TechFest 2013 Keynote
Rick Rashid
March 5, 2013

ANNOUNCER: Ladies and gentlemen, please welcome Microsoft Chief Research Officer, Rick Rashid.

(Applause.)

RICK RASHID: Hi there. Well, hopefully this will be a very fun day for all of you. For me, it's interesting. It feels like we've been doing TechFest for a really, long time, but it's actually only been since 2001. So, technically, I think, this is the 13th TechFest event for Microsoft Research.

Of course, in the first TechFest, it wasn't nearly as big an event as this. It was just for employees. We didn't take up nearly as much space. And we didn't even know how many people would really show up. I was actually one of the strongest skeptics about this whole idea of creating this technology show for our employees, because I wasn't really sure how many people would show up. I kept saying, “Well, gosh, it seems like it's going to be a lot of work, and I'm not sure it's going to be really valuable and worthwhile, but if you guys really think you want to do it, okay.”

Of course, what happened that year was we had, I think, about 4,000 of our Puget Sound employees come over a two-day period. And of course Microsoft was a much smaller company back then, so that was just a huge fraction of our Puget Sound workforce at that time. And everybody was really excited. I mean, the researchers came away from the event really excited about being able to talk to so many of our employees, be able to interact with them, and to sort of sell their ideas, and show off what they've been doing. And we got a lot of great feedback from the employees who looked at this and said, “This was inspirational.” They didn't even realize that Microsoft was doing a lot of these things and that we were working in so many different areas.

And so, from that beginning, TechFest was really born, and it has just grown since then. A few years ago we added this public day, where we are able to bring outsiders in to show a select collection of our demonstrations, and to meet some of our researchers. And so, I'm welcoming you here to that today.

I came to Microsoft in 1991 to start Microsoft Research. So, I've been running Microsoft Research from the very beginning. And over that entire period of time, we've had exactly one mission statement. And you see it up here on the screen. It's in three parts, and the order is really important. The most important part of our mission statement is that our goal as a basic research organization is to move forward the state of the art in the areas in which we do research. And the reason that's important is, if we're really not moving the state of the art, if we're not at the forefront of the areas in which we're doing research, then fundamentally I don't think we're that valuable to a company like Microsoft. We need to be constantly pushing the envelope, and have people that are both creating new technologies, and exploring the possibilities that those technologies open up.

Now, when we do a great job of that, then you get to the second part. Remember the first part said nothing about Microsoft. It just says if we're going to engage in research in a particular area, we're going to be the best we possibly can be, we're going to hire the best people that we can in that area, and they're going to push the state of the art.

When we have great ideas that make sense, then that second part comes in, which is we want to get our ideas out. The reason I went into the field of computer science was I got excited about the idea that my intellect, my mind could sort of animate these otherwise inanimate pieces of hardware that I could create things for others that would make their lives better, or help them in some way. And that's what our researchers want to do. They want to impact the world, and the best way to do that is through Microsoft's products, really getting our technology into the products, and into the hands of hundreds of millions, or even billions of people.

Now, all of that translates into the third part there, which is that ultimately our goal is not just to make sure that Microsoft and its products have a future, but really fundamentally make sure that the field of computer science, or the technology field, has a future. As a field, we need to keep moving the goal posts. We need to keep moving things forward. If we want to be able to continue to add value and continue to be important in the future of society, we need to constantly be improving the things that we're doing, and making sure that they actually have a future.

Now, you get a sense here of what's happened in the last 21-1/2 years. We've built a number of research labs around the world, with some of the most amazing researchers in the world. You see just some of the awards, and some of the kinds of people that we have.

To give you a sense of where Microsoft Research is today, we are the number one basic research organization in the field of computer science. We publish more research in the field of computer science than any other organization. We have roughly 850 Ph.D. researchers. That means we have more Ph.D.s working at Microsoft Research than many universities have Ph.D.s. So we have more Ph.D.s, for example, than Brown University. We have more members of the National Academy of Engineering, than, for example, the entire University of Washington. It gives you a sense of the scale at which we operate, and the important role that we play in moving forward the state of the art for the entire field of computer science.

You can see the breadth of our organization in terms of its spread around the world. And, again, people will sometimes ask me, “Why do you have all these different labs in different places?” and the answer is we're trying to make sure that we have the best talent we possibly can. If we only had research labs, for example, in the United States, we would probably be about half the size that we are today. In fact, about half of Microsoft Research is in the United States; about half of it is outside the United States. There's a limit to how many great researchers are created each year in the world, and we need to make sure that we're in the places around the world where we can capture those great people, and give them environments where they want to live, where they want to work, and they want to create the things that they create.

Now, if you noticed, there is actually nothing below the equator there. But we announced earlier this year, or actually late last year, that we will be setting up an Advanced Technology Lab in Rio in Brazil coming up later this year. That will be our first Southern Hemisphere lab. I was joking with Rico Malvar, who has been involved in helping to put this together, that we could have called it Microsoft Research Southern Hemisphere, but that kind of sounds a little too grandiose. We do have Microsoft Research Asia, which somehow doesn't include India. So, I guess we've been inconsistent in the past.

But, again, the key point here is that we have grown to become a significant player in basic research all around the world, and in every part of the world that we operate we work closely with the academic community in those environments. We work closely with universities, with other research labs in corporations, and especially with students.

Microsoft Research runs the largest Ph.D. internship program in the field of technology. We have over 1,000 Ph.D. interns working at some part of Microsoft Research every year. To put that in perspective, that probably means we have more graduate student hours working within Microsoft Research than any other academic institution in the world.

Now, you can ask the question, why is it that we do this? So, Microsoft made an important decision back in 1990 that led to the creation of Microsoft Research in 1991. They made a decision to invest in long-term basic research. Microsoft was a pretty darned small company back then with just about a billion dollars in sales. Now we've got a run rate of something closer to $80 billion. So it gives you a sense of how much smaller we were, and yet the company made a decision then to invest for the long-term, to invest in basic research.

Why did we do that? Well, there are lots of reasons why people talk about doing basic research. I've got some of them listed up there. And they all apply in a way. They all make sense and they're all values that a basic research organization creates. You sometimes say, hey, we're a source of intellectual property, you know, new product technologies. And we do a ton of that. You'll see that in some of the demos.

We generate about a quarter of Microsoft's patents; they come out of Microsoft Research. And those are the patents that are in some sense the most valuable in that they really are the scientific patents, the ones that push the bar.

The IEEE over a number of years looked at Microsoft's patent portfolio and rated it one of the top scientific portfolios in the field of technology, and really that's on the strength of a lot of the research that we do. So, that's really valuable. But, that's not in some sense the reason to do a basic research lab. In some sense the best way of putting it is probably to say it's a value that comes out of it; it's an outcome. If you have a great research lab you'll get that.

Problem solving: We have really smart people that know their areas really well, and they love solving problems. That's what they're trained to do. You bring me a problem, I just love solving it. I just love a hard problem to work on. And a lot of our researchers are that way, as well. And so our product groups do bring us problems. And that is an important value that a basic research organization brings to the table. We do help solve really tough issues that come up. But, again, I would view that as an outcome, not necessarily a reason.

Lastly the early warning system: a lot of people always talk about that when they talk about research. They often say, "Well, you guys are an early warning system. You can tell people about the future, because you're working off some number of years in advance of where the current products are." And that's true, too. We do do that. But, again, I would view that as really a result of the fact that we have a great research lab, not a reason to do it. I don't think that's a value -- any one of these things taken alone I don't think are valuable enough to say, "Hey, that's the reason why we should be paying for this very long-term enterprise."

The real reason I think that you do this (and I always say this whenever I talk about research, and this is true not just of Microsoft Research, this is true of research in general— the investments that the U.S. government makes in basic research, or the European Union makes in basic research, I think they have the same rationale, which is that investments in basic research are fundamentally about survival), it's about making sure that if there is a new technology, a new competitor, a change in the business climate, that you'll have that treasure chest of ideas, intellectual property, and people that can help you change, that can help you turn things around, that can help you engage, move forward in a new way, or to solve an important problem.

If you read the writings of Vannevar Bush, who helped to create the National Science Foundation in the United States, you'll see that same kind of rationale that he made for why the United States should invest in basic research. The U.S. had just come from a terrible war in which technology was a huge part of winning that war. And he made the statement that we need to invest and to create this sort of treasure chest of ideas and people, so that if there's a new war, or if there's a famine, or if there's a disease, that we as a society will be able to survive. So, I think that's really the key point here, that's why it's important to make these investments.

Many times the things that we do within Microsoft Research, at the time we do them, there isn't a really obvious value. I'll tell you a simple story from a long time ago now. It's one of the most important early stories in the history of Microsoft Research. Back in 1992, so one year after we'd gotten started, I think at that point we barely had two dozen employees, we developed—one of our researchers working with some others—developed a technology for dramatically reducing the working set size of 32-bit programs. So, the right way, if you're not really a programming language or operating systems person, what that really means is that we came up with a way of taking the amount of space taken up by the code of the program, that needs to be in-memory to execute, and reducing that by about a factor of two. Now, this is mostly a lot of math. I mean it was primarily just a lot of symbols on a blackboard. But, it was really cool. And we were really excited about this. This was a real breakthrough.

So, we went to our product groups—and I remember I was sitting in a lot of those meetings—and we had these great meetings. We explained what we were doing, we showed how it could impact the execution of programs, how we'd save a lot of memory space. And back then memory space was very precious. And there were great meetings, lots of good interchange, we met with a number of different product groups, but at the end of the meeting each group would say to us, you guys are just really smart. Okay. Look at all those symbols on the board. You guys are just smart.

But, we don't actually have this problem, right, because it's 1992 and we only have 16-bit code. We don't have 32-bit code. Now, we will have 32-bit codes, it's coming down the line. We know it's happening. But, by then the memory sizes will be a lot bigger, because they had all these curves showing how memory prices were going down by a factor of two every year and memory sizes were going up. So, we're really not going to have this problem, but you guys are really smart, and so, go back and work really hard. I'm sure some day you'll come up with something we need. I don't know when that's going to happen.

Now, those were the polite people. There was one group we met with and they actually had 32-bit code. This is the Windows NT group. And we had a great meeting with them. They had really great, smart guys, too. And at the end of that meeting the head of that group came up to me, he's a real big guy, and he looked me in the eye and said, “You're going to do this to our code over my dead body.” So, 18 months later he sent me a really nice note saying, “I really want to thank you guys for all the hard work. We couldn't have shipped without you.” And what had happened, and this was all lead-up to an important change at Microsoft. So leading up to 1995, which was the year that we switched to 32-bit Windows and 32-bit Office, with Windows 95 and Office 95, what happened was that in that period, from about '92 to '95, if you look at RAM prices they didn't actually change. They actually went up a little bit and kind of tailed down a little bit, but mostly they stayed flat. So, what the company was looking at, as it was about to ship its 32-bit products in 1995, was that the memory sizes of PCs back then was roughly four or eight megabytes. Just to put that in perspective now, most children's toys have more memory than that. In fact, all children's toys have more memory than that now.

But, back then that was what an Office PC looked like. Unfortunately, that was the amount of memory that Windows 3.1 worked great in, our 16-bit product. It wasn't adequate for our 32-bit products. But, we had been working on that technology going back to 1992 and we had it. Our product groups were then able to use it. It became a standard part of their build environments, and we were able to ship Windows '95 and Office '95 on schedule, and they were able to run properly on machines of that size.

That was important for a variety of reasons. That wasn't just important because it made us money. But, even more important, for a small company like Microsoft, this was a critical point in the history of the industry. We were going from 16- to 32-bit code. It was a big change. And our competitors didn't have this technology. WordPerfect and Lotus—they really didn't have it. And it made a huge difference in the marketplace.

If you talk to Ray Ozzie, of course he was at Lotus at the time, he'll tell you they knew that we had it. They knew we had this technology and it was killing them. They didn't know how we did it. And that's just one tiny example, one algorithm, from one, or two, or three people, that was a turning point for us as a company and a particular turning point in the industry. That's what a basic research group does. Those are the moments that pay off in some sense. We always have those constant payoffs; new technology is coming out all the time, all those other things I listed, those are the everyday payoffs. It's the big transitions that are incredibly important.

Now, I talked about our mission statement, moving toward the state of the art. I pointed out that we published more basic research than anyone else in the field of computer science. The technology transfer part is also very, important to us. This is just a list of technologies that have gone into our products from Microsoft Research since last fall. When I look back at all the research labs I've been in, doing this sort of thing for 35 years now I guess, when I look back at the research labs in the field of technology, I've never seen an organization quite like ours in the sense that both we're moving that state of the art forward, in a way that really almost no one had ever done before, certainly at the scale we're doing it. No one has ever done it before.

But, at the same time we have this intimate relationship in terms of getting our ideas into our products very rapidly, and having that kind of big impact. And again, it goes all over the map. Some of these things are individual pieces of technology that go into our products. In some cases these are whole products. If you look at something like Kinect, that's a whole new product that was created working with our teams in IEB, based around Microsoft Research technology.

So, with that I'm going to bring out one of our principal researchers, Curtis Wong, to talk about something hot off the presses. This is the hot off the griddle piece of technology transfer. Curtis has done a lot of things that have had a big impact in things like Worldwide Telescope that he'll mention to you. And most recently he's been working with our Office team to bring a new product, a new technology, GeoFlow, into our Office products. And we're just telling you this now. So, this is news, hot off the wire.

Curtis.

CURTIS WONG: Thanks, Rick. (Applause.)

It's so great to be here. Ten years ago Jim Gray, who was a Turing Award winner, was tackling the challenges of big data, to help astronomers get insight into the terabytes of data from the Sloan Digital Sky Survey. And he did a talk that I went to, where his last slide was “We need visualization and we know it.” So I got involved and started thinking, “Wouldn't it be great if we were able to build the biggest telescope on the Internet with the highest resolution imagery from ground and space-based telescopes and connect it into information all over the web, and even better, build a virtual camera in this environment that would allow astronomers and educators and even kids to create guided tours of the universe?”

And so we did that five years ago, almost to the month, and we launched it at TED. And so what I'm going to do here is… I'm actually going to show you a little bit. This is the Sloan Digital Sky Survey. This visualization is a million galaxies. In fact, it's a million galaxies from the Sloan Digital Sky Survey, and the data that Jim helped make possible is what's powering this. Now, even though they look like tiny dots, each one of those dots is a galaxy with about 200 to 300 billion stars. But, it's not just a visualization because you're seeing the large-scale structure of the universe. And if I pause and clicked into any one of those dots, I'd be able to connect into data, information, and imagery that's in a completely different cloud somewhere else and get full access to information about that one particular galaxy. Jim Gray helped make that possible.

What's interesting about this is, it's not a video, but it's essentially a path in this 3D environment. And there's hundreds of terabytes of data behind that particular thing. Now, we're going here into the Milky Way, and behind that you'll see about 100,000 stars, which are all the stars in our nearby portion of the galaxy. These are all rendered accurately in terms of space and location. With this kind of visualization you can see the structures that are impossible to see when you're looking at just raw data. With the Worldwide Telescope tour here, as you can see, all these half-a-million asteroids, kind of useful after that thing in Russia, but with the tour you can pause at any time and start drilling into some of this information.

So, we originally built Worldwide Telescope as a research project. It's been amazing, to see what's happened, because now we have about 10 million users, well 10 million kids of all ages, really, exploring the universe from every continent on Earth, including Antarctica. I hear from teachers all the time that send me e-mails, or find me somehow, and tell me about how they use Worldwide Telescope both as a learning environment, and also as a carrot to get kids excited. If they do well in class they get to drive Worldwide Telescope and show the class something that they've discovered and want to be able to explore. It's also used by astronomers. At the Harvard Center for Astrophysics it was recently used to visualize something called Nessie, which is a large-scale structure in the Milky Way, it kind of looks like the Loch Ness monster. That was visualized in Worldwide Telescope and announced at the last American Astronomical Society meeting.

We built Worldwide Telescope back then and to the point that Rick was talking about, I had this crazy idea of building this thing and I wanted to make it as a dot org, a free resource for science and education, and it clearly didn't have any particular connection to business at that time. But, it had a potential huge impact to inspire and educate people around the world. But, I knew that the technology that we were going to develop from this would some day be applicable to the things we do here at Microsoft, especially if you're looking at what's happened since that time, in terms of the growing amount of spatial and temporal data with mobile devices and things like that. How can we start to see things and visualize that kind of information? So, for the past few years I've been meeting with lots of customers to understand their data challenges and how we can apply some of these ideas that we developed from Worldwide Telescope, and help them gain insight into spatial and temporal data. So, after all, if we can visualize a million galaxies in the large-scale structure of the universe, we can probably do something for business intelligence.

So, I've worked closely with the Startup Business Group and Microsoft Office over the past year and a half, to bring together a solution we think you'll find interesting. So, I'm super-excited to be able to show you this preview of GeoFlow today.

This is a table of data, which is publicly available data from the City of Chicago. They have a web portal that anybody can get access to. And so I downloaded this data. As you can see it has information about some narcotics arrests from 2011 to the present. And if I sort of zoom down here you can sort of see that we have about 70,000 rows of data.

Let me do this one thing here and show you, this is a new add-in that is going to be available later. So, you can see I'll put GeoFlow here. And now you see this plug-in up here at the top. So, what I'll do is I'll launch GeoFlow and I'm going to create a new visualization here. Notice right here we have the globe and the first thing I'm going to do is provide some geographic coordinates. So, I'll take latitude and longitude. And these are some touch controls in case you wanted to pinch and zoom, et cetera. But, this is data from the City of Chicago.

So, let's go and take a look at the City of Chicago here. And we have this tradition that we used to do when we were building GeoFlow. Whenever we'd show new features in our feature review meetings, John Paine, who is one of our designers, would go up and say, "Boom," and then we'd all cheer when we see something new and cool, because we hadn't seen it before. And so it would sort of convey the level of excitement about that particular feature.

So, I want to kind of get all you guys going, involved with this thing with me, too. What I'm going to do is, I'm going to press "Map It." And so we have those 70,000 data points and BOOM. (Applause.)

Thank you. All right. And we're back in it. So, now those look like a lot of points, but they're all kind of blended together. There are quite a lot of them. And if I show them smaller, you can see just how many there really are. I'm going to go in and put a name in here. These are 911 calls. And let's start to look at this data. This is data from three years. If I take year and drag down into category, what it does is it shows three years of data. (Applause.) Seventy thousand data points.

Now, these are statistograms. Now what we can do is we can show them clustered as well. So you can see over here that 2011 was bigger, and it sort of got smaller in 2012. Of course, 2013 is only a partial year. But that's really not necessarily what we want to look at. Maybe what we want to do is instead of a column chart we want to sort of see it as a heat map. BOOM. (Applause.)

So, let's go back to column. What I really want to do is look at this by drug type. So, let's go drag this down into the category field. BOOM. (Applause.) Yes, okay. Now we're looking at a lot of purple. So, purple represents possession of cannabis, 30 grams or less. So, it looks like the police are spending a lot of time busting people with pot. (Laughter.)

Some of the other things we can sort of see here is light blue. This is heroin, white heroin. And that's sort of the realm of heroin. And then on the other side here, we see a lot of yellow. That's cocaine. But we also notice here that there are some of these tall bars here, which are real outliers. One thing I want to add here is that we have lots of different themes from which to represent the data. We have 21. So, if I look at this, I can see locations of what I'm looking at. We have other choices, say we picked a darker. Notice it changed both the data represented to emphasize contrasts so you can see that. We have a number of different choices of how you might want to represent that data.

But one thing you might want to be able to do is select and examine what are these outliers. So, I can hover over this thing, and I can see that there are 238 cannabis arrests of more than 30 grams. Clearly, this is the big box store for pot. (Laughter.) And what I can do is I can add an annotation in here. (Laughter.) Kind of call that out. Anyway, so that's embedded in there.

But the other thing you might want to really be able to do is provide a way for people to see the things that you're talking about in a way that's different from what you do with a static chart or representation. When you have a chart and graph, a lot of times what you want to do, that's just the beginning of the question. What is that? And what's going on there?

Here in this 3D environment, what you can do is you can create, as I showed you with Worldwide Telescope, sort of a guided tour into this space. And we can do that very simply. Say we wanted to start with an establishing shot here of this. And we'll just add a scene, and that's sort of the first scene that we have there. You can do things like adjust the time for this particular scene. Let's say we make it 3 seconds and we want to make it fast. There are a number of camera moves that you can do, either things like fade in or fly to, or different kinds of camera effects as you get to that data. We'll just say stationary.

And for this particular location we might want to just zoom in on this one location right here. And we'll do the same, we'll change that to three, and make that fast. And then we'll go to another place, and maybe look at something out here, another green column.

This is by O'Hare, in case you want another distribution point. And let's see, let's make this 3 as well. Okay, we're ready for another boom moment. Okay, let's do it. BOOM. (Applause.) The loudest guys are going to get a free swag of GeoFlow Post-It Notes. Think about it. Okay, so notice we're just moving the camera through, very simple.

So, when you start thinking about data and big data, especially in the cloud, how do you do analytics and representation of data in the cloud? Traditional ways of representing data, we would use spreadsheets, and you can hand the data to somebody. When the data gets too big to move, you have to move the analytics to the cloud. How do you do that? This idea of a guided tour is like being able to drive a virtual camera in the data in the cloud and send out an HD movie, maybe it's reduced data, but still it's a very large data set. And it's a window into the cloud. And that path is not a frozen thing. It's just basically metadata, and that metadata path can be shared with other people. Other people can annotate it, show something else, re-annotate it. It's a new kind of idea about how we might be able to visualize and represent data.

Let's go back here and take a look at this same data. The other key thing about this is time. We have temporal data with this, and it's really kind of a difficult thing to show. But in GeoFlow, we have the ability to look at things over time. So, let's go and select by day, and we'll accumulate this data, and we want to be able to sort of loop through it. I'm going to shrink this down. And let's play it. BOOM. (Applause.) It's somebody back there. Come find me. So you notice the data is moving through time as we're zooming through this; very simple.

I think I'm running out of time, because the hook is right here, but what I would like to just say is that I think GeoFlow is a really interesting example of how technology created in Microsoft Research that's purely for science and education and serving the greater good, we can leverage that same technology to develop innovative technologies for products.

And if you're interested in trying out GeoFlow, talk to your Microsoft sales representative about getting access to the beta, because space is really limited. So, thank you. Boom. (Applause.)

RICK RASHID: Thanks, Curtis.

Now, of course, one of the themes that you're going to be seeing as you go around TechFest relates to the thing that Curtis was talking about, which is this notion of big data. I mean, we're really in a world now where we can both store and access, but also collect data on truly a global scale. You can see we can even represent it now and display it and use it in presentations on a global scale. I think that's an important change in the way we interact with the world.

And I think another part of that is the fact that not only can we get all this data, but we can learn from it. We can have principled analysis that is actually based on data. In the old days people would talk about expert systems, but expert systems until fairly recent times, were really systems that captured human insight and human data, which is great. I mean, human beings would code in a set of rules, or a set of perspectives on things, but it wasn't necessarily principled, and it was often very fragile. Now we can collect large amounts of real data and take advantage of that.

At the same time, and even related to that, is another theme that you're going to be seeing as you go around TechFest, which is this notion of natural interfaces. So, at the same time that we've been building global scale data centers and filling them up with incredible amounts of information, and processing capability, we've also been giving the computers we carry with us and the computers that are around us, the ability to have the senses that we have; the ability to hear and see, know where they are, feel, even understand motion and shape. As we've done that, we've opened up an incredible new array of opportunities, and we've tied that into machine learning.

So, when you look at something like Kinect, when it came out, it was a revolution in real-time 3D computer vision fundamentally based on machine learning. It was so revolutionary that it spawned an entirely new wave of research in the university environment, because up until that point no one really had access to anything like that. They didn't have the underlying technology at that price point and that scale. But also we demonstrated how by using machine learning we could make that kind of real-time 3D computer vision incredibly reliable, reliable enough that it could be put in homes with children and dogs, and grandparents, and everybody else that's there, and still be able to work and deliver a huge amount of satisfaction and value and entertainment to people.

To give you a sense of how the world's changing, I'll give you a really, simple example that is funny; it's remarkably powerful. If you walk into Microsoft Research in Building 99, which is on the other side of 520 here on Campus, that's where our main research building is. If you walk toward the elevator, as you reach the elevator, the door will open for you. What's happening is we have a real-time computer vision system that has been trained to look at the movements of people around the elevators. And has learned over a period of time—and there was a long period of training—what it looks like when people want to take the elevator versus what it looks like when people are just standing around the elevator. So that as you walk up to the elevator, it knows that you're intending to use it. It's already called the elevator, and by the time you get there it's there.

It's so powerful that when the system was down a couple of weeks ago, there were just people standing in front of the elevator. (Laughter.) They had forgotten about pressing the button. By giving our computers senses, they can begin to anticipate what we really need, and by using advanced machine learning we can understand well enough people's intentions that it's reliable enough that they begin to depend on it. And I think that gives you a sense.

Now, another aspect of these natural interfaces is that increasingly our offices and our conference rooms are all outfitted with touch displays that we can interact with and communicate with. It's changing the way we at Microsoft are doing meetings. Certainly within the research environment we have a number of these now. But it also is opening up a whole new set of possibilities about how can we change the way we interact with these displays using natural interface techniques in order to be able to do better presentations and have better meetings and environments.

I'm going to bring out researcher Bongshin Lee, who will show you something called SketchInsight, where she and her colleagues have been working on exactly that.

(Applause.)

BONGSHIN LEE: As Rick just said, we are creating a completely new way for people to not only explore, but also present their data, leveraging more natural interfaces, specifically pen and touch.

SketchInsight is a data enabled canvas that allows you to create interactive data charts through simple interactive sketches. So, it extends the advantages of the traditional whiteboard by bridging the gap between machine and whiteboard.

Now I'm going to tell you a short story about global energy consumption to demonstrate how SketchInsight helps you create more engaging presentations. Let's first see how much energy we've consumed over the years, and say this factor here represents energy consumption, and see how it has changed over the years. The unit here is a billion kilowatts. As you can see here, the annual energy consumption has doubled in 25 years.

Why? One obvious reason is population, because more people consume more energy. So this time let's say this stick figure represents population, and find out the population change over the years. And, yes, the population has increased, but not as aggressively as energy consumption. The increase is from about 4.5 billion to 6.5 billion, which is less than 50 percent. So, let's consider some other reason like per person individual consumption. In this chart, this light bulb represents per person consumption and, yes, of course, it has increased.

And so far we've been looking at things from a global point of view, but there must be some difference between different parts of the world. So, let's look at the consumption one more time, this time across different parts of the world. And as you can see here, North America and Asia Pacific are the top two consumers. So, let's focus on them. And as we all know, Asia has a lot more population than North America. In fact, Asia's population is more than eight times higher than North America's. But if you look at the per person consumption, on average a North American person consumes about 7.5 times more energy than an Asian average person. But if you look at more recent years, then you notice that Asia started to consume more energy than North America now.

So, this is the end of my story, and I'm hoping that you are thinking of two things. First, we are consuming too much energy. So, we should probably cut back. And, second, more importantly, “Wow, that's so cool. How did she do it?” (Applause.) For an explanation, please come visit our booth later this morning.

But for now, I'm going to give you a sneak preview of how SketchInsight helps us find a story to tell by briefly switching to exploration mode. So, in this mode, we can load any tabular data such as a Microsoft Excel spreadsheet. And here you can see that I have several different data sets available. And using a small number of simple sketch gestures, with the system recognizing my text and matching it against the available data columns, I can perform a wide range of data manipulations to visualize and gain insight, saving any interesting findings for later presentation.

So, please come to our booth to hear more about SketchInsight, and try it out on a Surface Pro yourself. Thank you.

(Applause.)

RICK RASHID: That was great. Thanks. This is really cool. I was telling Bongshin yesterday that I really want to take this with me on a trip and do this for one of my presentations.

But, again, you can see the opportunity that now having the sort of ubiquitous displays, and being able to build upon much more natural forms of interaction really create both for telling stories, but also for exploring data and being able to analyze it.

Now, I'm going to tell one last little story. Some of you may have seen some of this already because it got a lot of play on the Internet back in November. Back in late October I was in China, and for the last two years I've been working with our speech recognition and machine translation teams with a single purpose goal in my mind: I wanted to be able to get up on stage in China and give a speech and have it translated into Chinese while I was talking.

It took a while, and in particular it took the development of a new technology, which is called deep learning, or deep neural networks. For those of you that are familiar with speech recognition, for a very long time now the recognition rates for general speech in general settings, so unrestricted speech in general settings, has been very flat, hovering around 22 to 25 percent errors, or one out of about every four words is wrong.

And that just hadn't budged for more than 10 years. And about two years ago, 18 months ago, some of our researchers, working with researchers in the University of Toronto, developed a new way of thinking about how to do machine learning that could be applied to this extremely difficult learning task—the speech recognition task. And this is what's called deep learning, or deep neural networks. And we were the first ones to publish a significant result in this new area and in particular a result in this new area. And in particular a result around speech recognition, where they were able to show a nearly 30 percent improvement in speech recognition error rates over previous techniques.

Now, 30 percent in a field where things had been like this for a long time suddenly seeing the chart go like that, that's huge. But, in particular it got the system to a point where the error rates were low enough that after experimenting with it myself I felt like I could get on stage, I could actually try this out, and I will say the team that was helping me were way more nervous about me being up on stage to do this than I was. I played around with the technology. I was feeling pretty confident about it. They were very nervous, because of course their jobs kind of depended on it. (Laughter.)

Although, one of the guys that was helping to run the demo, when we did the rehearsal the night before, just to test it out, because we're in this conference center in Tianjin, China. So, it was the first time we'd ever set the thing up there and dealt with the acoustics was on the night before. He was kind of nervous about it and he said, "If this goes wrong do I get fired?" And I said, "No, you don't get fired. Your boss gets fired." Okay. But, it was fun.

So, I got up on stage and I'm going to show you a clip, and the key thing to keep in mind here is that we're using a deep neural network technology developed at Microsoft Research to do the underlying recognition of my speech, so taking that to English text. We're using the Bing translator technology, which also comes out of Microsoft Research, to translate that into Chinese text. And then we're using another technology developed in Microsoft Research Asia, which actually uses my own voice and speaks Chinese, and takes the Chinese text and renders it into speech in my voice. So, I'll let you see the video. Can you roll the video, please?

(Video segment.)

All right. So, you had a chance to see that. Of course, I was looking at lights. I didn't really get to see the audience very much. When I was looking down there I was looking at the monitor, so I could see whether things were sort of working. At one point when I was looking off to the side I was checking the guy behind the curtain to see if everything was going okay with the demo.

There were Microsoft researchers in the audience, though, and what was interesting was how affected the Chinese students—there were about 3,000 Chinese students there—seemed to be. One of our researchers said one of the young women near him was actually crying. This was something they'd never seen before and potentially was going to be able to change their lives. So, it's pretty exciting work.

To give you a perspective on the error rates, by the way, they actually did the recognition of my entire speech and we had about a 7 percent error rate. That's really good. I mean to put that in perspective a human being doing the same task typically gets between 2 and 4 percent errors. So, we're getting very close to human recognition rates with some of this work and it's very, very exciting.

END