

Cognitive Computing, Security and Innovation's Promise

Speakers:

Dr. John Kelly, Senior Vice President, Cognitive Solutions and IBM Research, IBM

Moderator:

David Kirkpatrick, Chief Technologist, Technonomy

(Transcription by [RA Fisher Ink](#))

Kirkpatrick: John is basically Ginni Rometty's number two at IBM. He runs IBM Research. He's been at IBM since 1980. He's still going strong. He's responsible for cognitive computing at IBM, which is one of their central strategic priorities. Oh, you wrote the book with Steve Hamm, *Smart Machines: IBM's Watson and the Era of Cognitive Computing*.

Kelly: Before it was a popular subject.

Kirkpatrick: With a great journalist, Steve Hamm, who spent many years at BusinessWeek. So anyway, we're going to talk a lot about artificial intelligence and other issues that have been prominent throughout our discussions the last couple days. But maybe you should start, John, by talking about what the big picture is at IBM about what it's trying to accomplish as a company.

Kelly: So at 106 years old, we've seen it all, David. We've reinvented ourselves many, many times. I've been part of three or four of these major transitions since 1980, beginning with the PC and services and software, and now cloud and cognitive AI. So we're literally reinventing ourselves again, which I think makes us very unique. The companies that we've seen in the tech industry who have not reinvented themselves are gone. They're in the bone pile.

So we're doing it again, and we're doing it around three imperatives. One of course is artificial intelligence and cognitive on a cloud platform and through an industry lens, which means enterprise. And so we have chosen to focus on enterprise, on AI, on cloud, as opposed to consumer. Many of our clients serve consumers, but that's where we focused. The reason we believe that this is the right direction for us is we see an immense opportunity—we think it's on the order of two trillion dollars, which is twice the size of the classic IT industry—

Kirkpatrick: The opportunity for what? Cognitive computing at large?

Kelly: Cognitive computing in the enterprise. So decision support is the goal. Having been around the industry for so long, I think there's a reason why this is occurring right now, and we'll talk more about it, and it has to do with exponential curves. The first one we know about is Moore's Law—double transistor count every 18 months, double performance every 18 to 20 months—and look what that's done. It took us from the original IBM System/360 through the compute power you have in your smartphone.

The second exponential, which has resulted in all the great networking companies, all the Internet companies that you talked about today, is Metcalfe's Law. And Metcalfe's Law basically says every time you add another node on a network, the value of that network goes up as the square of that number of nodes. So it's an exponential curve. And that is why you see the Googles and the Facebooks and the social media companies gaining value, because of the expanding network. Now, the assumption is that the value and quality of everything on that network is good. We can discuss that later.

Kirkpatrick: That's not the assumption anymore, but it has been the assumption.

Kelly: That's been the assumption. As soon as that falls apart, by the way, Metcalfe himself said that exponential flattens out.

But the third one, which has me really excited right now, is that data is doubling every 12 to 18 months in the world. And if we can harness—

Kirkpatrick: That's an amazing stat.

Kelly: Yes. If we can harness that data in some manner, shape, or form, and improve our decision making, then we can learn on an exponential as humans—to the point of this conference—and we can extract value from that data that today is hidden in 80 percent of that dark data. And the way to do that is through artificial intelligence. So if we can use AI to harness the capability and the knowledge in that data, then our decision making, and we as humans, go up that exponential curve with it.

So that's why we're so excited about AI. It's not just another passing fad. It's not often you get these exponential curves, but we are in one right now.

Kirkpatrick: One of the things that's come up in a number of different ways on this stage, which I know is something you think about a lot, is this issue of, if that's the case, which I think is not a controversial point of view, what is the intersection between actual thinking, breathing people and that? And we opened our conference with a session on what we call the convergence of man and machine with the chief scientist of Alexa and Mary Lou Jepsen and two other really big thinkers about this. And Mary Lou Jepsen is actually doing brain reading technology using ultraviolet light beamed into the brain to measure oxygen uptake in the brain

cells and then ultimately do pattern matching with thoughts in the cloud so you can really figure out what people are thinking. And that's a kind of scary idea. She's working on more pragmatic things like MRI replacement, etcetera, too. But if we're going to try to take advantage of that doubling of data every 18 months and retain our humanity, what are our challenges from your point of view from that perspective?

Kelly: First of all, we often think of man versus machine. But every study I've seen, and all of our experience with AI, is that man and machine always beats a man or a machine. And I've seen it time and time and time again with Watson, that Watson will be trained by humans, it will hit a roadblock, it will get more human input, some perspective, it will start learning again, and back and forth. And I'll see the humans who are interacting with the system, whether it's a doctor, a lawyer, a tax accountant, call center operator, get smarter at the same time. So it's this back and forth between the man and the machine. So that says then that we have to do our best—and I always call it impedance match, the human and the machine, which says that we have to make the machine more human-like in the way it communicates.

For instance, the machine needs to understand is the human is understanding what it's doing? If not, it needs to train in a different way. It needs to explain in a different way. On a similar vector, the human has to understand what the machine is doing so that it can improve its intelligence going forward. So this impedance matching of man and machine is really critical. The way it manifests itself in our business—because we're focused on the enterprise, so we're focused on healthcare or legal systems or financial services—is getting AI and cognitive into the workflow of these professionals so that it becomes man and machine. And it doesn't take a neat gadget to do that, but you have to get it into the workflow of what the humans are doing in order to really extract the value.

Kirkpatrick: Would you go so far as to say that, maybe down the road, pretty much everybody in their work is going to be intersecting with AI in some respect?

Kelly: Absolutely.

Kirkpatrick: From the CEO on down?

Kelly: I thought this several years ago as I started to see what our AI system, Watson, was doing. I cannot think of an industry or a human activity or a decision that we make that can't be augmented by a machine. And I've seen this pattern now, sort of roughly speaking, where we as humans, a third of the decisions we make are good decisions.

[LAUGHTER]

Kirkpatrick: Okay. That might be high.

Kelly: A third of the decisions—it might be high in some places. Not this room. This room is probably two-thirds. A third are good decisions, a third are not optimal decisions, they're sort of

okay, and a third are bad decisions. And I've seen this pattern now in every industry and every human decision.

Kirkpatrick: That's a philosophical point.

Kelly: Whether it's a top ranked physician or someone in a call center and everything in between. And so I think the advantage here, to your question of is it going to be involved in everyone's decision, is it will make those top decision makers very optimal and carry their learnings for the other two-thirds, it will take the third that are so-so decisions and really make them optimal, and it will prevent, by and large, the bad decisions.

Kirkpatrick: Prevent most of them, really, do you think?

Kelly: If we can get it in there, and if the human-machine interface is correct, then it will prevent those bad decisions.

Kirkpatrick: Let me just quickly interrupt with that, because so much of what we do is interacting with other people. And let's face it, probably a large number of those mistakes and bad decisions relate to what we say or do with other people. Would you go so far as to say it would apply somehow in that scenario too? I know that's not your business per se, but just from what you know, do you think that will even change?

Kelly: Yes, I do. Now, this gets into another whole dimension of this conversation, which is that the machine has to be trained properly. I often say these AI machines are as dumb as a rock until you give them data and then they start learning on that data. So you need to ensure that you've got the right algorithms and the algorithms are behaving appropriately, that there's some transparency in what those systems are doing, and that the data you're feeding it is good data, unbiased data, whether it's healthcare data, or population data, or financial data. So making sure you're educating that system with the right information is extremely important. If you don't do that, then you're just going to propagate bad stuff.

Kirkpatrick: One thing that IBM is—when you talk about making the intersection between people and the machine more human in a certain way, I think about voice interfaces, which is one of the biggest things that's happened in your world. IBM is not particularly associated with voice interfaces per se. Many of your so-called competitors are—Amazon you are, Google to some extent. But I don't think of you as—or maybe you are. You're head to head with Microsoft and Google and Amazon and all three of them are focusing on voice interfaces. How do you think about that as a factor in this?

Kelly: First of all, I think what we're doing and what many of the companies you mentioned are focused on is different. They're very focused on a consumer, you know, put a device on a table and it's basically a microphone that does a voice translation into text and then does a text search on the web or something. That's fine.

Kirkpatrick: They claim it's AI in all three cases.

Kelly: Well, I guess there's some in the back but it's basically a voice interface to a search engine, fundamentally. What we're trying to do is much, much different, which is we—as an example, we have focused on healthcare as one of our industries. So we have trained Watson deeply for five years with Memorial Sloan Kettering Cancer Center on all of their data, on best practices, on what to do and what not to do, on clinical trials. We've read all the literature in oncology—Watson has. I say “we.” Watson has. I do self-identify with Watson.

[LAUGHTER]

Kirkpatrick: I see. So the conflation is happening in real time.

Kelly: It's happening. We're inseparable. But we have trained Watson and Watson will look a patient's records, thousands of pages. It will look at all the literature. It will look at what it learned from Memorial Sloan Kettering. It will consider the training it received at Memorial Sloan Kettering. And then it will go through and say, “Okay, we think for this patient that you should follow this protocol. Here's all the reasons. Here's the literature if you want to look at it. Here's the previous patients if you want to look at it.”

By the way, I stopped on the way out from New York yesterday at one of the country's leading cancer centers and spent four hours with a room full of oncologists, looking at some of the diagnoses. And what they want to do is—they're not interested in some voice thing talking to them. What they want to understand is how did Watson—for this patient and this cancer review board, how did it come up with this recommendation, and they want to probe on what literature did you find, Watson? Click through, read the article.

Kirkpatrick: It's all available.

Kelly: It's all available. We are totally transparent in how Watson was trained at Memorial Sloan Kettering. We are totally transparent in the algorithms we use, and it's not a black box. Any physician can click through the entire decision tree for how Watson reached that conclusion. We think in the enterprise that's important and for us, the man-machine interface is that back and forth on, “Watson, how did you conclude that? Well, maybe you should have done this. Maybe we should get a new MRI before we make that recommendation.” That to us is man-machine interface in an enterprise setting.

Kirkpatrick: So that transparency that you just described, is that going to be standard in the use of cognitive computing at IBM, in your opinion, across the board?

Kelly: Across the board.

Kirkpatrick: That's a quite strong contrast to the way AI is thought of in the consumer realm. Extremely dramatic contrast.

Kelly: Yes. So we've taken a position publicly, and we've come out with a policy around data and AI which says very, very clearly that your data is your data. We will take it, we will process it, we will have Watson provide insights back to you, but it's your data and those insights are yours. We're not building some grand model to use against your competitors, number one.

Number two is we will tell you when we're using AI. Number three, we will tell you how the AI was trained and where it was trained. And then lastly, we'll provide transparency as to how Watson made that decision.

Kirkpatrick: Have you ever thought that maybe you should get the leaders of Facebook and Google and maybe Amazon in and give them a little tutorial in that?

Kelly: Yeah, next question.

[LAUGHTER]

No, look, I think that transparency, David, in the world we're going into, we all have a right to understand how and where AI is being used and how it was trained. And I don't know if we'll get into this topic, but you can really corrupt an AI system with bad data or biased data or other things. And I think if we're going to use AI in our decision making, we need to know that.

Kirkpatrick: Really interesting. Just before we move on to other big questions and we go to the audience, you guys have done a lot of great stuff, but there has been a bit of a hiccup in the last year or two around MD Anderson's contract and the meme is sort of out there that Watson has been disappointing in some realms. How do you respond to that? What's the reality there?

Kelly: First of all, I'm thrilled with the progress that we've made. I don't think we've had real problems. I know all of those cases. There's some press that runs around, they'll find some doctor in Timbuktu that says, "Yeah, Watson didn't teach me anything new." Okay, well Watson's an oncologist, basically. I've had doctors say to me, "Well, I already know all of that." Okay, you're part of that one-third that always gets it right, you know, the top of your field. I'm targeting the other end of the spectrum as well.

Last week, I was with the CEO of Manipal Hospital in India and he basically said, "John, we have been waiting for this." There's 1.3 billion people in India. Three hundred million of them can afford to go to an oncologist, God forbid they have cancer. There's only 750 oncologists in India. So even the 300 million can't get access. He said, "John, if I had the money, I couldn't build enough medical"—

Kirkpatrick: Yes. This is a key way to think about things.

Kelly: So we're scaling knowledge. So yes, the third who are always getting it right are going to say, "Well, there's nothing new." Maybe some of the trailing third are going to say, "This is

threatening to me.” But, okay, we want to get rid of those bad decisions and make good decisions. So I think it comes with the territory. People are going to have concerns about it. People are going to question it. I just know what we’re doing. We’re doing it right and we’re reaching good decisions.

Kirkpatrick: From a revenue point of view, is it—you talk about it a lot, but it’s hard to tell what kind of a business is being created. Talk a little bit about that.

Kelly: Our business approach, again, is all enterprise. It’s part of what we call our strategic imperatives, which is growing double digits, very strong in our new businesses. We look at it the following way. We have chosen three vertical industries to go in big time: healthcare, Internet services and Internet of things/industrial. So those are big, big businesses right now built around Watson AI. And then we’re taking Watson AI and infusing it on other businesses. So I have a very big, two billion dollar cybersecurity business, number one or two in the world in cybersecurity, and we’re putting Watson into our cybersecurity capability to basically look for and catch the bad guys in that business. And then the same in analytics and elsewhere. So we’re infusing Watson in all of our software and services and then we’re building new industry verticals to drive IBM.

Kirkpatrick: So in other words, you’re saying it’s so central to IBM that’s it’s hard to even sector out where the revenue is because it’s affecting everything.

Kelly: Yes. As an example, we have clients on our mainframes, banking clients, who will not take that financial services data out of that mainframe. We took machine learning and put it into the central databases in our mainframe for those clients to do fraud detection. Because they can’t it out and put it on a cloud or put it on another server. So we’re basically infusing it in all of our products.

Kirkpatrick: So many things I want to talk to you about. As usual, we’re running low on time and I want to get to the audience. But when you look at the security problem, which is so massive at the moment—I mean, it really is one of the overarching societal challenges we face. Do you think that what you know about AI and IBM’s longstanding focus on security and deep knowledge—do you believe that we will be able to stay ahead of the bad guys and more or less keep society functioning sufficiently?

Kelly: Yes, but with this exponential curve comes a new set of threats, David. So if you think about cybersecurity in big animal pictures, we started with physical security, glass houses, lock the mainframes up in there, badge lock, blah, blah, blah. Then with the Internet, we went to firewalls, keep the bad guys out. And then most recently, we’ve gone to, okay, even the firewalls are permeable now so the bad guys are getting in. So if they’re in, we’d better find them. We’d better find the malware. And I mean, read the papers, there are companies that go two years before they find out that something bad is inside.

Kirkpatrick: Or political parties.

Kelly: What's coming is, in my mind, even more threatening, which is AI bots, AI malware, if you will. These are intelligent devices and things that will come into your network, into your system, change themselves, learn on the fly, learn what is attacking them, and my prediction is that cybersecurity is going to turn into an AI versus AI war, that the only way to defend from these AI things that are coming at you is with smarter, more intelligent artificial intelligence agents that will find, trap, destroy, and divert those things that are coming at you. Because they're coming not only at the rapid speed that we see today, but they're going to morph and change because they're artificial intelligence bots. So it's going to become an AI versus AI war. Therefore, we had better lead in AI in this country, and the companies that build those capabilities and use those capabilities will be okay. The ones that don't, as they say, you ain't seen nothing yet.

Kirkpatrick: That also really underscores the need for transparency on the good guys' side, except that it also suggests that has its risks, in addition, because the bad guys will potentially get access to the transparency. But rather than go down that path, which we could go on for half an hour, who has a comment or a question in the audience? Please identify yourself.

Long: Hi, I'm Jane Long. I applaud your transparency comments, but I'd like to ask you to expand on it a little bit. One of the things that happens in very complex systems when you try to be transparent is you can just dump a lot of data out there and it's very hard for, especially for common people, the public to understand it. And there's a concept called meaningful transparency which really describes what all that data means as part of the transparency process. Have you looked at that and could you expand a little bit on what transparency means to you?

Kelly: Yes, so we don't just dump our clients' data out there, or the results of it. We are very clear that we will share with them, the owners of the data, exactly what we're doing with it, exactly where it is, exactly what's touched it, exactly what operations have been performed on it. We have not, nor will we ever, ever take our clients' data and just put it out there or put a recommendation forward say for a cancer treatment or something and not be able to explain, or offer to explain, to that patient or that customer exactly how that decision was reached. That doesn't mean I have to show them all of the rest of the data. I'll just say, "Trained at Memorial Sloan Kettering. Here's all of the references. Here's all of the supporting evidence. If you want to understand the machine learning, I'll explain it to you." We'll be very transparent on that decision, and we can do that very easily and very well, versus having to do a data dump of everything. Because then I would violate the other principle, which is I do not share a client's data with anyone else.

Kirkpatrick: I want to ambush somebody. Dan'l, I saw you listening very intently. I'm curious if you have anything you heard John say that you think is particularly noteworthy, surprising,

you disagree with. And if not, you don't have to say anything, but Dan'l is a longtime Microsoft guy, recently retired, one of the deep thinkers in the technology industry.

Lewin: Yes, I have been listening and I think it's terrific, the way you're going about it and I think it's an appropriate and well-placed strategy for IBM. I do believe that it's one of the broader questions that you're talking about relative to security, gets into the edge case of geopolitical boundary lines and where the data resides and how you have to comingle data outside of these vertical domains for the broader use of artificial intelligence, the public and private data, as opposed to these industry specifics. But what you're doing and how you're going about it I applaud fully. I think it's great.

Kelly: Thank you. I should just remind everybody that it's all about protecting the training data, not just the training algorithms. So today, in a programmable system, you go in and you hack in and you do something to the programming. In these systems, you have to protect that training data because slight changes in that training data can cause that system to learn and do things that are orthogonal to what you think. So it's all about the integrity of that data and the source of that data. Now, your methods to protect it are totally different around the data versus around the program.

Kirkpatrick: And you also may have some real advantages over the Facebooks, Googles, Amazons, and Microsofts as well, in that you can really do geographic segmentation in a way that is much harder in their systems.

But anyway, here—we'll try to get a couple of voices on this floor before John answers and then we'll have to wrap.

Hill: I'm Zach Hill from the Future Project. I just wanted to take a step back and ask you as IBM, a 107-year-old technology company, I think it's easy to get swept away in the momentum and valuations of the Amazons, the Googles, and the Facebooks. But as somebody on stage mentioned yesterday, 10 years ago it was eBay, Yahoo, and AOL. What is your perspective on this current moment in technology, and what is the role of a company like IBM in the face of so much change? How do you endure, and what do you think is the perspective that you bring that's relevant that other people might not be able to have?

Kirkpatrick: Okay, answer that. That's a good one.

[LAUGHTER]

Kelly: Yes. Well, just in the course of my 38 years, I've seen them all come and go. I've seen these flashes and you say, oh, their tree's going to grow straight to the sky, and then the tree doesn't grow straight to the sky.

Kirkpatrick: And sometimes it does.

Kelly: And sometimes it does. But look, we view ourselves as someone that is a trusted source of technology, someone that you can bet on that we're not just going to be here this year or 10 years, we're going to be here 100 years from now. We're going to, as I said, reinvent ourselves to be here. And we are responsible. We build the systems that run the world's financial services. All of your bank accounts, all of your credit cards, most of your health records sit on our systems, and highly secure. We encrypt everything. So we view ourselves as sort of the responsible party in the enterprise. We sort of view the stuff going on in the consumer world as interesting. But don't forget, 80 percent of the world's data isn't searchable on the Internet. It sits either on our systems or in our clients' systems and that is probably the most valuable data in the world. And our clients say to us, "IBM, we'll put this on your systems, we'll trust Watson with this, but if you ever think we're going to put this 80 percent of the data out there—not in a million years." So the opportunity for us is enormous, just with that. And we take that responsibility seriously because of what our brand stands for.

Kirkpatrick: Okay. Was that Eric who had his hand up? We don't want to go without listening to Eric talk. Go ahead.

Eric Topol: It's great to hear you, John, and high regard for everything you and IBM and Watson are doing. But the thing that David mentioned about the meme and the very intense public relations blitz, the question there is it isn't tagged with the papers and the peer review journals to back up a lot of the things that you see in the full-page ads, like the doctor that reads 5,000 articles and then goes to see patients and that kind of stuff. And you have, for example, "60 Minutes," the segment you're on with UNC, where 30 percent of the patients at UNC cancer center were affected by Watson, but not paper yet to publish to back it up. So can you be more transparent about that, because the medical community would really appreciate it?

Kelly: Yes. So that UNC paper is going to be published in the next 30 or 60 days is what I was told yesterday. We have 37 peer reviewed papers that are out in peer review now to be released. Just as we're transparent with what Watson does and how it does it, we're going to be totally transparent, peer reviewed papers with leading institutions around the world. And I do think that that will help cause some of these guys to just go away.

But it is real. The results are incredibly impressive. Everything said in that *60 Minutes*—when you read that peer review article, Watson agreed with their docs on 1,000-person retrospective study, 98 or 99 percent of the time. In the one case it missed, it didn't have all of the lab results. And then 30 percent of the time, it found not just new possibilities, but things that humans agreed, "Yes, that was actionable, and I should have considered that." So that will all be in these peer reviewed papers and they're going to start rolling out 30 or 60 days from now.

Kirkpatrick: Okay, well thank you for that. I'm afraid we have run out of time. John, thank you so much.

Kelly: David, thank you, sir.