RAW COPY

AI FOR GOOD GLOBAL SUMMIT
AI + SATELLITES
COMBINING AI WITH SATELLITE IMAGERY TO TACKLE SDSs
MAY 16, 2018

Services Provided By:
Caption First, Inc.
P.O. Box 3066
Monument, CO 80132
1-877-825-5234
+001-719-482-9835
www.captionfirst.com

***

***

>> Hello, I'm from the United Nations institute for training and research satellite programme. If you were here for the previous session, you heard a talk from our manager, Dr. Elena Altieri. This is the poverty session of the sat lie imagery and AI track. Welcome to you all.

At our office, I imagine a team of analysts that work all day every day analyzing satellite imagery across different requests across the U.N. system, humanitarian support, security, Human Rights and other things. We have done there for the last 17 years, I have been in analysis at this point so we know a lot about satellite imagery and what you can do with it in terms of day-to-day operations of the U.N. system and of course also recently we have been learning about Artificial Intelligence along with many of you so we're excited to be a part of this event. I was involved in some of the organizing discussions on the is event and then I was asked to moderate this session so very happy to be here.

Yesterday there was a description of how the day two sessions would go, just to remind you, this is intended as a networking event intended to produce projects and partnerships. We have plenty of time for discussion here. We're really hoping that people come up with ideas and ask questions about ideas and try to connect with other people in the room to create partnerships to move these ideas forward.

To help seed the is up senior, we have brought together a fantastic panel to excite your intellect and we'll discuss poverty from an institutional perspective and how satellite data and Artificial Intelligence can support poverty elevation efforts again from the perspective of U.N. agencies and institutions, NGO, government offices, et cetera, et cetera. Our delightful Chair, Stuart Russell, he asked that we think ambitiously on the topics, not limited by the usual institutional, policy constraints we encounter in the U.N. and other systems.

I'm asking you to think in terms of star trek. Ask a computer a question and the computer tells you an answer. That's kind of the far off science fiction version of how we should be thinking, maybe more practically in the near-term. You think about how AI and satellite imagery can help you answer your questions instead of perhaps directly providing an answer just by itself.

I did see the discussions yesterday. I stole some tidbits to talk about from those sessions. An important thing, it is the Sustainable Development Goals are an immensal characterization length. These are kind of the eternal problems of the human condition boiled down to 17 Sustainable Development Goals. These are extremely hard things to meet. For example, gender equality, it is one of the goals, we really didn't do a good job on this with this track and hopefully we'll get better over time. It is important to remember it is within a framework of the 2030, these are long-term efforts which encourage you to think big, think about how these technologies are going to evolve over time. We'll talk about a few of the efforts here but there are a lot of other things out there that are not represented in the Summit or this room. I encourage you to search online, try to find other activities out there myself, other speakers in the room can tell you about other efforts as well, but there is an enormous amount of activity in this space right now and it is all very, very exciting.

Getting on to our speakers, so we have Bernhard Kowatsch who will talk about poverty and poverty elevation from the institutional aspect, how is poverty alleviated and how may data help. Then we have Marshall Burke, an assistant professor at Stanford and he'll talk about the seminal work he and his colleagues about extracting relevant information from satellite imagery using AI.

Then we have two fantastic representatives from within the U.N. system, we have a data scientist and project lead at UNICEF innovation and John Quinn, Artificial Intelligence adviser at U.N. global Pulse. Both folks will try to get at how an AI satellite imagery project develops within the U.N. system and why these are challenges and what their organizations hope to get out of the systems. Without any further ado, I shall hand the podium over.

>> Thank you. Before we start, a lot of you obviously wouldn't be here if you didn't know what the SDGs were, but just a quick test, who remembers SDG1? A couple of you. Okay. Another test, who doesn't remember SDG one, but good one, what's really -- what I have been asked to do, this is really a basic factor, what's poverty mean, why is it relevant and why are we talking about it with AI. A lot of times people think about is it something that AI is really helping you, a story to make this more trankable for those of it you that may not be on a daily basis exposed to the topics, because let's face it, if you're not coming from a background where you have been exposed to poverty, where you have been in a situation like that, a lot of people, it is a very remote problem when we talk about poverty. It is not something that maybe some of you experienced with the family, previous generations, but in the work, in particular, when thinking about the SDGs and to terrifying radoms indicate poverty once and for all, it is an important step to really think about how the empathy, how the people that are suffering from poverty, how they can actually be impacted. In the SDGs, which are the SDGs, poverty is goal number 1. Obviously it is interrelated. A lot of other topics we're talking about, access, having enough food, having access to water, a lot of aspects that was mentioned about gender equality, the interdependencies from the community SDGs, in SDG number 1 we want to eradicate extreme poverty, that's an ultimate goal with several sub goals related to that like people having equal rights and economic resources and I'm not going to read the whole deaf definition but what do we actually really try to achieve here.

Now, showing you an example here, and this is a study that's been done by colleagues at the World Food Program and what would it cost you to actually buy a plate of food in different places on this planet. You may not be able to read this but essentially if you were buying a plate of food in New York, if you were in South Sudan, you would have to spend 155% of your daily income on that one plate of soup. Let's just -- this essentially means that you have to spend more than your daily income rather than a plate of food. Obviously the facts of that, it is severe. Meaning that Realistically, what's it mean. You can't even afford a plate of soup a day right now. This is something -- putting yourself into the shoes of people who are effected by that, obviously in this particular case, it is not only the poverty that's income, agriculture, access to food, to prices, it is also about economic opportunity, maybe because displacement and war, you actually have to deprive those economic opportunities.

In the -- how are we really closing this gap, AI is one of those elements of how we can actually address these, not the only one. I think it is fair to say that of this, AI is not the single silver bullet we have changed everything and without looking at any other trends, there may be other technologies, other business models that we can leverage like blockchain, mobile app solutions, they'll have an immediate impact and as a system may be able to help us address this.

I didn't want to talk just about the basic elements of pure poverty, but let's look into what it actually means when we talk about satellites in AI. This is referencing a project we hear more from other panelists about the specific AI and satellite use cases. This is an example of road construction and using satellites, you can look into the impact of crop lands across the row. You see here, this is a picture of before the roads were built and then afterwards. Actually by satellite images, there is one of the elements, a lot of the work that's going into an infrastructure development, World Bank, others, international NGOs and governments work oing with the infrastructure development. This is interesting, now using satellite, Artificial Intelligence, we can -- we have a cost effective way of not only monitoring it, but also reporting on some of the benefits. We were talking about is this road that's been built somewhere, actually have having a positive economic benefit. Obviously the crop land itself, it is not yet telling you the exact income of the people that are there.

Another example, when thinking about Artificial Intelligence, we can use it in particular in poverty elevation and it is immediate shocks. It could be for how could we enter fine when it is about poverty elevation early on and to actually use it to identify structures, other aspects will hear more about the satellite use cases as well, and when we're now -- and this is because -- this is now we address this, Artificial Intelligence, obviously it can have very positive, but it could have very negative affects in terms of when you think about poverty in particular, you think about what's the number of jobs that are replacing the Artificial Intelligence. There is also other opportunities than to just highlight -- this is slightly off site, this is where an example of -- that's a product that's supported, virtual farmer market, and in reality, this is not right now using Artificial Intelligence, this is just a marketplace where farmers can sell crops. Very basic technology, it is essentially working with the mobile platform. The way you can change this, using satellite, the data that you actually uncover, it is using price information, informing people about what is actually the fair price that you have for the crops that you actually are selling if you're in surplus production. You inform people about how you can get access to information points and weather-related data, also farming advice for instance and think about the future where knowledge, as we know it right now, that can leverage that data in particular, it is not only accessible to somebody sitting in an office somewhere in New York or Geneva, but can be at the fingertips of somebody that's low cost, someone who may not be able to even read and write. Personally, I'm really inspired by the power of combining remote sensing and what's recognition, actually voice assess tans, talking about -- all of a sudden, we can address social issues like people may not have access to education right now and it may be sufficient for them to get quick grips in the knowledge economy using the benefits of satellite data and thank you very much for the attention, happy to talk more about that.

Thank you.

>> MARSHALL BURKE: Thank you for inviting me. I thank all of the folks that contributed to the work.

I corun a lab called sustain lab, sustainability in Artificial Intelligence lab, what we try to do, we use various sources of non-traditional data and in this case it is going to be satellite data to measure development outcomes on the ground, outcomes relevant to a lot of the Sustainable Development Goals we have talked about so far. Before starting, I want to take a slight step back and think about -- this talk will be about how we can use satellites, satellite imagery to better measure development outcomes on the ground, particularly poverty, we'll step back and say why is that a good idea? Why are we talking about this? How did we get here? Let me do that in four steps. Step one, and people have made this appeal in many ways so far, it is how do we normally measure poverty or related outcomes. The ways we do this, through household surveys, we go door to door, we enumerate long surveys asking people about assets or asking them what they consumed over the last year. We add that up, if it is less than 1.90 a person a day, we call that person four. These surveys are great, they reveal incredibly rich data about the ground and they're very hard to do. They're expensive. As a result, we just don't do them often. Here -- the best visualization we can make of the nationally representative household surveys measuring poverty or assets available on the ground in Sub-Saharan Africa in the last decade. What we find, in the median country, there is one survey over the last decade that's in the public domain that measures nationally representative samples and that can allow us to estimate poverty or some sort of asset wealth index of.

We really need the indicators we're not measuring.

That constrains our ability to do all kinds of things, track the SDGs, to monitor various programs and to evaluate interventions to figure out what exactly is working.

Everyone knows this. How can satellites help? Number one, humans, start with humans, turns out we're good at distinguishing well-being and imagery. Here I'm showing you two images. Anyone know of the image on the left here? So the weekend home of our President, on the right is Google image search. I can show you this and in an instant you can tell which of these locations is richer economically if not moraly, and the point being that we're very good -- I can show you two images and you can tell the difference. Rule of thumb, if humans can do it, you can train computers to do it. Here is an example of computers getting better at image recognition, think of dog versus cat, hamster versus weasels, the computer vision task that computer scientists have been working on for decades and with the advent of powerful new models in the last few years, it turns out humans can do as well or as in some other cases as the plot suggests better than humans looking at dogs versus accounts, it seems that humans still error 5% of the time with distinguishing dogs versus cats in this image. Computers seem to be beating us on the margin. This is a recent development. We have only very recently Figured out how to train computers to do this. The way we have done, this we have labeled huge datasets, literally millions of images, it has dogs, cats, we feed them into the machine and they figure out what's distinguishing between the dog versus cat.

Humans can distinguish things. If humans can do it, machines can probably do it. In certain domains we have seen a lot of progress.

Third, we have a lot of new imagery to play with. Again, this has been shown in much more detail, but a lot of new sensors went up thanks to Europeans and others. Plenty of things to play with. Our question, can you use the powerful new computer image approaches that have been developed on the taskers I showed before and apply them to this huge amount of new imagery that we have available.

S. this is four examples of work we have done, relating imagery to poverty directly or measuring important constituent of poverty in places like Sub-Saharan Africa. This is from agriculture, not measuring poverty directly. What we want to measure, it is plot Level maize productivity, can we predict the productivity of that plot at the reasonable accuracy in Sub-Saharan Africa. Again, the hope is if we can predict that, the majority of the the rural poor in Sub-Saharan Africa, they're agriculture producers and this is an important part of their income. The outcome we want to predict, we're using the 10-meters here, we have also used planet data, we have used various sources, land, we'll show you results on that in a second. Turns out that depending on the scale of the plot, the sensor doesn't matter much, and this is not fancy AI, this is relatively simple, something more like statistics or basic progression. What we find, satellites are off the shelf satellite vegetation indexes with important tweaks given what we can learn from the new sensors can explain a lot of the variation of plot-level maize productivity. We have examples of this spread out over Eastern Africa mainly showing, again, at the plot level, these are tiny plots. These are a tenth of a plot, the size of this room, we can make an accurate, surprisingly accurate I would say measurement of that productivity of that plot using imagery alone. Because we have this imagery everywhere, we can make maize maps and what you see on right here, it is that you have the scaled version of that. We apply this algorithm to all of the pixels and make plot level, pixel level productivity hubs and that's Example one.

Let's now try to measure poverty directly. I'll show some examples of this. Here we have measurements and we're very dependenn't on the ground data that's available, that's often collected by organizations like yours in the room. The agricultural data we collect ourselves and we have done this. The poverty results I'll show you, the ground data we have, it is off the shelf publicly available surveys. What we try to ask here, can we predict at a village level average consumption expenditure, average asset wealth using imagery alone. You have the village level outcome, the input, it is going to be various types of imagery, three meter data, we have gotten more fancy lately and now we'll use as the model these fancier neural networks, more modern computer vision approaches to extracting information. We have done this in Sub-Saharan Africa and we have results from India as well. The results, they're not perfect. We explain substantial variation in village level poverty outcomes across rural and urban Sub-Saharan Africa. This is before, we have done this now in 30 countries and working to scale it up. Depending on the outcome, we explain roughly a third to 80% of the variations from the village level outcomes using satellites alone. These are predictions I should say where the satellite has not gotten to see anything about that village. These are predictions for data that the machine has not been able to see. This is a fair test of how well this thing would work if you apply it to new situations without training it. Then you can make pretty maps. If you're willing to believe that this works well on the village level, you can scale it up and here is an example from Nigeria and we have extended this to the folks in the Nigerian government and this looks reasonable to them, whip I guess is good. That's Example No. 2.

This is an earlier version of the slides. We have is it in India, we have done it there too and the results there out perform what I just showed you in Sub-Saharan Africa. There is one constituent of poverty, agriculture, a direct measurement of poverty, here is another measure that we think of is very important of poverty itself and relevant to other SDGs, so can we measure from satellites various infrastructure measurements. Here is results for electricity and we have done it for water access.

Here again, we have village-level data and our input here, we have used the land sat results I showed you, imagery and we have it everywhere and going back in time which is nice. The model here, again, a fancier CNN. What we see, this plot, it is a little bit complicated but the headline number, it is up there. We have looked at outcomes for electricity and we predict with 85% accuracy whether or not the village has stable access to electricity. We don't only do well with electricity but with piped water and other baseline models you may use to predict access.

I'll skip over these. These are much older slides and I'll finish up there and very happy to discuss any details of this project in the discussion.

The take homes from my end, I think we're better at doing this. We're in the early stages, my view of being able to use satellites in the new more complex vision architectures to be able to measure these things on the ground but we're better at it and the next few years I think we'll see huge growth in our ability and other folks' ability to make better predictions here. In a lot of settings, it seems like satellite-based predictions are already comparable to the quality of measurements we typically have from ground surveys. I didn't go through that, in agriculture we show -- you go out, do a survey of agriculture and the observations, you have noise, and we show in some of these studies that the measurements you pull off the acceptsers are often no worse and sometimes better than the measures you have. The satellites can be better than off the shelf survey based measures and a big challenge is two-fold, in.

One, scaling these things moving from the pilot results I showed you to scaled products to indicators as Andrew -- as Andrew mentioned in an earlier session. Scaling these things up and making them useful to folks. Here, you know, working in the Ivory towers, it is where we really depend on interaction with folks like you. People who have specific questions, well defined questions on the ground but they need answers to. Where these sorts of estimates could potentially help. As well as folks on the ground who may have training data that will help refine these models, whether it is globally or for the specific instance they're interested in.

We're always looking for collaborators on this, please come talk to me afterwards if you're interested in collaborating. We have a stable of amazing graduate students that can work on this, free labor often, talk to us.

Thank you.

>> You have instructions on the table in front of you on how to submit a question if you want to do it that way or we can take questions after.

>> Thank you for having me here. It is a pleasure. I'm from UNICEF innovation and I'll talk to you about some of the work that people are doing around satellite.

The world is facing some new big challenges that are affecting all of us and becoming more and more global. I think we already talked about the challenges that we're doing today. I'll quickly go through them.

So 75% available population will live in cities. Displacement, there are 55 million children that are displace it had due to war or violent conflicts. We have Climate Change, Climate Change is affecting millions of people, including almost 500 million kids with food and security issues. Training of tomorrow, many kids are going to go to school and they're going to come out and they have to find jobs that don't exist today. What do we do with that? Pandemics, you do the global networks, connections, this is a spread that's faster than they were before.

You have all of the global challenges that are growing and you also will have technologies and industries coming out, the theme is trying to find the opportunity in between these needs and the industries, in particular, we're also looking into how satellites got help doing this. How satellites helped poverty and also tried to fight poverty.

I'll give you some practical examples of different things we're doing around satellites and poverty. So as was being said, the first thing we need to do, we have to map poverty, where is the most vulnerable people and how do they look like. And this is particular important in places like Europe and things change quickly. We show the image from 2012 to 16 and you have the changes in populations.

Unfortunately, the last sentence, it is from 2014 and had you have the conflict. We're trying to explore together with the Government of Iraq and the mobile phone network operator over there to try to combine both the mobile phone data and the satellite imagery to map poverty, not frequently and more quickly. However, poverty is a complex thing and that's -- we have the poverty index, et cetera, we're looking at the different factors that contribute to poverty. One of them, it is nutrition, food security, nutrition and for that, this is a key element. In Iraq we're looking a the how satellite images can help us to map that information piece of the poverty. We're doing this together with others and we're also trying a similar approach in Malawi, but over there we're trying to combine the power of higher resolution droning as well as satellites. We have had a drone corridor in Malawi and we're trying to explore and look into whether the higher resolution can help improve the methods had from satellites.

On these issues, they affect mostly specific people, those that are vulnerable, refugees, they're on the move, they have to be displaced to new locations. So we're looking into mapping the changes and the economics of this as such for refugee camps. An important part of these priorities that you're testing, it is crowdsource science, so we're putting these challenges to the Open Source and looking at whether the different scientists, others can contribute to solving the problem.

Once you know how to do these populations and how it looks like, you can start to look into the infrastructure and to make sure these populations have access to basic services like health, nutrition, et cetera. An example of this, it is the school map initiative that we're trying to combine resolution and deep learning to map schools around the world. We have started this work and we're trying to expand it in other regions.

These are dialogues, proofs of concepts, the question we face is scaling up the initial results and making them operational. As we were saying, this is a global problem and we have more vulnerable people, but we still don't have the capacity to take this proof of concepts and put them in practice and scale them.

I want to look at examples of what things you think should be done to improve this. The first, infrastructure, you implement the initiatives, you look at the satellite data taken it is sometimes put together, like the solutions, not to access but to look at the competent power to produce this huge amount of data and then the knowledge and science to create the models.

I think we should create the partnerships around all of this so that it is accessible and available for everyone to easily use it and apply it. This is one kind of example that everything has developed, it is an open-source platform where different partners and public, they can collaborate and come and provide datasets but also the science, community from researchers, academic institutions and you can collaborate to develop the models and create these for human kind purposes.

All of this, it was more external collaborations, but we also need to collaborate internally. Again, I think we already gave examples that were overlapping with each other, whether it is mapping, how do you make sure that we're creating this conversation and we do the work, you work together to solve these mobile challenges. For that, we have a new network that's been created to enable conversations submitted between different U.N. agencies as well.

I think that's it. Thanks a lot.

>> Great to be here. I'm John from global health. I would like to share a fun practical experiences we have had working on satellite imagery as it pertains to poverty and I would like to bring up a few of the limitations we perceive with the current methods to understand what we feel is possible at the moment and what is still difficult. We have been dwelling very much on the guidance we have from the AI researchers needing meat, need to know what's still to be done and that will hopefully address that. So we started getting into satellites and poverty through the lens of identifying buildings and roofing types, so as the global speakers touched on, poverty, it is a complex set of phenomena meaning different things and different ways to compute this.

The type of roof of the building, it is an important proxy indicator for well-being so if the roof, if it is that muched, metal, tiled, that tells us something. That's the kind of thing we can detect fairly straightforward with the existing object session so that they can say this is a village in northern Uganda and this is a thatched roof. To do this at scale, we understand the consequences of something like roofing types, we have the house surveys, censes capturing many things on the household, how many blankets you have, meals a day you're getting, access to education and the commitment indicators that come and expenditures. What we have been trying to do is quantify these things, it is a kind of thing you can give to an economist and they can use these numbers and where we understand the relationship with others as measured through this. This kind of thing, we think it is an interesting potentialsly quite large-scale used source of insight moving forward. I would like to draw attention to some of the few issue, or dwelling places, things that are small relative to the scale of the image. Much of the experience we have of detecting structures, it has been in the context of refugee camps, we have been working for sometime with the teams and they have an impressive analysis operation with people who understand deeply both remote sensing and the conditions on the ground and are able to apply various types of contectual reasoning, to be able to resolve the uncertainty, the ambiguity, a small blob of pixels, knowing what's happening on the ground. A few examples here, we have seen lots of pictures of buildings and processes, let's have a few more of those. These are results which relate to the methods of sections and we have are tried many, many different types of models, object detection methods, different sliding window methods, a number of different things, this is the best performing, these are structures detected in South Sudan and we can get a reasonably high level of accuracy. You get about the 95% precision and it sounds good but it is often not good enough. One thing, particularly when dealing with vulnerable populations, places in vulnerable regions, a small number of pixels may represent some important stuff. That 5, 10% error margin could sound good and look good when you look at some examples, but actually that could be an unacceptably high rate of error. Particularly in a pricing situation.

So what we have been looking at, how to get that up to the final level of accuracy, which is needed in this situation. We have examples of difficult stuff, where the models breakdown, sometimes these are pretty tricky for the current methods.

Let's think about where that 5 or 10% improvement that you really need to   get to to have these results be useful, where that comes from. Apparently this is due to the variation inherited, this is something that was right from the beginning, you can get something working beautifully on any one particular example but when you try to scale it up through multiple satellite sensors, different climatic conditions, different times, different times of day, then it becomes really difficult. There is a lot of variation, just in what the building looks like and what looks like a building in one place is what a stone or something else looks like in another place.

Understanding the context is I think really an important thing to look at going forward where we have the current is detection methods assuming all of the information that you need is in the pixels and when the number of   Ials is quite small, think of the bits of information in the four-pixel block and it is quite well and that kind of reasoning, that kind of mobility to use contectual knowledge, it is very interesting interactions.

We end up with things like maps of structures and we can compare where the detection structure is and where the experts have said that there were actually structures, it looks pretty good in practice, not good enough, another interesting thing I think that's kind of the methodology, looking at interactive, augmented, methods where we list the minimum amount of information from a human expert in order to get the right answer. So we have various models for this which I think are interesting to look into and then our objective can be how can we get the results of this efficient accuracy with the minimum amount of input from human experts. What is the most efficient use of that human experts, time perhaps measured in the number of clicks they have to give to get something that's accurate sufficiently in the emergency situation.

A couple of other points here, and just kind of what's coming up. We're interested in slum analysis, you have been looking at this, we have started to look at this for water provisions to understand the water supply and how much has been provided to slum inhabitants versus other parts of the city and going forward, this is something that we're looking to develop over the decades, how is the city itself developing. Of course, you can have all kinds of stuff, it just doesn't have to be built like this.

So just to wrap up, you know, a few comments on what's on the horizon, I mentioned this aspect of understanding the context of what's going on, something else, it is.

Coulding up for us, it is trying to use the data sources from the ground to have this multianalysis and we just had a talk yesterday talking about the radio mining system and we have been looking at understanding social media and this is the main social media to get some information on what's happening and what people's concerns are and reports of events, these are under reported and the mainly use here, this is useful for this and we can pick up the rumor, even the mention of disaster which is economically released vent. We have people talking on radio about things like this road has been washed away, the places are disconnected and one of our favorites, elephants trampling in the crops and this turns out to be a big deal, even to the extent which people are discussing what is the best type of poison to kill an elephants. You get a number of quite interesting and illuminating possibilities of things happening, which then can be confirmed with something.

Imthe satellite, it is refreshing in a sense. When you work with anything to do with people reporting stuff it is subject to all kinds of chills and misinformation campaigns and propaganda and what have you but very difficult. If we use that as a hypothesis generation and then you use the satellites to confirm what's happening, that's interesting.

I think I'll wrap it up. Thank you.

>> Fantastic, what's the best kind of elephants to kill the elephant -- just kidding. We have questions on pigeon hole, is anyone feeling verbal and want to ask anything? Love you all. Nice computer people.

An initial question from Stuart, how well do the networks transfer across cultures and geographies, for example, does a CNN trained on African data work in DL? I believe that's an exam question for you, sir?

>> Great. The shorter answer is, we don't fully know, we know within Africa, we have done the experiment in Africa and we have done it in Tanzania, applied it in Uganda, when you do that, the model performed less well than if you had the data to perform in Uganda. It performs well with but this is true through the settings where we have the data to look at it. We have not tried this model and applied it to a South Asian context, it is a good idea, something we want to do.

On the agriculture side, a lot of agriculture indecks that are indicative of the productivity, those are cross setting and you take an off the shelf model in the U.S. and ply it and you get the variation mother or less right. For a lot of the constituent of poverty, we think the underlying stuff we're picking up on the imagery, it is consistent across settings in a way that may help when you go to an entire different cultural setting. The full answer, we don't have a sense of how well this support is from across very different -- it is a good one.

Just to expand.

For those of us that don't do this, you point the computer at a different country?

>> Yes. What we do, we don't let -- we start in countries where we have data, and again, please come talk to us, we're limited, so again, I think I said before, the constraint is not imagery but ground labels and data, we scramble to get every last label we can find. Come talk to us. What we'll do, in the small set of labels we have, we'll divide them up in some that the models see, others that don't, we divide it by countries so the model sees Nigeria, not Malawi, we point that model to Nigeria, train it, get it and apply it to the images in Malawi and see how they do. That's how we evaluate this. Again, we're label constrained. In Nigeria, they have 200 million people and we have 2,000 labels, 3,000 labels in Nigeria. It is a fundamental constraint. The results I showed you, it is based on an extremely limited set of labels. This is the fundamental constraint of proving the models. Come work with us.

>> You have anything on the transferability of poverty analysis? Not on the poverty, we have tried to do that with the schools from Liberia and Colombia and the first didn't quite work I must say. With that, maybe most are more rain foresty types of areas, so we gave it a try. I agree, even on the context, it didn't work better or worse. So something that we haven't tried, well you're tray training to a different context, you have the output and you work with both of them. For

example, you train the data from Malawi as well as service station and maybe the output will be able to do more or less. We have want tried it yet. It is something we want to test.

>> Would you mind giving more information on how the school mapping kind of works at the technical level?

>> We're using the networks and we tried two different approaches, I explained it yesterday as well so it may be repetition for some of them, sorry for that.

We tried both as you were advised and unsupervised. The supervised caseworks for the cases where you have actually labels because you have good data quality ton that p in that case, we created the classifier that takes labels from schools and tries to learn the patterns that can be from the playground to the local Chair of the building to the community.

With the CNNs, we have a classifier that it will tell you whether you're going to learn or not.

The other one we tried to have, the supervised approach, we'll have labels in many cases, so for those, the worse case scenario, we tried to train clustering techniques to see if we can group different types of things and similarities to see if the schools are in similar groups as well.

>> It is all based on imagery, no additional --

>> Yes.

>> So if it is a very poor community, the school house is essentially a one-room sort of structure, you're probably missing things like that or is it -- do you bring that in in a different way?

>> It is a very informal school, it is happening in one room, we probably are using that. We're using the poverty school index usually with the funding from the government and the location of all public schools in the country. That's what we try to use for the training.

>> Interesting. John, I have enjoyed torturing you with the moving from one image to the other. What's been some of the biggest issues -- how disappointed were you the first time you tried to run the model on the new image. How do you kind of get past that initial barrier?

>> I think the disappointment doesn't occur straight after running the image. You run it, it looks pretty good. You know, you get nice looking pictures of what seems to be most of the images that are generated and it is really -- you know, when it.

Coulds to actually computing the numbers, I know we were looking at something like 97% recall and sort of falling a bit short of that.

I think -- I mean, this, what we're trying to do, it is learning from many examples what actually can be expressed by some fairly simple reasoning rules. The analysts on the team are very good with shadowing, how there is internal sha he dos and there is no roof for example. We can learn that with these models which we have the examples, there sometimes are and in this kind of situation, and it can sort of try to learn just by seeing many examples, the internal shadows, maybe we can see that, it is an active dwelling place with the roof.

I think what is -- this is what's tricky, you try to build in something of that nature into the models which are inherently boxes that convey that kind of reasoning. I think that's difficult to have that moment.

>> Interesting.

Then we have a question on labels. If you're using pictures, for example, there is obviously a bias in where the pictures are from. There is probably endless pictures of cats and dogs, much less pictures of what you are actually looking at. Can you talk a bit more on that.

>> A project that's accelerated that we're supporting, tech for food, it is actually something where we train vulnerable people, people that never used a computer before in.

Abouting part of a global Digital Economy. So we have -- with some pilots, now with Iraq, young adults, they would never touch a computer before, after 8 weeks of training, they're multitasking, like amazon, charts, data cleaning,ish picture levels and so on. What we have done for our drone AI project, we have used these images and actually have had our students while being trained to actually label the images so they're doing labeling, the images, why we're already doing that, and I must say that in our context at least, a lot of times when talking about identifying like buildings in a disaster zone for instance,ish it will look different, we talked about thatched roofs, metal roofs, if you're in a disaster zone, a hurricane looks different than had a landslide or an earthquake and these other types of things, you have the act Raleigh training data, can you reuse things that are similar? Yes, I agree. In particular, if you're talking about models that are supposed to work and elevate the aspects of what you're to do with it, maybe it is just to say that it is not only about the training, that's what you were alluding to, if you have 97% accuracy, it may not be good enough. The request he is, what are you doing with data? Is the data an input in the broader process, this is typically when we look at vulnerability analysis, this is something that we take about an input into other layers of information and household services on the data.

I'll plug one of our own projects briefly, we worked on crowdsourcing and the photographic harvesting for a few years and we have a project called geo tag x and it lives on in the project called into mc, a collaboration with the University of

Geneva and uses a python library called pipe also that you can basically develop into a system so the crowd volunteers, they can basically go out, harvest photographs for you. It is possible to basically perhaps compensate for some of that bias that we find in the available media.

There were audience questions. Take it away, please.

>> Sometimes the food, it is there, it just doesn't get to the people that need it the most. That could be in the poverty stricken Zones, it could be in the disaster related situations, et cetera. Let's look at all of the wonderful stuff you have presented. Let's put a logistics and movement lens on it. How is the type of work that you do help us with that kind of a problem? Sort of looking at movements of goods, movements of food, logistics cross-border and whether the food gets to the people that need it the most and can imagery help us with that? Thank you.

>> BERNHARD KOWATSCH: I think -- let me first say, I also mentioned when speaking before,ish I think there are other technologies that are also helpful when you think of poverty and hunger in the core of our work.

I think logistics challenges, personally, I would see IoT, Internet of Things, sensors, blockchain application and mobile solutions, helping us way more in that particular sense of talking about supply chain tracking or network optimization.

That's not to say that you can't use satellite data to track -- to have other things, you could probably mention congestion at ports, cross-border, things that's definitely working. I have thought about the costs of things, how you're actually able to -- can you identify a track on top of it, is it that track or a other one. There are challenges. To a specific question, I think it is different layers of the honest transport and optimization, when talking about availability of food, we saw examples of how to actually track production, crop production, all of that. That's where it is getting more interesting and, yes, you can also then look into maybe not just assessment right now, but also predictive models and things that are -- where you can't see right now what will happen in the futures, you can actually work on more systematic type -- when talking about social safety nets, prepositioning of foods and what it can actually do to address this. How is using publicly available data from satellite imagery with an incredible if degree of recognition that you just mentioned, what we can observe, it is more expensive than instituting digitalized supply chains, blockchain across the Least Developed countries and emerging markets? I would be curious to hear from the gentleman from Stanford University, you say you play around with a lot of different models and look for challenges and examples.

Is this something worth looking at, from the movement of goods, food or other related, freedom of transit for example and what extent it is actually respected as you move across the different countries, is that a challenge or is it already something that maybe one of your teams is looking a the?

>> We have talked about folks that look at exactly that. Imagery may or may not be the best tool for this. There are other indicators that you hear about here. If food is scarc, the prices are high. There are other ways to think about measuring prices locally, via the mobile phones. I think it is a great question. We're happy to throw imagery at whatever problem there is, if there is way to evaluate the predictions, if WFP, you, whatever, they have training data on the data scarcity and we're happy for our students, models to look at that. This is a specific hard problem. Whether it works, I don't know, you can imagine things you would look for. Hilled be speculating whether or not it would --

>> There are fundamental issues with the imagery. The height resolution imagery where you're able to see an enormous amount of detail, if you're lucky, maybe you get one of those every few days. Planet has imagery collected every day, but perhaps you can't see the level of detail there where you can concisely identify this road is out of service because of landslides, it just cut it off or something like that. We have done the network analysis after a disaster, for example, also during conflicts when you try to see the roadblocks and checkpoints that have been set up that are going to impede access. Pull get the imagery collected often enough to get that realtime network analysis. I think we're not quite there yet in terms of imagery availability. There is annoying details like the high resolution, it is collected in long strips of 200 kilometers long and they're north to South so if the road is east to west, you're at a bit of a disadvantage on that. It is something we're looking at. We have talked about not analyzing the full road network, let's keep our eyes on certain markets. If the markets continually have vehicles showing up every day, which is something we could see in the daily imagery, then we know that market is active. If the market suddenly appears empty, then something else is going on there. It is -- I would think it is one of the hard problems since the field is somewhat in its infancy from the AI perspective, there is a the lo of moving Parts to it that it is something that we'll do manually hat UNICEF. We did one for WFP a few weeks ago on a long road and in the Congo to share the imagery with them that was a 75 gigabyte image download so it is a big amount of imagery, a bill challenge that way. It is something that -- a big challenge but it is something that we talk about.

>> Back there.

>> I think this probably echos the question, you went to the floor so I thought I would bring it in. Talking about the misuse, the negatives of this potential earth observation which is fantastic and you're doing great work to try to elevate poverty and meet the SDGs. I wonder if any of your organizations have done any work thinking about which information, whose hands this information may get into, including the media, how they may perceive communities, migration patterns and how it plays to a negative, a potential negative for community and society, and if not, who do you think is best placed to take on that role at least thinking through potential negatives of this information becoming publicly available.

>> You're touching on an important point. it is something that's very on top of people's mind when doing that type of work, at least in regard -- I'm sure others will say it too. There is a difference between what you will make available publicly, there are examples, other colleagues have talked about where data has been available publicly, has been sanitized, also has had negative consequences for population groups, so on.

I think -- this is -- it is data privacy, but it is identification, what are you able to share. This is where some of the challenges come in this terms of publicly available data versus what's not publicly, what happens with the analysis, how do you share the data. I think one aspect that's important, it is to further data sharing agreements, who can share which data with whom. Actually respecting the privacy and also with the specific protection thought in mind.

>> I think that's really an excellent point. thank you for the question.

This is particularly an issue where you're dealing with vulnerable populations. That's related to a point which came up in the previous session about information and balances and it may be that there is a lot of information about that. It is in their hands. I think the data that we have tended to deal with is in the 50-centimeter level. Even with that, there are important privacy concerns, but this enters another dimension here that we start looking at the data, we start to see people and perhaps identify individuals. I think we need to think about that to be ready for it. Our approach in global policies, we do a privacy risk assessment for each project we would propose and think about what data is going into that, what will come out, are the privacy risks proportional to the benefits. That's something we have been working on having a framework of, which is available at the moment to help to judge some of the risks of that sense of proportionately.

>> We do a very bad job I think within universities thinking carefully about this. It is a critically important question. I live in fear of creating a data product that -- it is easy to imagine the good things you can do with it, hard to imagine the bad things, I think we don't have a systematic way of an internal mechanism to do, it is really important, honestly, we don't have that at Stanford right now and other organizations, universities don't have it really either. It is in a systematic way, there is human subject review and that's for a specific research task not for generated products and we need better guidance on how to do this.

>> With all of the things that have been said, two types of data, the satellite data itself, it is owned by -- we have much to say about that, but there should be regulations from governments, et cetera, and the other one, it is the insights we may get from the satellite imagery, I don't think they're different from the type of data that we have been used to deal with for many years. We have performed household service work for so many years and if we do one with data, it is a long time and we have -- it it is the processes, so from that, we have all of these processes already and we have shared, it is aggravated then and we don't -- we have the VBU in the vulnerable populations.

>> I'll zero in on the bias part of the question. At UNICEF, we do traditional manual analysis and we use a lot of automated methods to support us. At the end of the day, it is one of the trained analysts reviewing something. For example, if one of the more common things we work on, it is floods, floods are the most common natural disaster. Water is very visible, it absorbs most -- a lot of energy. So at a certain level, it is quick and easy to extract that water from the satellite image. However, if you just do kind of the 5-minute solution, you have it done automatically, you're missing a lot of the individual pools of water that may be just as important in flood situations. We have had these debates at UNICEF where it is -- we could cover so much more if we do the 5-minute solution, but then what if your house was in one of those puddles of water that we didn't locate? What if a small pond had formed in the road that you were trying to get down to get to your house? So the bias issue for us, it goes directly to the -- to how we do the analysis.

I see AI generated products a a trend analysis at this point, more than 98% accuracy, that the human analyst should be going for. I'm certainly not saying we're perfect at UNICEF so we get exhausted, make the owe comparable mistake and a lot of people look at the data and are very, very quick to call us out if we get something wrong. The bias for us, it is really kind of the critical part. None of our analysis is done on an area that there isn't some sort of a United Nations mandate to look at or countries themselves requested the assistance, we're less worried about the transparency issue.

We have a couple of hands up over here? In some ways the panel could be renamed from AI satellites to ground trooping or something of that nature. My question is, particularly to the people that -- the idea of what ground troop data is out there, what ground troop opportunities of collected data have been out there that's not been applied to the problems, and if there is a capacity for changing the collection of current datasets, NGOs, U.N. agencies that are out there getting ground troop data, if you make it more explicit, something of that nature, what is currently being collected to tweak to have models have more global ability to have the transfer of learning, what you have talked about.

>> Maybe to start, I think one specific challenge, even before another challenge out there, it is -- a lot of the analysis that we would do as World Food Program, where we operate and what the type of operations are. You have different data needs in particular when there was a disaster or a landslide, a hurricane, on going conflict versus maybe it is protracted conflict, maybe it is middle income country, it is mainly about empowering farmers to get connected to markets, something like that. A challenge, it is a constant challenge, what data is systematically available across countries rather than one. It is a challenge with different agencies, NGOs, companies maybe, and it is not that you have the same dataset in every single country and you don't need it, a lot of times, it costs money unless you're building an algorithm for the satellite images and it is available everywhere on had this planet. I think we're not yet there unfortunately.

I think it is a great question. There's a lot of data sitting out there, how do you collect the data, that's incredibly useful for this -- these sorts of problems that's not been Lynn rated. It is sitting in government's desk offices. We have ways to allow researchers -- liberated -- we have ways to work with that data. The census department, you go in a bunker and work with the data, you can do it. We predict privacy, we have ways to do this. I'm certain there are huge datasets sitting out there, so the map I showed at the beginning of my talk with the consumption expenditure, poverty surveys, none of that is is in the public domain, we can't access them. They're referenced at least to the sub district level and often, that's what you need. I think you're dead on, it is easy for me to say. Any survey that's done, they're done on tablets anyway, let's Geo reference them and have a way where outside folks can work with that data. With he know how to protect privacy in the setting, let's scale up those methods.

>> I think there's a lot of data but most of the times it is in a structure that's really hard to use. There is plenty of PDFs, different surveys that maybe they don't have an ID that you can triangute but you have the data and it is sometimes impossible to use it. Trying to systemize that, it would be extremely useful.

>> AUDIENCE: Thank you very much. I want to talk on behalf of women. I want to talk about the issue of education of AI too all the countries. Talking about all of the countries, there are those who are less developed countries and I want to touch on Africa region, where you have a plantation of tea, coffee, any other crop, and it is taken care of in by many countries, the farm, its being undertaken by the women, maybe because of illiteracy, or what, but when it comes to introducing AI technology, maybe when doing the harvest, any other thing, what happens to them, the human because it is the only way they improve their livelihood, it is through labor in the farms. What happens to them. That's one.

Two, how do you mitigate such challenges? I want to just request if there is any kind of survey, research that's been done on how you will get challenges that we shall face through -- from the us least-developed countries. Three, it is about AI purposely in this meeting, AI is for good, is there any other time whereby AI is for bad? Maybe I can touch on the first one. What I was talking about in the presentation, a specific example that was supported, the farmer market, it doesn't use AI as of now but essentially connects small farmers to traders, it is not a small farmer right now, but it is elite farmer, the person with the smartphone and aggregating from the small departments from around. I think what's important from that, smartphones will become widespread, that will -- maybe one day in the coming years, you will see that everywhere on the planet, you may have a smartphone -- I'm just -- you know, we can have green emissions and I think we're looking at the pathway that everyone will have a smartphone. Specifically for AI, I do see that it can be enabled more than currently people are most likely to read and write properly, people don't have the education and they have better accessibility. You can address this in multiways, one is a way that you have a person that's a had literal person with the smartphone, the knowledge, its the agent with kind of -- we have seen that model in a lot of countries. The other is using the AI, inclusion of people with and especially when you get into the voice recognition, you have the automatic response system, this is where I have seen lots of opportunity for even doing -- when talking about leapfrogging.

Maybe we don't have to wait until everybody can read and write. Maybe we can use the benefits of AI to integrate some of the more vulnerable people. I believe that we will go in that direction.

>> An issue relevant to this, it is the problem that we have, the lack of diversity among AI researchers. Where AI technology is being developed in a very specific demographic geography, then the technology that could result from that, it could be either not helpful or harmful in ways that we had not anticipated. So that's clearly something to address and I think we need to look at the composition of the panels within this very conference, for example, to see that that's the case. One thing we have looked at through our lab, it is having in that case as much of the team as possible, which is almost all of it, of people who are local native to the country. That really works a lot better. It is no the just better in terms of in principle, but you get better technology as well as a result. Understanding the implications of technology, on a community, it is difficult in the best of times. It really requires people to understand the context and be able to do that.

>> I think most of what we're discussing would essentially have benefits to the ladies on the tea farm, you're looking at how to predict when the crop is in danger of failure and financial mechanisms can be brought in, how to predict if severe weather conditions, something, may adversely affect them. If we get to the topic of robots harvesting the tea, then I think that's where you're thinking much more about a threat.

I want to go to the pigeon hole, we did get good questions. If an aid agency had replaced all of the thatched roof with tin roofs, does that mean that the village is no longer poor? Can we see causes of poverty, lack of education, water, transportation, et cetera, getting to the root cause.

>> I was hoping we would address that question. Yeah. Yeah.

The risking of miss independenterring this analysis, one of my -- misinterpreting this analysis, a favorite of mine, a friend from Red Cross was involved in analysis from an island, they were looking at roofing damage as an analysis of the damage of the hurricane. They looked post-hurricane at the entirety of this region. The majority of roofs were missing from the houses, looked as though a catastrophe and the Red Cross was extremely concerned. Then it transpired in this place, they're used to hurricanes, they have tin roofs and as a hurricane is.

Coulding on, they remove the roof. They wait until the hurricane passes and then puts the roof back on again. If we are to take these things at that level, indeed we could come to the completely wrong conclusion.

If -- to measure poverty, you look at proxies, if you look at the household surveys, the censuses and the information that's assessed there, it arrives after decades of trying to figure out the best proxies to ask about, and any one of those, indeed, they tend to be affect rather than causes.

I think when we start to get into the causes, that's a rather deep question, which is difficult to look at with the technology and indeed sometimes we may over technology with things, we have a risk of doing more with technology than we really can. The better thing is really monitoring at this stage particularly with satellites as we're trying to understand what's going on and the level of a few basic things which is visible from space.

>> Should they build a school, provide agricultural advisers, should they provide higher called seed with better yield, put in irrigation canal, what should they do? You know, can you answer those questions based on satellite data or is that too difficult?

>> I can give our view of this. If we use AI we have better measures and more comprehensive measures and then in a research sense, we can relate to the other potential drivers of these outcomes.

I think there is still so little understanding of the efficacy of the interventions, largely because we don't have measures of the outcomes. So if we have better measures of the outcomes, we can do those analysis that will inform the policy decisions. My view is that better measurement of the outcomes we're talking about here, they can answer the kind of question that you're posing. It is hard to answer it.

>> Maybe just to add to this: Coming back to a statement earlier, the comment earlier, this is not the only thing, particularly with mobile technology, we have implemented programme, we use a smartphone and C cards to replace paper, like the paper records with child nutrition, child clinics, by the doctors, essentially rather than just writing on a piece of paper, they just already key it in a smartphone. That information is actually giving you lots of insights, not only on outcome level data but it could also capture other rational. I think when you bring some of the satellite data together with some of the written data monitoring, other aspects that you may have, that's where the future and triangulating is and you create the different data points and that comes in. (captions will resume momentarily).