FINISHED COPY

ITU AI FOR GOOD GLOBAL SUMMIT
POPOV ROOM
PLENARY 6
AI FOR COMMON GOOD AND SUSTAINABLE LIVING

GENEVA, SWITZERLAND
8 JUNE 2017
1400 CET

Services provided by:

Caption First, Inc.
P.O. Box 3066
Monument, CO 80132
1-877-825-5234
719-481-9835
www.captionfirst.com

***

This is being provided in a rough draft format. Communication Access Realtime Translation (CART) or captioning are provided in order to facilitate communication accessibility and may not be a totally verbatim record of the proceedings.

***

>> REINHARD SCHOLL: Thomas, yes, our moderator is Thomas Wiegand. And the theme is common good and sustainable living.

And can I ask the speakers to come up to the podium? Yes, thank you.

>> THOMAS WIEGAND: So I have been observing some of you guys through the sessions and I have a request to you. Please follow it. Let's all for one minute shut our eyes, not talk to anybody, one minute.

Can you do this? Okay, let me start. Now, please.

(Pause.)

>> THOMAS WIEGAND: Thank you, good morning.

So we have a session now. I hope this shut eye for one minute helped. To quote a famous Austrian-german author, everything has already been said, but not by everybody. We try not to follow that. So we will try to bring additional information, new stuff to the discussion. And we are having this session. We are preparing the breakout groups for those topics that will follow this session. So we will have a breakthrough groups on common good, AI for prosperity, and

social good data following this session. We have breakthrough groups on sustainable living and promoting healthier citizens and smart citizens communities. Those groups are supposed to produce actual results that can be seen as an output of the discussion here.

What we are going to do is start with the talks. I'm delighted to have two great speakers. The first speaker is Dr. And Professor Fei-Fei Li from Stanford University. She is currently on leave. She is the chief scientist of Google Cloud. And among the many achievements, and I can't read them all to you, she has very early understood how important it is to collaborate in this field of AI by creating Image Nets and I'm looking forward to her talk.

>> FEI-FEI LI: Where do I drive the slides? Good afternoon, everyone. It's quite an honor to be here, first time in Geneva. What a beautiful city.

So I am going to share with you a little bit of what I have been doing in the past almost 20 years of my AI career. It is mostly focusing on computer vision and machine learning. Okay, good. So the visual system is one of the oldest and most important sensory systems for animals. In fact, more than half of our brain is involved in visual processing. It is being 540 million years of evolution. Today vision is one of the most important components of human intelligence responsible for many activities that are important for us such as navigation, manipulation, communication, entertainment, work, and food seeking and all this.

So animals from simple ones all the way to complex ones, almost ubiquitously have a visual system. In fact, a fun fact in evolution, that vision was said to be responsible for one of the most important evolutionary events in the history of animal kingdom, which is the Cambrian explosion 540 million years ago during which time the number of animals quickly exploded from just a few, handful of different species all the way to many, many more. So it really underscores the importance of visual intelligence as part of the overall intelligence system.

At the same time the evolution of visual machines has a much shorter period of time. Just like AI as afield is about 60 years old, computer vision as an important area of AI is about 60 years old, or a little less. So even though we have been progressing rapidly, the technology of vision is still just at the very brink of making important contributions to our society. We still have not yet given sight to most of our visually impaired fellow people. We have lots and lots of cameras from space all the way through the airplanes and drones overseeing our home, the earth. Yet we don't have a comprehensive technology to map out everything we are seeing. Or in swimming

pool scenario, today we still don't have a very robust system that can detect drowning swimmers. Or in healthcare there is increasing demand in terms of visual processing for diagnosis.

So many areas of human life and society today demand computer vision technology and the goal, I think, as a computer vision scientist is really to develop algorithms that can shine light on to this digital world.

So with that in mind, today I am going to share with you two research topics -- actually, in terms of the progress of computer vision, we have come a long way, as Thomas introduced earlier. Image Net has become a benchmark for the promise of computer vision and contributed to the progress of machine learning and deep learning.

If you just look at the image classification task from the first year that image, that challenge was rolled out, which is 2010, all the way to about two years ago, the error rate has been steadily decreasing, especially thanks to the breakthrough technology of deep learning, to the point that it is hard to distinguish the performance of a human versus the performance of machines.

In fact, in the field of computer vision, we have many areas of research not only in object classification but also object segmentation, object detection, human post, 3D recognition, scene parsing and so on. This is a very lively and thriving field working on the basic technology of visual intelligence.

Well, today with this audience, I would like to share with you actually two applications areas that we're applying computer vision technology to. One is in AI assisted healthcare. Another one is for visual senses. And then to conclude I'll also discuss a topic that is very dear to my heart which is about education and diversity.

So let me start with the first topic, which is what we call the Guardian Angel, AI assisted health care. Here is the Guardian Angel hypothesis. Artificial intelligence technology can help better the workflow of healthcare. We hear a lot about diagnosis using AI technology, especially in radiology, pathology, and so on. But there is another huge area of healthcare that is really important, which is workflow. Workflow happens everywhere from emergency room to operations room, from ICUs to primary care to pharmacy or home care.

So workflow is an area where patients, we care about the quality of treatment, the safety of patients as well as innovating technology to cut down the cost. Here as an example we worked on one really important problem. Is someone else driving my slides?

If yes, can you let me drive my slides?

(Laughter.)

>> (Speaker away from microphone.)

>> FEI-FEI LI: All right, thank you. So this particular problem is hospital acquired infection. It is actually a really prevalent problem in all healthcare systems. In America, one in 25 patients get infected by hospital-acquired infection which might have actually fatal consequences to patients. Costing America about 35 to $45 billion per year.

One of the most important factors for hospital-acquired infection is hand hygiene. That is a leading driver. The lack of good hand hygiene practice by medical personnel, from nurses to doctors, is one of the biggest drivers for this issue.

So it is a really hard problem to solve. In the past --

(Chuckles.)

>> FEI-FEI LI: What happens is for decades and decades we, the hospital system used in-person audits. We have secret shoppers that try to assess the quality of this. Recently RFID technology is used to try to replace human audits, but RFID is very coarse. It has a lot of low noise signal to ratio and disrupts the workflow.

With the recent advancement of computer vision technology, especially sensors that provide privacy, which are depth sensors, we have proposed a computer vision system that can track the movement of clinicians in the hospital without intruding in their privacy. So we don't have to reveal who they are but be able to track their hand hygiene movements.

This is noninvasive. It is continuous. It is unbiased, very, very cheap and powered by AI algorithm. This is a collaboration with Stanford children's hospital where in a particular hospital unit, this is the floor plan you are seeing, that we put these sensors close to the, in the hallway close to the hand hygiene dispenser. We can monitor the movements of the clinicians.

Long story short, we used a human tracking and then deep learning recognition system behind our technology and our performance of these hand hygiene activities recognition is actually really high, in better than many of the state-of-the-art systems.

Here is an example of what you see as the system detecting a person and his or her moment of hand hygiene activity. So one of the most exciting aspects of this technology is it's continuous. It's cheap. And it can not only, hand hygiene is just the beginning. In the hospital and patient environment there is many, many activities that requires continuous monitoring from patients themselves as well as clinicians. This work is opening up possibilities in using AI to assist the workflow environment of our healthcare institutions.

And these are the students and doctors who collaborated with us in this particular work. So now I am moving on to the second work I want to share with you, is actually census. Census is what government uses to get an understanding of what our society is. Government spends billions of dollars every decade or so to do a census. The United States just did a census in 2010, spending about a billion dollars per year. This is where we get basic information of demography, income, education, race and all this.

A few years ago, a couple of years ago my student and I started thinking about a really crazy thought. Can we do census for free using data, big data available around the world or in the country in America? And we thought about Google Maps. Google Maps not only have GPS information. It has pictures. It has millions and millions of pictures of our streets on Google Street View. And there is one object that is prevalent in every city, every neighborhood in America, and that object is a car. A car really has a lot of information. Every time you know a car, you can look at its make, model, year, and get a lot of information from that. So we thought of this crazy project that we are going to go to the 200 most populated cities in America and download Street View pictures from Google Map.

We downloaded 50 million images from the 200 cities in America. Then we are going to detect every single car on the street and use a computer vision system to recognize the details, the make, model, year of the car. We are going to use that to infer the social makeup of America's cities.

This is the project we did. We called it visual census, demographic prediction using 50 million Google Street View images and the cars in them. I'm going to skip what exactly we did in terms of car recognition. Suffice to say we used a deep learning system to first detect a car, the presence of a car and put a bounding box on it. Then use another deep learning system to recognize 2657 types of car ever manufactured by human society after 1990. And the recognition accuracy is really high. With that, we can get a lot of information about cars. And then we can ask questions like can cars predict a neighborhood or city income? And I'm going to just skip this and say we use a lot of features of the car.

But the long story short is on your left the actual zip code and income, average household income is the data collected by household by U.S. census. On your right-hand side is the predicted average household income by using the car information we get from the Google Street View. High correlation. Here is another example. Okay, I'm going to skip this.

We even used car to predict voting results in the 2008 election, Obama versus McCain. We can go down to zip code, precinct level -- is someone else clicking?

Here is a map of the Democrat versus Republican voting results by precinct. On the left is the actual result. On the right is the predicted result with the car information.

So in summary, here is something fun for you. What car predicts Obama voters? It's sedans, it's highly environmentally friendly cars. It's highly friend environmentally friendly city, it's Obama city.

With Republicans, they are highly correlated with trucks and SUVs.

With this data we actually -- can we go back one slide? With this data we actually can predict cities and their environmental, how friendly they are to the environment. We can predict crime rates, we can predict segregation of the cities. We can predict the different races. It's a very, very fun study using massive available online imagery data.

Here are the students who collaborated with me on this project.

Do I have time?

All right. So the last work I want to share with you is a recent nonprofit foundation we put together called AI For All, focusing on the diversity issue in AI technology. So we are all here because we believe AI will change the world.

As an educator, the question that really is very prominent in my own mind is who will change AI. If you look at the demography of technologists, computer science technologists today, we have a very alarming trend. In academia or in America, industry in general, we have decreasing trend in terms of the percentage of women taking on computer science related jobs or getting computer science degrees.

In Silicon Valley the companies, the major companies who have released their data in terms of gender are showing a very unbalanced trend in terms of the number of women working in these high-tech companies versus men.

And the picture for under-represented minority is even more grim. We don't have much data on that, but we know that we have very few African Americans, Hispanics and so on working in computer science in general as well as AI.

So about half a year ago, a little more than half a year ago the Obama White House held a science conference during which there was a panel. We discussed this. Why do we need diversity in AI? Just to put it very concisely, it is first really important for the labor market in economics. AI and computer science jobs are increasing, not decreasing. We need to include more people. It is really important to have creativity and

innovation in our technology. Study after study has shown when there is a diverse group of people working together, the results are more innovative and creative. It is important for our social justice as well as the reflection of our moral values.

So recently with the help of my former students and colleagues, we put together, launched a nonprofit called AI For All focusing on educating the next generation of AI technologists, thinkers, and leaders. And the mission for AI For All is to bring diverse voices of AI so that the next generation of the technologists can reflect the composition of our general population. And the focus of AI For All is through education programmes. We use, partner with universities to hold K through 12 programmes, mostly summer programmes focusing on AI education and humanistic centric AI teaching. We hope to create an ecosystem between students, current technologists in industries as well as the public and policymakers. And the inaugural universities that are participating in this programme includes Stanford, Princeton, Berkeley and Carnegie Mellon.

Some of our founders and advisors come from diverse background of entrepreneurs, Silicon Valley, CEOs, educators and policymakers. We are thankful that Melinda Gates and Jensen Huang are the seed funders for AI For All. Whether we are doing education or technology, it is just the beginning. It is a very exciting time for AI and computer vision. I'm happy to talk more offline. Thank you.

(Applause.)

>> THOMAS WIEGAND: Thank you very much. Our next speaker is Antoine Blondeau. I hope I said that -- oh, probably not. He is the cofounder and Chairman of Sensient, a San Francisco based startup in AI. In his career he was the CEO of Digima, a company that powered DARPA'S project that ended up to be Siri. So, Antoine?

>> ANTOINE BLONDEAU: Thank you. Can I use this? Perfect.

Hopefully the little clicker here will work better.

Good. Let's see if we can put the slides up.

Okay. So as I was coming over here earlier today I was reflecting on what Marcus shingles talked about yesterday. That is that the, if you want to go from linear to exponential in this industry, you have to involve startups, because by definition, let's see if it works. It doesn't work.

So by definition, the startups have in their DNA the notion of leap frogging. In a startup, if you don't go exponential, you will lose against legacy and against the competition. You have to have in this type of forum startups that are very happy that they are represented.

Go back to the slide. Hopefully we will make this work this time.

I have even less luck than you do. What is that?

>> We don't have your slides.

>> Okay, we are back to ...

>> ANTOINE BLONDEAU: Yes, Thomas, why don't you go ahead and I'll make sure that my presentation is up. Thank you.

>> THOMAS WIEGAND: I hope my slide is going to come now. Otherwise we will have two crisp presentations. Okay. How are we doing? Can you show my slides? AI is easy. AV is difficult. As we have now learned. Maybe we should have done the shut eye now.

>> Can we bring the robot back out for entertainment?

>> THOMAS WIEGAND: Do we have questions for Fei-Fei? That may be something we can use our time for regarding her presentation. Anybody has a question? Yes, please.

>> AUDIENCE: Yes. Good afternoon. I'm Zoltan from the School of Tropical Medicine. It is a question about what potentially is the bill from some of the excellent work including on the census that in very different environments. So in countries that don't save and have license plates. So it is a fascinating question to think about how you could create a mechanism to enable the free access of panel data and to help build countries resources around better data. So there is potentially a question about the power of this technology. But just a reflection about how would you translocate either the census work or your work in healthcare where it is community-based and there aren't any cars or license plates.

>> AI is easy, AV is hard.

(Laughter.)

>> FEI-FEI LI: Technology! Yeah, so that was an excellent question. So I think the moral of the visual census work is not about how special cars are. I mean, they are kind of special, but it is about how powerful visual data can be in understanding our society.

So one very exciting recent trend we see is, there is starting to be a lot more work using satellite imagery especially focusing on areas of the work that there is not much road to access and there is too much social unrest for social workers to go in. So my colleagues Stephen last year published this science paper on using satellite imagery to predict poverty level in, I believe, African regions.

Again, we don't have to focus on cars per se. Satellite imagery would be a major source of data.

>> THOMAS WIEGAND: Okay. So is it working now with the slides? Because we kind of have to move on. I understand -- oh, here we go. So now it starts working. So let me do my presentation now.

So I have a few remarks I put together and I have to say that they changed over the course of this conference.

One thing I thought was missing somewhat is the digital twin, which is something that we have everywhere. And by training I am a communications engineer. So I have been active in this field of communication down here. But if you want to create additional trend of the real world and the real world contains ever everything that interests us, humans, health, cars, energy, materials, you name it, factories, we have to have sensors and produce data. We have to have computation. We have to have communication. All these things need to be seen together.

As I mentioned I'm an engineer. I want to optimize things. So we found that we have to work on them together.

So the digital twin is this concept that actually turns real process into digital ones. It has this effect that we are all aware. For instance, the digital media industry has changed fundamentally. The reason was distribution and copying became effortless. It has been costly before, no effort now. Same with digital stores and product display. In Germany we have the term for the industry 4.0. We are making a digital model of the factory and simulating certain things that happen in the factory so we can optimize the process, maintenance, et cetera, et cetera.

And so the principle is once a costly process of society, business or science is turned into digital, the resulting benefits from just this part of it already can have a profound effect on everything else. And change the model and the existing players are rolled out and new players are rolled in.

So there are obstacles for this to happen. It hasn't happened in some areas. For instance, I will say medicine, for example. And it could be that there is these obstacles, like privacy, security, safety, to mention a few. But there are more. It has to be clear how this plays together. And AI, I just see this as a part of the whole digital twin aspect. Of course, those of you who make robots also have to add the robots into this picture as well.

So this is the first thought. Second one, I want to talk about some work that we did on machine learning and communication. So we are working at my institute. I am from the Fraunhofer Heinrich Hertz Institute in Berlin. So the 5G network, you have heard about this. It has a number of big promises. So 1,000 times through-put. One hundred times the devices that includes sensors. Ten times battery life and one millisecond latency. Hopefully not all of those requirements should be achieved at the same time. Still, each of them is a big challenge. Think of the energy consumption. The energy

consumption per bits stays the same if we increase the throughput by 1,000 times. Do we need 1,000 times the energy? If it's at 4 percent, that would be meaning 4,000 percent of the energy for just data transmission. Something has to be done in a revolutionary way here.

So what this is supposed to enable is mobile has Internet, car to car and car to infrastructure communications and industrial wireless. All these things are something that we are working on.

And we have been most recently, when working on the machine learning part of it. So machine learning is widely used in designing these systems. So we use machine learning as an engineering tool. We are solving problems with machine learning that we have found to be difficult in the past. There is some areas where it has been known to be used and there are other areas.

Communication engineering design and machine learning are something that really go together. And these two fields are already converging: Communication and machine learning are converging.

To give you an example, a radio map is a function that relates to a geographic location, to a radio system parameter like path loss. So when you are covering an area with your, with this mobile communication, the reception is different. If you still have FM radio you are standing at a traffic light. It starts to, reception starts to get bad. You move a little forward, reception is better. You can estimate these forms.

These path losses, this is a map. You can then make a 2D view of it. So we want to have an online reconstruction and prediction of these radio maps from user measurements. And we've done this in Berlin. This is the original path loss map. This is our result. And the people who asked us to do this found this useful.

Another area where we have been using machine learning is video coding standards. The ITU has been very successfully launching video coding standards. H264 is every bit on the Internet and H265 is becoming relevant. H265 and H266 is something we are planning now. Only the decoder is specified and realtime encoding is less so. What do we do?

By the way this is the games for the CODEX. We measured video quality using EEG, some work we started years ago. An's electrode at the back of your head can tell you how strongly you perceive the distortion introduced. The coding algorithm becomes so complicated we created a learning algorithm that creates an algorithm to do the video encoding.

We did these two things and so we took the standard. We made the subjective quality and realtime encoding and put it

into the product and we are winning every competition so far with this product. And our first broadcast was sky champions league in Germany. My home club, Berlin, actually played one to one, which I was very happy about.

So the last thing I want to talk to you about is interpretable machine learning. We have heard that -- well, let me retract a little bit. As a communications engineer, I know when you deal with Mother Nature, systems fail. So the coverage is only to a certain percentage. So you have to, when you have machine learning and AI systems deal with Mother Nature, they will fail and you will have to know what they are doing in order to design them and certify them. So the question is we are doing big data, Fei-Fei's Image Net. And machine learning and come up with annotation, potentially very important one. The question is then do we trust the machine? Like this one? Well, what we have done is we worked on a method to interpret machine learning algorithms. On the left-hand side you see the classical model prediction. On the right-hand side you see the data, the model, the interpretation. You can have a human inspection whether it is actually finds what you want it to find and you have verified predictions.

So what we will do is take a neural network and revert it. We would map the classification value back through the algorithm on to the input signal. It's a picture here that gives us a heat map that will tell us what in the input picture led to the classification results. For a picture that would be pixels, but for a self driving car it would be this traffic light let the car to drive to the left or whatever and we have successfully done it. Here in the picture the shark is detect understand and will give you the triangle fin or this method gives you positive and negative evidence.

So we are currently applying this to real systems where those decisions matter. And again it also is something that is interesting to see how it is learning to, for instance, optimize this Hitari application here.

So to conclude, the digital transformation is, I think we need to see it as a bigger effect and need to also ask ourselves, okay, computers are scaling in data centers. Communication, hmm, it's okay. Scaling, but it's very difficult. The demand for more data is ever increasing. The spectrum is limited. And that's a big, big problem. And sensors, sensors are just not really scaling exponentially as everything else. And artificial intelligence without inputs and without communication may need to see the psychiatrist or the psychology doctor, right? So I think it's a big issue. We have to worry about the trade-offs, digital transformation versus privacy, security, safety. Machine learning and communication

are converging and lots of people are actually actively in those fields.

By the way, our machine learning work is influenced by our communications work. So we are worried about nonstationary machine learning. We worry about the spec projection because we want to use it for communication.

The video compression, that's the work I'm doing and interpretable machine learning is important on approaches like that. We need to know what the algorithms are doing and why are they finding certain evidence.

And my thanks. And if you have questions, that's my EEG photo.

(Applause.)

>> ANTOINE BLONDEAU: Okay. Hopefully everything is going to be working. And let's see leap frogging. Are we good?

Three, two, one. Yes?

>> Click towards that.

>> ANTOINE BLONDEAU: My name is Antoine Blondeau, Chairman of Sensient, based in San Francisco, California.

I was talking about the importance of startups in this process of defining what sustainable goals can be for us as a community. It is very important that we are involved. I will try to give you a few examples of what type of leap frogging can be done. To move, I think Peter Lee mentioned yesterday that we are in the autism phase of AI. What can be done to make it real for the world.

First of all, I am a firm believer that it is not about the data. If you think about who we are as humans we don't go around and collect data for the sake of collecting data. We don't go around and make predictions for the sake of collecting predictions. No. Everything we do in this world, whether it is collecting data, making predictions is in the service of making decisions. And so AI has to be about this. It has to be about making decisions. Therefore, AI has to be embedded in the workflows and the processes, and it's how you get to self evolving, self-learning system, really augmenting human intelligence. We learn by making decisions, whether they are good or bad. We learn from feedback. That's how we become intelligent. Therefore, it's about decisions.

Let's talk about this a little bit. Seventy years ago roughly soon after the Korean War ended, a colonel by the name of John Boyd at the Air Force was staffed with analyzing the dogfights between the Chinese fighter pilots and American fighter pilots in the Korean War.

He came up with this decision making framework which is still very much in use for those of you who have dealt with the U.S. military, as I have. You will recognize this. What does

this loop framework stand for? Remember, it stands for observe, orient, decide and act. Observe is taking your senses, eyes, ears, to understand the circumstances in which you find yourself. Understand your current world.

Orient is taking the abstractions that you have made from this very first phase and compare that against what you've learned so far. Your knowledge, your training. And then once you've understood your second senses in context of what you learned, then you decide. Decide is obviously selecting what you believe is the optimal strategy. Then you act. When you act, obviously the world has changed. Therefore, you observe again.

And it goes through that loop. Now, what Mr. Boyd found that is fascinating was that the most successful pilots were not the ones that were the best at any of these steps. They are not the ones with the best of eyesight, for example, or not the ones with the best marks in training.

The most successful pilots were the ones who are going through the loop as often and as quickly as possible. In other words, Mr. Boyd was discovering, was telling us that surviving, adapting, learning, is key to success. And so we do it every day. But what is interesting, you can port every single phase or every single step of this process in software. Therefore, you can develop a self learning system that learns from action and decisions and learns from feedback. That is what AI is about fundamentally. True AI.

Now, to do that, you can do it today on a small scale. So targeted advertising, the business of Google, for example, uses this methodology effectively, but on a very narrow scale. The game is now to move from narrow scoped, narrow domains to broad scope, much higher dimensionality, much more complex problems. Problems that we as a society care about. To do that, you have to scale AI.

I was talking about leap frogging. Oops, I'm sorry.

So this is an example of leap frogging. This is taken out of a corporate presentation at Sensient. Ten years ago when we decided to start this company we thought about this notion of how can you scale AI to the hilt in order to achieve the resolution of problems at the much higher level, much higher level of obstructions and much higher level of complexity.

So this system is distributed around the world. 4,000 compute sites, 2 million CPUs, a lot of GPU cards. And if you can distribute AI across this massive infrastructure you can begin to solve real complex problems and make a big difference.

Now, you can make a big difference around many things because if every aspect of our society or every aspect of our economy is based on human intelligence as it is, if we can

augment or replace human intelligence with performing AI, then theoretically every business is an AI business.

So that's the theory. Of course, as I said, it is still in the first inning. It is craftsmanship, not industrialized yet. It is going to take time for us to get to this. We have been successful as a society before. When it comes to trying to take inventions and turn them into industrializable processes. Electricity is a commodity now. It wasn't like this 120 years ago or 150 years ago, but it is now. The computer is a commodity now. It wasn't like this 40 years ago.

What will it take for AI to become this? The biggest issue we have and it has been discussed times and times again is the lack of talent. Half ago which was produced or released last year by Deep Mind, Deep Mind Google, Alpha Go is roughly 20 or so Ph.D.s over two years to architect two novel architectures for neural networks and (indiscernible) in the middle. Forty Ph.D.s man hours of hard work. I can tell you it is not easy to manage 40 Ph.D.s for two years. It's headaches.

How can you solve this and industrialize this? You can do it by doing something like this. Which is effectively creating a method of heuristics approach in this case using evolution. So Fei-Fei mentioned before evolution is effective, right? It gets us to intelligence. Can we use evolution or AI effectively, a form of AI in this case evolutionary computation to evolve architectures, just not hyper parameters, topology of neural network, can you create intelligence effectively? That's what we do across this large network. If you do this, these 40 Ph.D.s, we don't need that. We do this with one intern, three months and 5,000 GPUs, right? That's the idea of scaling AI. That's how you begin to industrialize it. You see evolving intelligence.

So where can you apply this? We are working with MIT, in this case the Media Lab, to enable their so-called cyber ag project. This is a very cool project where the Media Lab is embedding in a container, they call this a food computer, actually they call it. They embed the whole environment for growing a crop within a food computer. So you have temperature controls, lighting, nutrients, humidity, what have you.

We embed AI in it, black box optimization. We are effectively able to, using something that Yoshua Bengio talked about yesterday, the long-term implications of time series. We are able to progressively evolve a process by which qualitatively and quantitatively gets to better outcomes. So we thought basil, we are producing the best ever made, interesting. AI is producing basil, the best ever made, quantity and quality wise.

Just an example. You obviously use the same types of technology to address wellness. Same thing. Can you optimize this black box called human wellness? You can. Just like you can optimize agriculture, same thing.

Something that I really care about, also with MIT. This one is C cell, a project we did together. Can we save your life in an ICU? We talked about before, someone talked before about the issue of mortality in the U.S. Was it you? The biggest cause of mortality is blood infection, typically caught on site. You die roughly one out of every two times. Better not contract it. The biggest single symptom of a blood infection is massive drop in blood pressure. Today there is no way to predict it. You can just react to it. When you react to it, you have one in two chances to not make it. So can we predict it? We can.

That was interesting. We actually run this across 6,000 patients. We are able to predict the on set of sepsis 30 minutes ahead of time, which is enough time to preempt and prevent it, 91 percent of the time.

Now, the call to action here for this community is, how do we get from this in the lab to this in the ICU? Because today if I want to take this, which I know is going to save lives in the U.S., I'm going to have to fight for years with the FDA to get it. There is not a good reason for this, if we are talking about exponential versus linear, I would like to have it when I get to the ICU, God forbid. That's a call to action. We have to understand how we can make this happen fast, not slow.

Oops. Is it going to work this time? All right. And this is the last thing I'm going to talk about. This is related to privacy and I think a subject which is near and dear to many folks in Europe particularly, where today for lack of a better word Facebook, Google, Alibaba, et cetera, are watching you. Everybody is watching what we are doing every day. There is no barrier against this. Anybody recognize what is on the right-hand side? What is this? Anybody watches science fiction movies?

>> Great System of Her.

>> ANTOINE BLONDEAU: Exactly, in the movie Her, it depicts a near future era where these devices instead of, they are called smartphones but they are not very smart. The smarts is here. This is an access tool. So her depicts a future where this device or any device, the network is real smart. As smart as your personal assistant to the power of N. And when that happens, and if the person is independent, but not quite independent, but your proxy supposed to your proxy as opposed to Google or Facebook's proxy. You can ask your proxy to do your shopping, your services, whatever you have, without having to expose yourself to the world of brands, to the world of

commercialism out there. You can shield yourself from that world. You can have that proxy under your control.

Today there is no business model for this and no startup as yet has come up with this system. But there is an opportunity here for this to happen. That's why the big guys are very focused on focusing the hardware in your hands. They know if they don't have access to it, this will happen. Their ability to price their services will go away.

So if we do that well, not only consumers will be in good shape, but also a bunch of industries will be. I think that is the opportunity that we have as a community. If we can industrialize AI, this will happen. And we will all be better for it. Thank you very much.

(Applause.)

>> THOMAS WIEGAND: So we will have a few questions. So I will ask each of you one question. If you want, you can ask each of the others one question back. And then we hand it to the audience.

So please be quick. Fei-Fei, my first question to you is I see great results but they are all 97, 98 percent results, which is great. If it took us 50 years to get to the 97, 98 percent, how much time would it take us to solve those extra two to 3 percent? Because I'm a communications engineer. My experience is that it takes you 50 years to get to those and it takes another 50 years to get the other 2 to 3 percent.

>> FEI-FEI LI: One thing I learned as a scientist, don't make predictions.

(Laughter.)

>> FEI-FEI LI: I don't really know how long is it going to take to close the gap of 2 to 3 percent. And I guess you are specifically asking about Image Net kind of benchmarking, right? So from a basic science point of view I think by the time that you are 2 to 3 percent to perfection, you kind of are moving to the more engineering development phase of that system and question rather than basic science phase. So I think now a lot of effort is put into, is actually happening in the industry in terms of closing that gap. So I don't have a good prediction, but I think in terms of image classification there is a lot of brute force by data and larger and larger models. We are going to see that the progress soon, but the question to me is not the 2 to 3 percent but really what is the next important question that we ask and try to tackle.

>> THOMAS WIEGAND: Antoine, you are learning, using an unsupervised learning method to solve problems. How do you see the energy aspect of all of this? I mean, we humans perform a thinking task with a certain amount of energy. And these

algorithms use energy. They use a lot of energy. How do you see the energetic aspect.

>> ANTOINE BLONDEAU: Eighty watts, our brain consumes 20 percent. If you want to take the processing power and apply it on a silicon basis and deliver the same, you are talking a lot of megawatts, many, many, many, many.

By the way that, if anybody is concerned about the use of the human brain in the near future, that's the salvation. We still have some ways to go because we are very effective. So yes, energy at the end of the day, energy is the final frontier for these machines. Particularly a self-learning machine will want to be able to access more computer time because it knows that's how it satisfies its functions better. So if the underlying agenda in your question is are we in a situation where we are going to waste a lot of energy doing this, it is possible we will waste a lot of energy. So my recommendation is always to compare -- to look at the problem we are trying to solve and see whether it's worth the energy we are going to spend. If it's worth it to spend brain energy to do that, that's better than wasting silicon energy.

>> FEI-FEI LI: I'll ask you a question. Community standards are important from 43G to 4G to 5G. Do you think there needs to be standardisation of AI? If yes, what is it?

>> THOMAS WIEGAND: Well, if you want to make it mainstream in a fair way so that lots of players can participate, then I think standardisation is required. And for example, very pragmatically I see that the use of machine learning in communications is a subject that could well be standardized. The network management resource allocation, user location, those are all things to actually lower the energy per bits that you are spending because the energy costs are the highest costs an operator has these days. It is not the people. It's the energy.

So yes. I believe it may be a good idea to start thinking about a regulation also in the space of autonomous systems. It would be nice if they work in Germany. If you cross the border to Switzerland, it would be nice if they continue to work the same way and you don't have to switch the car. There are lots of topics that need to be discussed between the government and industry in order to get this working.

By the way, I think a United Nations organisation would be a great place to actually do this discussion of standardizing, for instance, and regulating autonomous systems.

>> ANTOINE BLONDEAU: I have a question for you. You talked about women in the workforce and particularly in the AI space. I look at my company. We have the same problem. We have very, very few women. We are trying, we are looking.

What do you think can be done? Is it really the company's issues? Is it something in the upbringing that creates this problem, lack of interest? What is it in your view?

>> FEI-FEI LI: Oh, boy, that's a really big topic. We can talk for days about this. And I don't claim I know it all.

So I am mostly in higher education. When I look at the lack of diversity in AI and stem, I think it's across the board. There is the pipeline issue, so starting from K through 12 that you can already see there is attrition. In the U.S., MIT, the entrance class is 50 percent women. And Stanford too. But by the time you see graduation and how many women take the CS degree, it's less than 30 normally. By the time you reach grad school and faculty, the attrition is really high. So the pipeline is an issue.

There's cultural issues. We talk a lot about from the early age classrooms all the way to industry there is still quite a bit of cultural issue.

There is also in AI For All we particularly focus on a very important issue that I think is under focused, which is what is technology for? And this is coupled with the culture. In Silicon Valley there is a very strong sentiment of the geek with hoodies, you know, that kind of stereotype is that you do technology because it is cool, which is good. I mean, that's a very positive thing.

But people have different reasons to do technology. Especially people from diverse background and all walks of life. The mission for the technology is really important. So in AI For All we focus on the human centric mission of the technology. So to inspire girls, to inspire under represented minorities, because many of them have choices in their career and they want to go for something that is inspiring and it is for some benevolent cause.

So thank you for raising that. It is a big issue that we all should think more about.

>> THOMAS WIEGAND: Any more questions? You're good?

Okay. I think we are good. We can take maybe one or two questions from the audience and we have to close the session. Please.

>> AUDIENCE: Can you guys hear me? Are you guys using any of the techniques learned from the block chain technology such as Central, Crypto Economics, so on and so forth? No?

>> THOMAS WIEGAND: The answer seems to be no.

>> ANTOINE BLONDEAU: The answer is no, but there is potential especially around distributed computing and distributing resources and the securing of intermediate results in the workflow for block chain technology to be applied. Together with AI.

But not today.

>> THOMAS WIEGAND: Any more questions? Yes, please.

>> AUDIENCE: It is not working? Oh, it is. It shows red.

The same question but if I can ask you each to give a comment, what do you think would be the one thing or things that would be the most enabling of things from your perspective to allow poor countries to benefit from your research and your innovation? It is an unfair big question, but is there one thing from your individual disciplinary perspectives that you think would be the most enabling thing to do? Between to enable poor countries? I think it's education.

So if you allow me not to talk about sectors but around the infrastructure, in this case the human infrastructure, ironically AI is first and foremost about human intelligence right now. If we could focus the digital framework in some of these countries, is there on site or giving them the opportunities to have some of their skilled HR go and study abroad and come back, right? I think this would be to help enable an ecosystem, that would be in my view the most useful approach.

>> FEI-FEI LI: I really agree with that. Let me just add another dimension. I think democratization of data. AI despite the fact that AI eventually will go beyond the big data technology, but right now very effective AI technology relies on big data. And I would like to see more democratization of data so that the play ground is. Much more leveled and a lot of people can access that to develop AI technology for global usage.

>> THOMAS WIEGAND: Excellent. I think this is the end of this session. You have one minute before the session of break and now you have another ten minutes of break and the four breakout sessions will happen. I hope those of you assigned to those know where you need to go. And I think we see each other again afterwards at 5:00 o'clock. Is that correct? Yes.

At 5:15 we meet again to have the results of the breakthrough groups. Thank you, bye-bye.

(The session concluded.)

***

This text is being provided in a rough draft format. Communication Access Realtime Translation (CART) is provided in order to facilitate communication accessibility and may not be a totally verbatim record of the proceedings.

***