is. I'm back to Google. Then I'm lost again. Then I need to go to see the doctor again.

So we partner with NHS in the UK because they have arguably one of the best data sets in terms of information, basic information, safe, clinically assured, written for people 11-year-olds plus, so that you can read about what you need to know without having to look up the information on Google.

Then once you know what to do, we have, and this is how we monetize the platform, we created a thing called one stop health, which is a sort of Amazon of healthcare, where we curate a list of services, products, apps, hospitals, clinics, on-line platforms, that can help you with the specific problem, like in the slide if you have hay fever, we have a partner that can sell you hay fever medication on line. This is how we approach the problem, and it's entirely about information.

It is just giving the right information to everyone. Now trust

is one of the main topics here and I'm glad it was so prominent, because I realized a year ago that trust was a super big problem. Let me check my time, 8 minutes, I'm going to do one more minute. What I've done in the last year to build trust, because I realize that even if I give you the best information in the world, you don't know if you can trust me. You trust your doctor because there is a plate behind the wall that says the doctor went to university. You say great, I can trust him. I don't have that. Nobody has it in AI.

I'm building trust in the only ways I can. First thing we did, we introduce clinical safety management system in the company, like a hospital, a health company, very few start-ups have luxury too do that or want to invest in this because it cost money. We were one of the first to get the CE mark in April. Then appoint a clinical advisory board because I'm not a doctor. I'm a physicist. When I don't know what to do I can ask my advisory board. We are lucky to have someone on the board who keeps us honest and straight. One point about transparency and being able to trace decision-making, we build a model to say you have migraine, then you ask me but who tells you that I've got migraine, how can you make sure the data you put in your system is correct. I'm announcing today that we have an exclusive partnership with the British medical journal who is going to be our validators of data models. This is one of the first validation announcements I've ever heard, where you have an independent third

party with a lot of trust working on the analysis of the data to prove its quality.

The fourth obviously for users is very important, transparency, so we are GDPR compliant but we are trying to go further than that, so we have a dashboard where users can see the data we hold about them and they can delete every data point in the system.

What do we need while we are here? We need some money, some help because we are a small start-up, I decided not to go to venture capitalists, because I don't like their model. I want to build something scalable that lasts, doesn't get sold to somebody just for cash. We would like to find the government or an institution that is interested in doing some pilots of this at an information level, and help us also ideally translating some of this corpus we have, the luxury to have a license with the NHS we can translate in 14 languages, because giving the right information to people is the single biggest thing you can do to improve primary care anywhere in the world.

Thank you.

(applause).

>> EFFY VAYENA: Thank you. Moving on to Rafael.

>> I'll do my best to stick to the time but I'm half Spanish, half Swiss, so I usually struggle. My Spanish side is pushing more. I'll do my best. Apart from that, I work at the university of Geneva at the faculty of medicine and most specifically I work at the institute

of global health, where we have a team looking at problems and health challenges, public health challenges, at the human animal ecosystem interface. I'm pleased to see this morning some of you already mentioning about one health approach, so that is great, seeing the link between one health approach and AI.

Over the next minutes, I'll try to go quite fast, I'd like to talk to you about this major and current humanitarian crisis. I'd like to talk to you about snake bite. This is probably something new for most of you. The approach we bring to the problem, we try to build this snap project, snap idea, building the first medical decision support tool for snake identifications, based on artificial intelligence and collaborative expertise.

Snakes, definitely scary animals, very interesting animals at the same time, but very scary, and I would say more, the numbers and the figures in terms of snake bite global burden of snake bite are scary and alarming. Over 125,000 people die every year as a result of snakes bite. This particularly affects poor countries, in Sub-Saharan Africa, Asia, Central Asia, southeast Asia, Latin America and other parts of the world, it affects workers, women, children, it is a humanitarian crisis.

At the same time, not only snake bite kill people and snake bite and venoming, they also cause dramatic consequences for these people in terms of amputation, disability, psychological trauma, stigma, etcetera, incapacity to work for example.

The good news are that last year, after many years of fight of some key players such as doctor without borders and many others, W.H.O. finally accepted to include snake bite and venoming as part of the neglected tropical diseases list together with rabies and other major diseases. Today we are going through a crucial and very interesting and promising political and scientific momentum around opportunities for tackling and combating this humanitarian crisis.

Most of the work, W.H.O. is working on a roadmap for the years to come, and most of the work is of course focused on antivenin development, making sure these antivenin are safe and well distributed all over the world but same there is a huge window of opportunities for additional innovation coming from the DPO world from artificial intelligence.

I'd like to talk about what is the role or at least what we feel could be a possible opportunity for artificial intelligence in this domain.

Here a quick example, this is children, a kid, boy that was bitten in the face by a snake, South Sudan. Here the snake, so think carefully what is the critical, at least one of the most important question for a clinician when confronted to this situation and to this situation. So one of the critical questions is this one here, or this question, so first of all is the snakes venomous or not? Secondly, what type of snake and what type of venom, has the snake injected. What type of clinical manifestation can I anticipate in this situation?

Some of these questions are extremely complicated. Especially because, first of all, medical doctors are not trained, they are not snake experts at all, this is a field that is part of zoology, animal biology so it's a completely different field. At the same time, we are confronted with a huge variety of snakes, huge diversity of snakes at the global scale. These are numbers from the World Health Organization. Over 3,000 species of snakes including venomous and non-venomous but we need to know if it's venomous or not. We have 600 species that are venomous and more than 200, 219 species that are highly venomous or very important from a medical perspective.

The problem, the third is that many of these species and snake bites require very specific treatments and very specific way of taking care, managing the patient. Knowing which snake we are talking about is quite critical.

Here our idea, so basically the first medical decision support tool that first of all will help support clinicians in this first step in trying to identify, trying to give a name to that snake. We have already identified, and I won't go through in detail so don't worry, we have already identified a number of scenarios, validated with colleagues from NSF and other key experts in the field of snake bite, where we believe that such a tool could actually work.

To give you an example, in many parts of the world, Nigeria, for example, over 20 percent, Myanmar over 20 percent the victim or bystander will kill the snake after the bite, they will take the

snake, bring it to the clinician, put it on the table, because they know it is important. That is what is happening in many parts of the world. The clinician could take a picture of the snake, but more important than that, you could encourage the communities not to kill the snake because that is actually quite dangerous and many secondary bites occur while killing the snake.

That is where opportunities where we could take picture of the snakes. At the same time, in India and Nepal, we have this health workers that go immediately through with motorbikes and other service they go to the place where the bite occurred rapidly, so we can imagine they also take pictures of the snake for identification before going to the clinical doctor. In India as well, I'm happy we had another presenter in India, they have a whole network of snake catchers. These are people that go to your house, take the snake and remove the snake so it's, they are already doing this, taking picture. So a number of opportunities that we are currently exploring.

In the process of identifying that snake or identifying that image and trying to give a zoo logical name, we have imagined a three level validation process. First of all artificial intelligence, secondly we imagine a volunteer thinking approach so more crowdsourcing citizen science approach. And of course gold standard, which will be snake experts, professional herpologists.

Basically the next steps are main objectives that we foresee in our project for the next months and years. We need to build a massive

open data or image repository of snake photos. Secondly we will be working based on those images on development of a artificial intelligence algorithm. 3, we want to demonstrate that the crowd is quite good and efficient in their capacity to identify snakes. Finally we will introduce this into a app and ultimately we would like to go beyond a purely zoological tool because so far I've been presenting what is the name of the snake, we would like to go beyond, move to a monitoring tool, etcetera.

I'll briefly go through objectives 1 and 2 mainly. For objective 1, we need to build this photo repository. This is critical. We already started looking at this. We recently published this paper where we identify a number of open data sources where we can find material where we can find images on snakes. To give you a specific example, this is a platform called I naturalist and the same way you have people that like birds and go out during the weekends and like bird-watching and take picture of birds, there is a huge community out there of people that like snakes. They go out there during the weekends and take pictures of snakes.

This is happening spontaneously on the web. I don't know how many of you, how many snakes you have seen in your life maybe, but if you live here, sorry the numbers of small but these people here, these are some of them, 214 observation of snakes, 209, 206. They go actively out to search for snakes. This is happening out there.

We cannot neglect this. Facebook, there are groups where people

share images on snakes and they like to get the challenge of trying to identify the snake, because it's something that is not quite easy. We are frankly collaborating already with this group of snake experts in India, so it's a Indian snake organisation and they are leading this network of snake rescuers, snake catcher that I mentioned before. You call them and they come home and pick the snake, remove the snake from your house. They will be on board in terms of getting pictures and trying to identify snakes.

I'm happy to say, NSF, doctors without borders is formally endorsing the project, they are on board with the global network getting images from the front line. They are exposed to snake bite frequently working in areas of South Sudan, as I showed the image of the little boy. They will be feeding the system and helping us in different stages of the process, and particularly also in the clinical validation and implementation of the tool.

I'm happy also, this is part of the collaboration with the team at EPFS. We started preliminarily working on algorithms. There is potential to work on this, but there is a lot of work to be done. I would like also, because this is ultimately a humanitarian challenge, we would like to make this as open as possible. We believe that it's of course a very good team, but we need to open to other teams, Silicon Valley to also get involved in trying to get that very good algorithm out there.

I'll stop here. I thank you for your attention. The project will

start officially in September, this year, I have the key player, the key partners. But I'm happy to discuss and to involve others. We very much need help. Thank you.

(applause).

>> EFFY VAYENA: Thank you, Rafael. Today, already in the morning sessions and this afternoon, there was reference often to the ethical challenges that many of you who are trying to develop these applications are facing, also reference to privacy. Our next presentation comes exactly right here to present and discuss some of these challenges. Dafna, you have to introduce yourself.

>> DAFNA FEINHOLZ: Thank you very much. First of all, I think I'm Mexican, so -- (chuckles) that is part of the programme part? No, I'm sorry for the delay. I was in another meeting that took longer. And then it was impossible to reach this room without the help of one of my colleagues, I wouldn't be here. I'm chief of the science and technology section at UNESCO. I'm going to share with you ethical ideas or some of the ethical dimensions that we thought is important to take into account when we develop these things. Please tell me when I'm on time. One of the most important things by listening to previous speakers is how are we trying to deal with artificial intelligence with everything that relates to life and health in particular, that is so complex and so difficult. I want to give two examples that will try to introduce why I think ethical reflections and cultural dimensions are so important when we develop projects

with AI artificial intelligence. One is about particularly related when people go to doctors.

We conducted in Mexico, before I joined UNESCO in Mexico I was part of the National Institute of Health. We conducted a survey a we wanted to know when people go to doctors and why. We conducted the survey among doctors and the population and it was completely different, the results were passionate because we saw that, for example, for clinicians, health means like a line from being very sick to being healthy. And for the population, it has nothing to do. So people would go to the doctor when they have some, I don't know probably when they are coughing with blood, but before that, a headache must be probably because you are having a bad day, and you are not in good mood. But that is not a health problem. So you don't go to a doctor for that one.

The other one is the snakes, so in the north of Mexico there are lots of snakes, but there is a cultural perspective very important, because for them, for the people there, the soul of the person is, if a person is bitten by the snake, the soul of the person is kept by the snake. So they don't want to kill the snake for that reason.

But in the clinics, the doctors do not understand that. So there was a case that a mother came to the clinic with the baby and the snake, and of course people in the hospital were crazy, right? Like please kill that snake! And the mother said there is no way. And of course, the snake ran away. Can you imagine the crisis? But you

will stop laughing because it was so terrible, at the end of it, is because when they found it, they kill it. But the mother committed suicide, because she felt that she was not responsible enough to keep the soul of her child, because it was in there.

So until the child was, that was cured, then that snake could be released, so this is just to give an idea, this I wanted to take advantage of this, that I didn't have in mind, none of these examples, when I prepared my talk. But I just took the advantage of this to bring to why I think the ethical and the cultural dimensions are so important when we discuss health in general, and when we want to add now artificial intelligence to it.

So, the basic considerations that I will try to go through are what is a good society? I mean who will decide that, even health is not so clearly decided. What should health promotion contribute to a good society. So what would be this good health, how might artificial intelligence base help promote or enhance autonomy of citizens, moving from the perspective of persuasion and coercion, I will try to give some examples.

I am always a bit care, I mean worried that with this new models, we are trying, we are speaking about patients more as consumers instead of patients, consumers of apps, consumers of many other things, and I'm a bit worried about this move. How are the benefits distributed of all the things that we are doing, and the different projects that we are listening to.

One of the most important things I guess in the ethical considerations is the integrity of data. We are building all the programmes and all the applications based on data. But the data can be incomplete for many reasons, maybe because some of those that we can computerize are not computerizable. So we are only using information that is prone to be used to be translated into numbers.

It of course can be biased in many ways, by culture, by age, by all the criteria, I don't need to go into that, and it could be misleading of course as well.

So how can we ensure that we are using the best possible version of the data to create models and algorithms? In addition, physicians must be able to critically access the source of the data used, because of this inherent biases and the algorithm itself. We saw that there is already kind of an example of what you said like trying to make sure that the information that you are using, but there are also errors when some algorithms are being processed.

So I think it's also that is why it's important that there be the possibility of critically assess the source.

Now when we do the algorithm decision-making, the benefit of analyzing, of replacing the physician intuition with algorithm needs to be performed. It is not only about the information. It is about how the doctor looks at the person, sees the person, feels the person, interacts with the person. That is what the clinic is about. Information is basic but it's not information.

The value of the mutual consensus between physician and patient is very important, and the role of the family, how are we going to put that into any of the algorithms that we are creating? I mean family plays a major role in decision-making in health decisions, in particularly, in many countries like in continents like Africa, like Latin America, in many countries in Asia, so how are we going to do that?

In any case, the decision that might be suggested by the algorithm on the application should not be the only one. It should be taken in consultation with the patient and the physician should have the last word. It should not be left to the application to make the best decision because of all these reasons that we are saying.

What about empathy? Can we really put empathy into the algorithms? How many things you would say or you don't say to your doctor according to how do you feel when you are, I mean feel about yourself when you are sick and you go to see the doctor. Okay. So empathy is also very very important.

And for mobile health applications, for example, now we have a third party which is the one who is creating these devices. I think we have to take into account always that there is a power relationship between patient and doctors, and now we have another and a third party who is the one who is creating this device.

How do we know that these recommendations that are going to be provided are really in the best interests of the patient? The sharing

the benefits, efforts need to be directed towards both developing and developing world, the population most in need of interventions will have access really to all the things we are discussing? It is important to, not to create more inequality, particularly in the poorest regions in the world. To what extent patients will have access to know how their information is being gathered and used and have an understanding on that, what will happen if they are not going to be having a saying of that?

Are we going to really, there was an example of a person that probably the device is suggesting that this person might have a suicidal conduct, are we also going to go and spy and do something, going to his house, what are we going to do?

In privacy, of course, we have all these problems of data ownership, confidentiality, and about language, what I also wanted to say is language is one thing, but, language is very important but language is not only words. I'll give you another example, and I think I will be -- so this is in fact one of the values of scale up, it's the cultural perspective is the issues of technological access and cultural congruence, we did the algorithms. And I think with my previous examples, I tried to convey this message. Thank you very much.

(applause).

>> EFFY VAYENA: Thank you. Okay. So we have about 20 minutes, 25 minutes for discussion and questions. I have so many questions.

I suppose many of you, so let's mix it up, maybe I'll give you first the floor, and then we can have -- just identify yourself and then your question.

>> I have a quick question for Lukasz and maybe a longer question for Dafna. For Lukasz, as part of Stanford's AI sort of a lab, are you looking for problems where you can have more novel, interesting new approaches to artificial or machine learning algorithms, are you looking to apply existing algorithms to problem. The question for Dafna is you talked about how the physician needs to be sort of the final word, but one of the things that is consistently being brought up throughout the day is that there is a lack of doctors. How do you manage the need to, sort of using a person as the end stop while providing sufficient healthcare to everyone, which is I think one of the goals of this.

>> LUKASZ KIDZINSKI: Can you hear me? Okay, good. That is a very good question. Thank you very much. So, when I was talking, the health AI group we realize there are many solutions available already, like the app or algorithm for detecting skin cancer was developed in one of the labs at Stanford. When you talk with people who develop the algorithms, they often say there is no viable business model for deploying this algorithms, or there is some legal barriers for that.

Actually, we are also interested in just the deployment part, so many people already have algorithms, but then there are people

from law school and from GSB, Graduate School of Business at Stanford who think about how these methods can be deployed in practice. On the other side, we have researchers who are excited about research and pushing research boundaries further.

So there are people from both worlds, trying to talk to each other. That is the overall idea.

>> DAFNA FEINHOLZ: Is this working? Is it? Okay. Thank you for the question.

Yeah, first of all, I think the whole idea of my intervention was not to say that we should not use them at all, and I believe that there are very good uses of artificial intelligence and technology in health. They are already proven. And of course in some cases it's definitely, I mean the idea is that these devices could help and be allies of the doctors. And in some specific cases, if they can replace them because there is really no way, and the phones in Africa, the cell phones in Africa have been useful for that to communicate, to locate people that are not there, that's fine.

But what I mean is for very complicated decisions, I think it's very important not to see -- and as we develop more this technology, just to make sure, same as we are discussing for example for self-driving cars, decisions that have life, when life is at stake, humans have to be at some point responsible for it. That was the only ....

>> Thank you. Thank you very much for the very interesting panel.

I have a question actually geared towards for Dr. Anandan, countries like India has large health information systems, and they are just to preface my question, because I think it has to be understood in that context, India being a decentralized insofar as healthcare delivery is concerned which is federated in most countries that are the case, the data gets accumulated from the states upstream and comes to the national level.

But they are also monitoring key health conditions and snake bite is one of them. Snake bite is in their list and they are monitoring. And the states report upstream to the central. I think what we are trying to solve in that India context right now is to put AI back into the platform that they have already launched the large scale realtime data collection system. But the question here is that where do you want the country to think about having an AI component in the information system. At the level of pattern recognition, at the level of New Delhi for example collecting data from all 36 states, or at the level where response is needed which is on the ground, but in both cases, algorithms are different in the way in which what they want to achieve.

So what would be your technical advice to those countries that are interested in, okay, we like AI, and we are interested in it, and we want to utilize AI at the national level, how would you go about advising them? So I want to give the floor.

>> I won't speak about snake bite, I'll leave it for him. But

I think to answer your broader question, the answer is I think where it's most appropriate. If you are thinking about early diagnosis or intervention for individuals or even small groups, the technology has to be available at that level so it can readily be deployed, because even though the data might be centralized by the time the data goes back from the central to the end node it's going to be too late and so on.

There are multiple kinds of uses, use scenarios. For example, if I'm looking at apps to enable front line healthcare workers to give better advice, either it's for health conditions or for things like nutrition, child nutrition, child welfare, it should be there at that level. I think data that is collected there duly validated and stuff like that, needs to go back to algorithms at that level.

On the other hand, if you are thinking about predicting broad trends in health or disease prevention or a upcoming epidemics especially with seasonal changes, things like that, they have to be more at the district or state level, because they do vary a lot. India has a diverse kind of climates, depending, so and a centralized approach may not be useful.

But there are broader indicators and factors more to do with trends and policy level changes that have to be made in order to maybe make some changes, that might be longer more at the central level. The thing about AI is it can operate at any level. It is not limited. We don't need to think about it as one shoe fits all but rather adapt

it to the specific kind of usage scenario and use case, and bring the data to that level. That is how we are thinking about it.

>> Thank you very much. I think it's a very interesting point. I very much agree with what has been said. I think that two components are not necessarily independent or they should probably be complementary.

I think the case of India is interesting because a lot of snake bite occur in rural areas, but there is a lot of urban snake bite as well or peri urban snake bite. This is a interesting observation. Definitely, and that is partly why we very much want to collaborate with NSF for example. They are very much at the front line where snake bite is occurring. That is where we probably need this tool to help in identification. That is the first question before going into the clinical management. The tool of identifying the snake should be there.

Then if you go to the policy central level, as I said, we really want to go beyond the purely zooo logical question, first objective is what is the snake, when you know what is the snake comes all the clinical management aspect and epidemiological questions. In the long term, we can imagine the tool could provide information on where snake bite is occurring and giving a understanding of epidemiology of snake bite. If you go to the central level this is where you decide where to disseminate anti-venom. You can decide from the central level where you are sending which anti-venom where at the national

level, in a big country such as India. It is very complementary anyway I think. I don't know if that replies to your question.

>> Good afternoon.

>> Thank you. I'm from W.H.O., my question is more for the panelists, everyone, but Lukasz, I'll take you up on the offer of give you problems and you have the solution. But my question is more for the panelists. I've been on the technology side, have frustration of doing things but not being able to deliver it on the ground. The question is how do you see the U.N. and the government agencies, what do you see them doing specifically to make sure that this is fostering on the ground from the foundation side, from governments, how do you see the main base and what can be done more to help AI go to the next scale level.

>> My name is Anne. I think it's important to increase the awareness and knowledge about this, for people who are working in the authorities. That is really important. They should understand that. This is a tool that can be used, but we also may need physicians who will use these tools and understand that it will not give the right question, right answer every time, because it's predictions. And you have to understand that statistics, it's predictions, and it's not always giving the correct answer, even if you can kind of sometimes predict the things better, but it's a tool you are going to use.

>> Clarify, to your point, you are saying that it is important to inform and raise awareness within the government and within the

state institutions, or to the outside? Or both?

>> I think it's important for both. But especially for the authorities to understand what it is. And like the physician and clinicians should also get knowledge about what machine learning is, what is deep learning. The best thing that I have experienced is that when the physician also have a PhD in artificial intelligence, kind of have an understanding of both. So that is very valuable.

>> Can I -- I think this is a fundamental question, because particularly in healthcare, there is obviously a very high barrier to pass before you can get these things out in the market. I think in the institutions where there is a national level or international level, they should take proactive role in first of all secting the criterias for these technologies to be viable, because today there is no framework. So everybody is out there, everybody launches AI, you don't know whether it works, it doesn't work, how do you measure it, how do you decide if it works. There is no criteria. Today when I go to the doctor, there is a criteria, there is a certificate on the wall and that is why it works.

It needs to be done fast because if it takes ten years to create the framework, we are going to waste a lot of time. That to me is the priority.

>> I have actually a fairly specific idea about this. One of the things is that the technology capability or the solutions lie often within academic experts and industry experts. The mainstream

industry has literally cornered a lot of the AI and machine learning talent.

However, corporate for social responsibility reasons or even individual reasons, many of them are eager and ready to help. But they don't know how to take it to implementation. They don't have the incentives or the opportunity to do that. There is not a marketplace. There is not money.

Governments on the other hand who own these, own these communities and people don't have, don't know how to think about technology, don't know what is right, what will work and so on, and then there are also social sector organisations that are very good at working with governments and making things happen, that tend to be somewhat technologically not very advanced in their knowledge. Then there is also social enterprises, many of them we heard speak today, who have a particular innovation or technology, once again getting it to the end node.

Thinks about data ecosystem, this is going to require some kind of common shared data ecosystem for that particular. I think the best way to do this is for governments to actually initiate centers of excellence for AI for social good, or something like that, in different parts and make it something like a public/private partnership. Because they all have to come to play and they all have different incentives and it's only the governments that can initiate or catalyze such a thing. But I think it is not a good idea for

governments to run it.

>> I want to add to all this awareness raising and this stakeholder involvement, it's the population itself, because the users, if some of people of the government are not aware, imagine the population. And I think they are the end users. And I believe it is important even to elaborate the product and then to make sure, we need to understand how they understand this and how will the user use it or not. I would just include them definitely.

>> I think if I can follow up exactly on that, I think also the approach that I tried also to bring where the citizen has a role in our case collecting images of snakes, this is happening. This is something just happening out there, we cannot neglect the thousands of people that do this every day, as I said and it's happening globally.

So this is something that should be integrated somehow in terms of data collection, data analysis.

>> I'm very pragmatic and basically researchers are pragmatic. If you give us a list of problems, actually people will solve them. You may not think that they are solvable and people will sit and solve them.

One example of that is in 19th century, fames mathematician gave a list of 20 problems or ten problems that were completely unsolvable and within 50 years people solved them. If we can create a list of 20 problems that we believe that AI can approach, then people will just do that, I think. I think it will be five months.

>> I'm so tempted to give you a regulatory question to solve (chuckles).

But before actually we move to the next question, I was a little bit stuck to your point that today we don't have standards, and I assume you meant approval standards for technology that goes in AI technology, which makes me ask the question, you mentioned your project is a medical bot, and you license it or put it in the market as such. So what was your regulatory approval process? Did you go through a medical device process?

>> There was none basically, I'm making it myself. We are self-regulating ourselves because there isn't one. In the UK, every medical institution is controlled by a thing called the CQC care quality commission that doesn't regulate medical apps. Why? It regulates everything else. They need to do that fast, because people otherwise, because the public is fundamental in this but the public doesn't know. They don't know how to judge. Imagine I have the best app in the world that knows everything about your health and every time gives you the right answer. How do you know it's the right answer? You have no idea because you are not a doctor. You need the framework.

>> It's an important point here, right, somebody needs to develop the framework, I know in the UK they were trying to develop a framework for medical apps, and or the gray zone apps which is not exactly, I asked somebody that question earlier, because I was trying to understand how do you regulate, how do you go through market approval,

what are the verification standards, in healthcare we have used for years and years, are we going to use something else now. That is a take home message. Somebody needs to ... (chuckles) somebody there needs to help develop those standards, I'm not naming names.

We had one more question here, and then one.

>> Thank you. Thank you so much for a very excellent presentations, all of them very interesting and we understand that of course AI is a tool that should be used for good purpose, like health. My understanding as everybody else, it mainly depends on algorithm to reach to the final decision or the treatment or the diagnosis or whatever you need to go. But the doctors in the decision, they depend on the education, the training, on their years of experience, on everything they have seen, all the cases they have seen.

So my question actually is, what happens, and this is mainly in the bioethics has ever been discussed as there are differences in the conclusions between doctor and AI, if both didn't come to the same conclusion, whether in the diagnosis or whether in the protocol of treatment or the type of intervention to proceed with, so my question is if it has ever been discussed in the bioethics, and what we should do when we reach this point. Thank you.

(someone speaking off microphone).

>> It is actually quite related to that.

In terms of, okay, you say I walk into a doctor's office, I see a EMT certificate and that means I have trust. We can measure the

accuracy of that, right? Because we know roughly the error rates of human doctors in different fields.

Now we can very accurately measure the error rates of AI, we are increasingly finding out that those are better. So at which point, at which point then are we going to have regulatory framework in place, where we say as a human, you are not allowed to do that without the use of the technology, because it has been shown that you are worse at this.

>> If you don't mind because I spend the last two years thinking about this answer, so hopefully I can shed some of my -- whether you agree or not, it's the data from emerging markets in India, because everybody talks about this mythological empathic fantastic doctor that knows everything, that looks in your eyes and says, darling, don't worry, go home. The reality is that in India, if you are unlucky enough to be able to see a doctor to see for two minutes, they don't ask you, they ask you how are you and they give you something and they give you the right advice 30 percent of the time, 3-0, 42 percent of the time, they give you unnecessary or harmful advice. Okay? This is India, from the statistical economist last week.

Obviously, it's a very difficult point. I think that to your question I think it's going to be detected by the liability, because people do things based on fear generally. So the doctor will look at what is the most risky scenario, if they can be sued, they will do the thing that doesn't get them sued. Generally that is what happens

everywhere in the world. You know? So that I think answers that. In terms of the decisions support, it's going back to discretion, how do you get the doctors to adopt this, this is the key, because actually, I believe that for AI and health to become mainstream it needs to pass through the doctors, for the reason, they have the plaque on the wall, they have the trust, and the AI in health used by consumers, sorry, patients, actually, people, because they are all patients, they are people, so one-to-one matching (chuckles) yeah, so people, we need somebody we trust to tell us we can trust AI. It is the only way. We cannot educate people about AI. It is going to take centuries. We need to get somebody with trust to tell us Matteo, take this up, if it tells you to come see me, come see me. Why? Because doctors can be more effective, they can make more notations, make less mistakes and make the same money. It needs to be good for them. So the doctor needs to have an advantage, a benefit.

>> I think the solution is coming from over here. Lukasz.

>> I slightly disagree with that. I would rather believe more in algorithm, especially building this algorithm.

>> Computer scientist at Stanford of course you do.

>> (overlapping speakers).

>> Key point is that trust cannot be regulated under no condition. You can do whatever you want, whether you trust or not, your trust is up to you. There is no way you can regulated. What can be regulated is liability and accountability and stuff like that. If you think

about it, you will get a different answer.

>> Please go ahead.

>> I think the physician role is, it depends on which culture you are in. Like in Norway, when I used to go to the physician there, okay, I come with a problem, they say you should do this, this is what you should do. When I went to U.S. to Mayo Clinic and you get a appointment with physician, he say these are your options, 1, 2, 3, and you have to think about what do you, how much money do you have. The role of the physician is very different in which country you are.

>> But I'm having a little bit of a, if I may, use my position here to challenge you that trust is not regulated, it cannot be regulated, would you distinguish then trust because of course maybe something is trustworthy but simply don't trust it and nobody can force me by law to trust it. But would it make sense to say there are certain things that can be regulated, that will enable trust? For example, one of the things, if we trust the healthcare system often we do that because we are, or a drug or procedure, is because we validated through various methods, we came to a point to think that there is enough there, that is reliable thing or it's competent or whatever the criteria.

So I think, I guess, yes, trust in and of itself you cannot make it legal but the elements --

>> My point, if you can objectify the criteria of performance

in some way, percentage of errors or whatever, you could measure them. You can only standardize things that you can measure. But I think the more important question here that she was getting to which is really to me was about accountability, responsibility and liability, those I think are important questions to discuss.

>> But liability is a sort of side effect, let's say, right, if something works or doesn't work it comes first and then we have the problem that follows from what you said if you compare --

>> Ideally it would be things that are trustworthy but it's not that way.

>> EFFY VAYENA: So we have now a lively discussion. We will go a little over time if that is okay with everybody (overlapping speakers).

>> First of all, I do think that it is important to regulate, and we need that. But I also agree that it's not enough, because also those doctors that have their diplomas they make mistakes as well. So I also think that trust, because of that as well cannot be trusted, but and again, it's about liability and not about trust, and that is why, because the doctors even very well-informed can make different decisions, or mistakes.

So but that will be a problem of being a good or a bad doctor, but not comparing the two.

>> Very quick on it, how much can we also think about the fact that trust could be something dynamic, in the sense that this, the

capacity of collecting data, capacity of analyzing data and potential for improvement of all this artificial intelligence tools will evolve rapidly. The idea is how do we keep up with that trends and making sure that trust is actually something dynamic. (overlapping speakers).

>> It is cultural because as you said in the U.S. and Norway, but if you go to many Latin-American countries you have to see the person. You trust the person, not the machine.

>> I want to piggyback off of a lot of comments about liability. Especially in a lot of these apps that are marketed directly to consumers, how would the liability change to the doctor, for example, if there is a patient who has a runny nose and headache, that sounds like allergic rhinitis and the app might be telling them that but in the rare case that it's a CSF leak, who carries the liability for that?

>> I don't think this panel will answer the liability question but let's collect the last three questions and see what the comments are.

(off microphone).

>> Mine is not a question. I want to add to the discussion to structure it a little bit. We have a saying that if you hear a gynecologist that says is very experienced we laugh about it and test how many children he has, that is where his experience come from, but you are talking about doctors without demonstrating the

understanding that the kind of innovation that we are trying to put in place is to solve the very same problem you are grappling with, that where you have got 30 percent of a doctor making the right decision, is precisely because you have that doctor having to make, because he is generalist, having to make a decision on neurology, nephrology, when he is a generalist.

Now, and that is why in Matteson you have compartmental isation. In one of our hospitals in South Africa, doctor is a orthopedic surgeon but he does only right knee and nothing else. The reason he is doing only right knee is precisely for what you are talking about, that he is aware that he cannot be as good as AI and that is why he does only one knee. When you go to healthcare level and you have 1 percent doing only right knee and I'm a generalist, you are going to get 30 percent. When we begin to innovate, we innovate for that 30 percent. When you want to regulate, we get the guy who is doing the right knee, to regulate, and I go in between IA and the technology that's being used on the ground. So that you can have regulator tells you what exactly are you supposed to regulate, because the person that you are innovating for is supposed to do all these multi disciplines while he is not a specialist.

>> Thank you. Last question.

>> Thank you. Building on that, the question is how do we, how does one enable regulatory access, because if this was a discussion about medicines, it would be a very different discussion. The

questions about ethics, the questions about access, the questions about patient engagement, physician understanding, would be very very different. What we know is that regulators are not sure, they take a precautionary approach, do no harm. The regulatory system for drugs isn't there to prevent people from getting medicines, it's to prevent harm.

I think in a way, taking a page out of that mind-set in what is a very established regulatory system question, how do regulators in which ever level that needs to happen around this type of appliance, this type of technology, need help. Some regulators in some countries are very good at some of these countries, software is diagnostics. Some have no understanding of it whatsoever. I think it's not right, earlier on I suggested maybe W.H.O. should potentially take a leadership role in this. I don't know, thinking back actually, if that's correct, with respect. It may be that a national government has to ask W.H.O. by the way.

But I think there is definitely something about how does the community here and more broadly help the regulatory system I suppose think about how access can be expediented and how they can be part of the solution.

>> Brief comment though.

>> Actually, there is an observation here, which is I think one thing that is unique about these AI applications and technologies is that they have the capacity to measure their own performance on

a continuous basis. They can actually, you can record everything that they do. If you build them properly with the right amount of transparency, you should, the system should be able to say everything about exactly when it was wrong, with respect to some practical outcome, why it was wrong in terms of its internal models and interpretations. That factor is rather unique and could be used in some way in the regulatory process.

>> Do I have any other quick comment? Otherwise --

>> One on trust.

>> You have one on trust.

>> Just to highlight, very short point, when it comes to healthcare, you need trust to be coming from somewhere else, because we are like children. A small child trusts everybody. You need to protect them, because you can abuse the children. It's the same with healthcare, is that knowledge gap between people, patients and the doctors is so big, that many times the doctors can abuse that, but also AI companies can do that now.

That is why you need a pretty strong trust framework, because we need to be protected, because nobody knows anything about health.

>> Standard development for that will be crucial. You had something?

>> Maybe perhaps a quick note on what is going to happen in the next session. I think that is quite important to note. There are identified proposals that are being pitched since morning. I think

some of you who know who you are, otherwise we will seek out some of the ones that we have identified, and those who are wanting to partner with those individuals who have pitched, including if there are any financiers if you will, want to talk to, or entrepreneurs or individuals who want to culture and continue those needs to be grouped. That is one of the critical pieces of this. Would you like to go a little more --

>> I think so. We have heard some pitches since morning that have come on board. But we only give a portion to everyone who has not had a chance to pitch also. If you have something that you want to pitch out as a project, come grab a sheet, we have a template to fill in. See if you can get stakeholders involved. If you have stakeholders, we can discuss it in the 4:30 session. We start off discussing different stakeholders, trying to point you towards people who can benefit, whether it's funders, stakeholders in terms of technology. You have one person who is there for solution, goes to him for sure. But find the stakeholders. What we would like to do is by the end of the day is see which of these are more viable and how we can continue supporting these over the next year.

That is the objective. That is what we want to say before we go into break. We are already into break for 7 minutes. We start the next session at 4:00.

>> EFFY VAYENA: Thank you very much the panelists and all of you for staying over time.

(applause).

(break).

(standing by).

(standing by).