

THE 16

The Evolution Revolution

Speakers:

Benjamin H. Bratton, University of California, San Diego

Ray Kurzweil

Vivienne Ming, Socos

Moderator:

David Kirkpatrick, Founder and CEO, Techonomy Media

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

Kirkpatrick: Ray Kurzweil is one of the world's leading inventors, thinkers, and futurists. He's been called by PBS, one of the 16 revolutionaries who made America. And this is amazing, this paragraph, "He was the principle inventor of the first CCD flatbed scanner, the first omnifont optical character recognition, the first print to speech reading machine for the blind, the first text to speech synthesizer, the first music synthesizer capable of recreating the grand piano and other orchestral instruments, and the first commercially marketed large vocabulary speech recognition.

He's gotten Grammys. He's a recipient of the National Medal of Technology, in the National Inventors Hall of Fame. He has 21 honorary doctorates. He's written five national best-selling books, including "The Singularity is Near" and "How to Create a Mind." He's cofounder and chancellor of Singularity University and now a director of engineering at Google for the last four years, heading up a team developing machine intelligence and natural language understanding.

And what we're going to be talking about is very much where AI is headed, what AI means in society, whether there is a singularity coming, what it would mean if it does happen—obviously an idea that Ray has been very closely associated with for a long time. But I wanted to just start, Ray, by asking you what is the main thing you're working on right now?

Kurzweil: Well, I've worked for 50 years or more on AI, actually went to meet the two leaders of the two principle schools, the symbolic school, which is Marvin Minsky, and the connectionist school, Frank Rosenblatt, who created the first neural net. This was in 1962 when I was 14.

Kirkpatrick: When you were 14? You went to meet both of them when you were 14?

Kurzweil: Yeah. So Minsky became my mentor for 54 years, until his passing recently. And so there was a lot of rage about thinking machines based on the perceptron, which supposedly simulated a human brain. And it could recognize printed letters, so I brought up printed letters in Courier 10 and it could recognize them. But then I showed it a different typestyle and it

THE 16

didn't work and Rosenblatt said, "Well, don't worry. If we feed the output of the perceptron to another perceptron, another neural net, and the output of that to a third one and so on, it will get more intelligent and generalize and be able to recognize abstract patterns." I said, "Oh, did you try that?" and he said, "Well no, it's high on our research agenda."

Well, he died nine years later, in 1971, never having tried that. It would be decades later that they finally tried these multiple layer neural nets, three or four layers, and they did get a little more intelligent. It could then recognize multiple typetypes. But as recently as five or six years ago, the accusation was AI can't even tell the difference between a dog and a cat. It turns out the essence of a dog and a cat is pretty abstract and subtle, and it's at layer 15. We couldn't go beyond three or four levels as of a few years ago because of a math problem. The information as you went from one layer to the next would sort of aggregate in a small part of the space. You need what's called a convex error surface. I won't explain what that means.

A group of mathematicians, including one of my colleagues at Google, Geoff Hinton, solved that math problem, basically taking the information after each layer and spreading it out. And now you can go to any number of layers. So the program that won the Go Championship recently was a 100-layer neural net and 100-layer neural nets at Google and other companies can now recognize a dog and a cat and thousands of other categories of images and do it better than humans.

That actually accounts for the tremendous surge in interest in AI. There is a problem though, and that's what I address in my most recent book, "How to Create a Mind." There's a motto in the deep learning field that life begins at a billion examples. And we have a billion examples at Google and other companies, which is one of the reasons I'm there.

Kirkpatrick: Google has a lot of examples, yes.

Kurzweil: Like dogs and cats. We have lots of dogs and cats. We also have a lot of language, but it's not annotated with what it means. And we don't even know how we would annotate it. So that's the big challenge now in AI and it's what I'm working on and it's what I talk about in my book, how can we learn from less data? Humans can learn from a small amount of information. Your significant other or your boss tells you something once or twice, you might actually learn from that. Humans don't always need a billion examples to learn.

So that's the big frontier. I'm working on that. I'm working on understanding natural language. Long term, we'd like to move capabilities like search and language translation from keywords to actually understanding the meaning of what's being written about.

Kirkpatrick: When Google comes out with something like Google Home, are you at all involved in that?

Kurzweil: Well, I have a research team that's working on this general problem of language, so we're involved in language tasks.

THE 16

Kirkpatrick: So maybe a little involved with that one. That looks like it could be pretty cool. By the way, I never even told my people, we have one coming at the office as soon as we get home. Because they just shipped them, I guess last Friday.

So are you still thinking a lot about the singularity, and what's your current sort of estimation of what that means?

Kurzweil: Well, I have two key dates, 2029, when computers will actually understand human language at human levels. That's what we call a Turing complete task, meaning it needs the full range of human intelligence to do that. So to pass a valid Turing test, a machine couldn't use some cute natural language processing tricks. It would have to actually be at human levels of understanding. And I've said consistently we'll achieve that by 2029.

Now, we can already do a good job if it can read at levels that are not quite as good as a human but make up for that by reading more documents. So Watson, for example, from IBM read 200 million pages, including all of Wikipedia. It might read one page and say, "Ah, there's a 56% chance that Barack Obama is President of the United States"—not for much longer, but—and you might read that page and if you didn't happen to know that, conclude that it's a 98% chance. So you did a better job than Watson of reading that page. So why is it that Watson could get a higher score than the best two human players in the world at "Jeopardy," which is a broad knowledge and language game? Well, it makes up for its weak reading by reading more pages.

So that's the threshold we're at now. So what I'm saying is by 2029, they will read at human levels, do that on billions of pages, and also be human in other types of tasks. One after another, tasks that we used to associate with humans have fallen to AI. You know, people said, oh, well computers will never play chess and so that happened in 1997. They'll never play master levels of Go, they'll never drive a car. But one after another, these kinds of tasks have fallen to machines. They'll be able to achieve the full range of human intelligence by 2029.

Now, I wrote that—actually, I gave a range around that date in my 1989 book "The Age of Intelligent Machines." I said 2029 in "The Age of Spiritual Machines" in 1999. Stanford held a conference on this startling claim and we took a poll and the consensus of AI experts is it would be hundreds of years, if ever. Twenty-five percent thought it would never happen.

2006, there was a conference at Dartmouth called AI@50 on the fiftieth anniversary of the 1956 Dartmouth conference that gave artificial intelligence its name, and the consensus then was 50 years. So it was moving closer. I was still saying 2029.

Kirkpatrick: The 'it' being—what is the target exactly?

Kurzweil: Computers achieving the full range of human intelligence.

Kirkpatrick: And is that what you call the singularity?

THE 16

Kurzweil: No.

Kirkpatrick: Okay. Because I want to make sure we define that—

Kurzweil: But that's a stepping stone.

The next step is—I mean, I talk a lot about exponential trends, so the exponential growth of information technologies. But another exponential trend is miniaturization. I mean, this little computer is actually a billion times more powerful per dollar than the computer I used as an undergraduate, and it's also a hundred thousand times smaller. We'll do both of those things again in the next 25 years. So in the 2030s, we'll have robotic computerized devices the size of blood cells. They'll be intelligent. We'll have millions of them in our bloodstream. They'll keep us healthy by extending our immune system. They'll go inside our brain and feed signals directly to our brain as if they were coming from our real senses, our eyes, ears, tactile sense, providing full immersion virtual and augmented reality from within the nervous system. And most importantly, they'll connect our neocortex to the cloud.

Now, I mentioned this is a billion times more powerful than the computer I used as an undergraduate, but that's not the most interesting thing about it. If I want to extend its capability a thousand or a million-fold, it communicates wirelessly today with the cloud and we can access the whole of human knowledge which doesn't fit in the phone. We can't do that from our brains directly. We do it indirectly through these devices. We'll connect our brains directly to the cloud, 2035 scenario. Not just to do capabilities like search and translation directly from our brains, although we'll do that, but to actually extend the scope of our neocortex, so instead of having 300 million pattern recognition modules, which is my description of how the neocortex works, we can have more. We can have a billion or two billion. And then our thinking will be a hybrid of biological and non-biological thinking. However, the non-biological part is subject to the law of accelerating returns. The cloud is doubling in power now every year as we speak and that will continue.

We got more neocortex two million years ago, you might remember. We got these big foreheads. That gave us additional neocortex. We put it at the top of the neocortical hierarchy. That was the enabling factor for us to invent language, and art and science and conferences on technology. No other species does that. But that was a one-shot deal.

Kirkpatrick: The most important thing of course.

Kurzweil: That was a one-shot deal because our bigger foreheads became a challenge for childbirth and if it had continued to grow, we wouldn't be able to be born. This next expansion, where we connect wirelessly our neocortex to synthetic neocortex in the cloud will not be a one-shot deal because the non-biological part will grow exponentially.

THE 16

So now, the singularity, if we do the math according to my formulas of the law of accelerating returns, we'll multiply our intelligence a billion-fold by 2045. That's such a singular transformation that we borrow this metaphor from physics and call it a singularity.

Kirkpatrick: But it's a billion times more capability than the human brain can achieve on its own.

Kurzweil: Unenhanced. And we're already not unenhanced. That's a double negative. I mean, we are enhanced through these devices, but we'll directly multiply the scope of our neocortex. Our neocortex, I describe in "How to Create a Mind" as a series of modules. We have 300 million of them. This has actually been recently confirmed by the European brain reverse engineering project under Henry Markram. He's noticed a module of 100 neurons that's repeated 300 million times throughout the neocortex and they're all basically the same. And there's no plasticity, no change, no rewiring within each 100-neuron module, but constant change between the modules, and that's consistent with my thesis. And we have 300 million. We got this additional shot of it, as I mentioned, two million years ago with these big foreheads. We will connect to synthetic neocortex. That's something else I'm working on is basically simulating that basic process. We don't understand it fully, but we have some clues and we can begin to simulate it.

By the 2030s, we'll have very perfect simulations. We can then connect to the cloud. And just as our phones today multiplied our intelligence by connecting to the cloud, we will do that directly from our neocortex.

Kirkpatrick: You mentioned before that we have connected devices of some sort throughout our body in the millions. I remember once I interviewed you when I was at Fortune, probably 15 years ago at least, and you were talking at that time about each individual synapse you thought would have some kind of connection to an external information—

Kurzweil: Well, not every synapse. First of all, the region of the brain that's significant and where we do our thinking is the neocortex. Neocortex means new rind and it's the outer layer of the brain. But it developed so many curvatures—you know, you've all seen that image of the curved surface of the brain—that it's now 80% of the brain. And it's organized in a hierarchy, so at the bottom level of the hierarchy, I can tell that that's a straight line and that this is a curved line. At the top of the hierarchy I can tell, oh, that's funny, that's ironic, she's pretty. You might think those are more complicated, but it's actually the hierarchy below them that's more complicated.

The hierarchy is organized like a pyramid, so the top of the hierarchy has relatively fewer modules and at the bottom we have a huge number that do these simple features of visual images and so forth. The ones that are intelligent, that organize technology conferences, for example, are relatively fewer. [LAUGHTER]

THE 16

Kirkpatrick: Simone does that, actually. [LAUGHTER]

Kurzweil: So we will connect the top layers to the synthetic neocortex and add additional levels of the hierarchy, so we will become funnier and more musical and so forth.

Kirkpatrick: That sounds good. [LAUGHTER] Before we bring up the other panelists here, talk a little bit about how you ended up at Google. Because a lot of us found that surprising. You have always been kind of a pretty much self-employed free agent/entrepreneur. Explain what happened there.

Kurzweil: I was a serial entrepreneur. I started probably ten companies, five of which I've sold and five of which are still going. So my most recent book I mentioned, "How to Create a Mind," which talks about how the neocortex works and has a very particular model, including the algorithms that I believe run in these 300 million modules, and how to then build AI based on that. And it's different than the neural nets somewhat in that it can learn from less data because humans do that. That's kind of the unique feature we still have over these deep neural nets.

I gave an early version of it in 2012 to Larry Page. He liked it. I met with him and asked him for an investment in the company. I had just started to develop these ideas and I asked him for an investment and he said he would invest but, "Let me give you a better idea, why don't you do it here at Google? We have all these great resources. We have billions of images of dogs and cats, for example, and we have lots of computers."

Kirkpatrick: Especially cats, yeah.

Kurzweil: I think there's actually more dogs.

Kirkpatrick: Is that true? Okay. I prefer dogs myself anyway.

Kurzweil: So we had an agreement, a meeting of the minds, and so I took my first job for a company I didn't start myself later that year. So it's been about four years now.

Kirkpatrick: And is that the majority of your time now, at Google?

Kurzweil: Yeah, I'm fulltime there. I mean, I've got a few other things I'm doing, like writing some books.

Kirkpatrick: Including a novel, I believe, right?

Kurzweil: Yeah, I wrote a novel called "Danielle: Chronicles of a Super-Heroine" that was actually about superintelligence and it basically addresses the question of what would happen if a child, a young girl actually had superintelligence. And she does some remarkable things, like cure cancer and bring peace to the Middle East. She becomes the first democratically elected president of China at age 15. And there's a companion book with it which is called "A Guide for

THE 16

Super-Heroines and Superheroes,” basically how you can be a Danielle, and that’s actually longer than the novel, because it takes time to explain this to people that aren’t superintelligent like Danielle. [LAUGHTER] So that’s going to be published by WordFire. It should come out next year. And I’m writing “The Singularity is Nearer,” which will also come out by Viking, which has done all my nonfiction books, and so that should be 2018, which will be 13 years after “The Singularity is Near,” which has held up very well but a lot has happened over the last decade.

Kirkpatrick: It’s had extraordinary influence, there’s no question about that. So you enjoy working for somebody else?

Kurzweil: Well, Google is a very unusual company. It’s got 60,000 employees but it really does have kind of an entrepreneurial spirit. Remarkably, it’s very much bottom up. All the engineers and people in the other departments kind of decide what they’re going to do. They’ve got some guidance in terms of what the goals are of Google, and somehow it all self-organizes. It’s not top down like, “Okay, we’ve got this project, you’re going to do this and you’re going to do that.” So it’s very creative. It’s sort of moonshot, ideas are very well publicized. It’s a bold company. For example, it bought this little app with kind of cute little cat videos for a billion dollars and people said why would they spend a billion dollars on a little app that has cat videos? Today, YouTube is bigger than network television. So they do a good job of really realizing these visions.

Kirkpatrick: Well, let’s broaden our conversation a little by bringing up Vivienne and Benjamin. Vivienne Ming is chief scientist at ShiftGig and formerly of Gild. She’s finding ways to use technology to connect directly with job opportunities, so people use it to connect their skills and talents directly with job opportunities. She also wrote a piece in our magazine that’s really quite interesting. It’s a fictional piece imagining what would happen if an AI were to replace a financial analyst and it sort of envisions that very afternoon and ends on a somewhat Trumpian note. So you were somewhat prescient because your main character ends up more or less heading to a Trump rally.

She’s also cofounder and managing partner of Socos, a cutting edge educational technology company which applies cognitive modeling to align education. And basically, she’s a theoretical neuroscientist—is that the right terminology?

Ming: That’s what we call it at UC Berkeley.

Kirkpatrick: What does that mean?

Ming: That means we’re lazier than all the other neuroscientists. We just make fake brains on computers and teach them how to do things and pretend that that tells us something about how real brains work.

Kirkpatrick: I like that definition.

THE 16

Ming: If you want to make it a little fancier-schmancier, think like theoretical physics for the brain. Start from first principles and see if you can get insights about cognition, sensation, emotion from those insights.

Kirkpatrick: But it's interesting, you're applying that to education and job/employment issues. I think that's fascinating.

Benjamin Bratton, who we can thank in part Pradeep Khosla of UCSD for helping connect us to, is a really interesting thinker about the brain, about cities, about artificial intelligence. He's coming more from a design mindset. He's a professor of visual arts at UCSD and he's director of the Center for Design and Geopolitics. But you'll find he's really thought a lot about AI. And I guess I would maybe start, Benjamin, just talk about how you think about what AI is, what you think about the idea of the singularity, any thoughts about what Ray's been talking about or what's important to you when you think about this concept.

Bratton: My work is really about the culture of AI, let's say, the kind of social milieu in which, as AI emerges, that it's able to take root. And so while the technological issues that are involved in this are enormously important and interesting, and I happen to think that as a kind of meta technology, that AI has the capacity to transform our societies and our economies, our cultures at an absolutely fundamental level, from identity to what counts as governance to all of the above.

Kirkpatrick: And governance is something you've written a lot about.

Bratton: Yeah, and it continues to be an issue of some intrigue.

Kirkpatrick: We've been talking about it in the last 24 hours.

Bratton: Yeah, and I think the question of AI plays into this story about what happened yesterday quite directly.

Kirkpatrick: How so?

Bratton: Well, let me get to that. Let me go a little bit step by step. I mean, one of the things that my work around AI has been focusing on is let's say a suspicion of our anthropocentric biases around AI, and it would go back to Turing's 1948 and 1950 papers by which he introduces a notion of AI—

Kirkpatrick: This is Alan Turing.

Bratton: Alan Turing, right. Which is really as a sort of sufficient condition to AI that it would perform thinking in such a way that it would mirror how humans think the humans think and that in the reflection and recognition and empathy with this performance, that we would grant this in fact is an intelligence in this way. And I think the bias is toward the presumption that we would make sense of AI, that we'd use AI, that we would recognize AI in the world and

THE 16

understand its conditions primarily in terms of this conception of the anthropocentric predispositions has been one that has limited the conversation, the social conversation, the political, the cultural conversation about the context in which AI may flourish in ways that have been detrimental.

So my interest is in opening those conversations in such a way that we can have the questions, but how should and might AI innovate these institutions in ways in which we would want to.

Kirkpatrick: You mean you think the anthropomorphism has limited our ability to imagine a wider variety of types of AI?

Bratton: And to even recognize it in certain ways. I mean, if we think about it in this way—and as Ray would attest to, one of the interesting things that has come out of our research in evolutionary robotics and bottom up work as well is that the capacity for an embodied form of AI to sense the world around it, to see it, to smell it, is ultimately inextricable from its heuristic capacity to manipulate and interact with that world around it as well. And so ways that we have made distinctions between—going all the way back to Kant—between sensing and thinking turn out to be in certain ways less functional as explanations.

Now, if we extend this to the landscape scale—and of course, our urbanism in many ways is being defined by this distribution of sensor networks that exist at this scale. Again, where the sensing stops and the thinking begins is not always so clear.

Now, at that scale, the multiplication of these sensors, this as well, and sometimes those sensors may be being accessed by multiple AIs at different points in time and AI may be accessing different ones of those sensors in different ways. And so we could say that let's say the organism to species to niche to single to phylum dynamics that have structured our own evolution through these processes of a single tetrahedral sort of form with a single array is not the way AI is evolving at the infrastructural scale.

So I think in one way we're not recognizing what the AIs that we have are doing and where they are because they don't look like us and they don't respond to us in this particular kind of way. And I would also very quickly say, because this is where Vivienne jumps in, is that—all of which to say there's also conversations less about thinking of AI in terms of mind and consciousness and much more of AI in terms of automation and AI as an institutional problem, something that actually has an infrastructural capacity, and infrastructural capacities to me are in many ways much more thorny and much more interesting.

Ming: I was just going to say, even the discussion earlier today talked about how GE handles these real time optimization problems, which I think is fascinating, but sort of falls back onto this very—there's a central processor, the data from these various sensors are coming together into one location. I think the earlier discussion about how you would have these distributed, where all of the processing was happening locally, where you can, commercially available, get

TE16

deep neural networks embedded right on chips doing image recognition and so forth. So what if we're talking about a world in which all of that local processing is happening, that chip is pulling everyone in, it's recognizing all of you, it's picking up on the sentiment being expressed by your facial features, and then it's trading that information out. So instead of it being the Internet of things based on a platform, it's a marketplace in which that thing is now bidding out that information to your smartphones and to a security network and to the smart lightbulb system. And what's interesting there, to me as a neuroscientist, is now I'm talking about a whole bunch of local processors that are pulling in information, processing it and then distributing it openly to their neighbors, which is starting to sound a whole lot like the incredible mess we have inside of our heads.

And what's interesting about it, and ties in with your story, is when I think about what that is, it's not an autonomous car. It's a distributed autonomous transportation network over an entire region, which is not fundamentally aware of what's happening here and here necessarily at the same time if it doesn't need to be. You know, when you look at the way brains work, we're very aware of a very limited set. In fact, we're usually constructing explanations for our actions after the fact based on very limited amounts of information compared to the actual neural circuits that were involved in those decisions. I expect to see much more like that, autonomous minds which are turning parts of its internal mining structure on and off in response to its predictions about likely commodity prices for what's being mined out of it—again, at a scale, and possibly with communication limitations internally in the mind. That means there's no single system which is directly controlling it and certainly no sort of embodied bot I would expect someone to be having a conversation with.

Kirkpatrick: So both of you are sort of talking about this possibility of AI being more of an emergent phenomenon that's kind of something that happens somewhat organically as intelligence is more and more widely distributed in systems, whether it's in a city or potentially in other contexts.

Kurzweil: I can offer a brief comment on that. I mean, I agree with what's being said, but it's not either/or. We've always used machines to do the things which we're not very good at. So, you know, we can't even remember a handful of phone numbers. If I ask you to recite the alphabet, most people here could do that. If I say, okay, recite it backwards, that's a trivial operation for a computer but you can't do it. So computers can do things that we can't do and one thing they're very good at, and getting better at with the law of accelerating returns, is being able to operate at vast scale, and so we have these exponentially growing sensor networks, which is kind of the senses of this super mind—and we have very limited sensory capabilities in comparison—and it can do things that we can't do and that's always been the case.

However, human intelligence excels in certain things and computers have not been able to understand language at human levels. It can't read a novel and give you a good summary of it

TE16

and write a decent review of a movie and so on. That is really what I've talked about. And it's not either/or. The 2029 date is when they really will match human abilities in these kind of ever growing smaller category of things that we excel at. But my view is it's not an alien invasion of machines coming from Mars to displace us. We create these tools to extend our own reach, and we're already smarter because of these brain extenders.

Kirkpatrick: But you don't disagree that AI, something you would recognize as AI could emerge from say a connected urban infrastructure where intelligence was distributed at the edge in enough density that there was some kind of ultimate emergent intelligence that we would see as a new form of intelligence?

Kurzweil: We create AIs with particular tasks but then we put out millions of them in a distributed fashion and so it will create a kind of community that has its own personality. But that's true of humans too. You know, every audience that you and I speak to has a different personality and different groups have a personality that's kind of a super mind. In fact, we created communication technologies so that we could amplify our cognitive abilities by communicating with each other.

But you know, a vast amount of our knowledge is embedded in language and so as soon as computers can actually understand that, we'll unlock that knowledge and be able to then look across millions of documents and gain new insights.

Bratton: It's precisely because I think that humans and AI do think quite differently from one another that putting so much on the anthropocentric sort of models but there's a human criteria by which at a certain date or in a certain way it will achieve this particular threshold is missing quite a bit. To me, in terms of the philosophy of this—and let me speak to this briefly. To me, I see AI in this way as a bit like what in philosophy is called the other mind problem, ways in which you try to—there's a way of thinking and being embodied in the world that is different enough that the question of what our references may be that would allow us to communicate are unclear in this way.

Kurzweil: Well, I'd put a different wrinkle on that. I think computer intelligence has been different and so it can do things that in fact we're not very good at, from trivial things like reversing the alphabet to playing chess and seeing patterns in financial data that humans could never see because we can't look at a billion numbers and understand it. However, there is some unique attributes of human intelligence that's desirable for machines to master, like really understanding language and using more than simple natural language process tricks, really understanding the meaning of documents and being able to converse with us and gain insights by reading a million novels. That we're going to actually have to emulate human intelligence in order to do that. At least I think that's the best strategy, to actually understand how humans understand language and then recreate it that way. At least that's what I'm—that's my thesis and that's the direction I'm working on.

THE 16

Ming: Well, it seems like the strategy that was emerging, you know, the work on cortical processing units or, similarly, the work looking at massive datasets, sort of what you might call the Watson approach. There's a number of different academic labs that have all sort of followed this belief that if we can stuff enough newspapers into this thing's head, one day it's going to wake up and have an opinion about the way the world should be, and an assumption that the limitation is scale and not something a little bit deeper.

Now, I'm a deep believer both in artificial intelligence and in AI augmented intelligence. I actually am looking forward to talking about—you know, my lone remaining academic field of interest is cognitive neuroprosthetics, and there's this cool stuff like neural dust, at UC Berkeley, Lawrence Berkeley Lab, these little nano machines that exist today transduce local field potential into RDF signals, embed boosters in your dura just below your skull, and then this stuff transmits out and now we've got—

Kirkpatrick: That's Kurzweilian.

Ming: Yes. And now we're talking about a system with thousands of three-dimensional sensors. It turns out the data, the input problem is fundamentally harder than the readout. But that readout potential is right there and the opportunity to merge together what machines and humans are great at in a really productive way I think is exciting, although I think it is just as singularity prone as any other possible emergence of superintelligence.

You know, I actually have been saying, probably well inspired by you, that in 15 to 25 years, these sorts of technologies will fundamentally change the definition of what it means to be human. And if we're not careful, it's only going to do it for your kids.

Bratton: They already have.

Ming: So they already have. The intelligence of your kids is already a function of your wealth. There's all sorts of things which have changed people by virtue of how they interact with technology.

Kurzweil: Actually, you don't have to be very wealthy to have these technologies. There's two billion smartphones. You know, a kid in Africa with a smartphone can access all of human knowledge.

Ming: But who uses such a thing? You know, there was a great line out of "The Second Machine Age" talking about the exciting acceleration of smartphone usage around the world, and isn't that great? That means kids anywhere in the world can spin up an AWS instance and run an R data analysis model on large datasets. That may be a true statement, but I don't do that. And this is my world. Like what person spontaneously becomes that, and we haven't built a world in which they were actively creating the artists and the scientists, instead of just imagining that it will emerge because AI frees us from work.

TE16

Kurzweil: I'll give you an example of what people actually use these technologies for. Like my father, who was a composer, couldn't hear his compositions without occasionally raising a lot of money and hiring an orchestra and then he could hear his orchestral composition. Today, a kid in her dorm room can use her mobile phone or tablet computer and a mini keyboard and create a whole orchestra and then also use automated software to generate a walking baseline and so forth. And so it is amplifying creativity. We were just chatting about my daughter, who does illustration and uses all kinds of intelligent computer tools. It's used in every area of life. And also, "Gee, who said that quote?" and then you can look that up. And so that's how we use these things. It is an amplification.

Kirkpatrick: So that's augmenting our intelligence right now.

Kurzweil: Exactly, yeah.

Kirkpatrick: But Vivienne, I wanted to go back to something that I know you think about a lot, which is this question of how AI gets integrated into human society. And you've touched on it a little bit, but I wanted you to go at it a little more directly, because I know you worry that we have the likelihood of getting, one way or another, to something that we really will consider a different form of intelligence than ourselves, and yet we may not be prepared to live alongside it somehow.

Ming: You know, technology, AI or otherwise, technology is just a tool and it always will depend on what we do. I clearly believe in it, since this is what I do for a living is I build technology, particularly machine learning technology. But I don't—and this is a hard won lesson from the ed tech world is it doesn't matter how right you are or how amazing your technology is. If you've built an educational technology, you just helped no one. I'm just talking about the history of the field. We have a long history of not making any difference in anyone's lives.

So how do you confidently go in and build technologies that really actually change people's lives? You have to focus on the human side. So in this context, one of my genuine worries is—our debate seems to be—and I realize this is a trivialization, or misinterpretation perhaps of how some people have characterized some of your comments, but it tends to be this dark AI god which will destroy us versus a tech utopia where, simply by inventing technologies, the world is this amazing place, it becomes Star Trekian and we're all fully realized as a result because now we don't have to bend over and pick strawberries anymore—which none of us have ever done, but.

So I think what actually is the case is the question is as these technologies develop, as they are played out in the world with greater and greater impact and accelerating every day, if you're doing the same job last week that you are doing today, be really worried, because someone like me is going to come along and think, "I bet I can automate that, faster, cheaper, and better." And we're looking at the hyperinflation of work in which you're going to be doing a different job every week, every month, something new. One single job description: adaptive creative

THE 16

problem solver. And if we aren't building those people, then the AIs we're building will never be fully realized. Which is to say that they will slowly take over a lot of cognitive labor in some very useful ways but we won't then magically take up the mantle and solve the big complex problems which we are richly positioned to solve. I think we're going to be worried, quite seriously, about what do we do with a billion young men around the world with nothing but time on their hands.

Kurzweil: Let me address that, because you bring up a lot of valid issues.

And I've written a lot about promise and peril and the existential promise and peril, so biotech could save our lives but also destroy the world, and we've had the Asilomar Guidelines, which have worked actually quite well. But I had a dialog with Christine Lagarde at the annual IMF conference two weeks ago on just this issue, and I think actually, what we saw last night has been influenced by this. It's much more automation that is causing economic insecurity than China or Mexico and so on.

But this is not the first time in human history that we've done that. I mean, how many jobs circa 1900 exist today? And if I were a prescient futurist in 1900, I'd say, "Okay, a third of you work on farms and a third of you—actually, 37% work on farms and 25% work in factories. But I predict in 100 years, by the year 2000, it'll be 2% on farms and 9% in factories, so it's a reduction overall of 7:1. Eighty-five percent of these jobs will go away." And everyone would go, "Oh my God, we're going to be out of work" and I'd say, "Well, don't worry, we're going to invent new jobs that will replace them." And people say, "Oh really? What new jobs?" And I'd say, "Well, I don't know. They haven't been invented yet." It's not a very good political argument. It leads to economic insecurity. It actually happened though. We've multiplied jobs significantly. Even as a fraction of the population, we've gone from 30% in 1900 to 44% today. The jobs pay 11 times as much in constant dollars per hour—

Kirkpatrick: We the people who are employed?

Kurzweil: Thirty percent of the population was employed 100 years ago. It's 44% today. The jobs pay an average of 11 times as much in constant dollars per hour as compared to a century ago. And we've moved up sort of Maslow's hierarchy because the jobs are at a higher educational level and are more satisfying. So for example, we had 52,000 college students in 1870. So what you do was not a meaningful part of the economic profile 100 years ago. Today there's 20 million, plus five or six million faculty. So that's about 20% of the workforce is either a student or a professor, and that's just higher education. So what are they doing? They're studying poetry and art and science and brain science, something that was unheard of 100 years ago. And there's lots of other examples like that.

But it remains an economic insecurity. We're actually moving in the right direction. Even over the last decade, the percentage of the population working has gone up. The wages have gone up in constant dollars. However, people look and go, "Oh my God, I'm driving a car or truck.

TE16

What's going to happen to my job?" and there's lots of different categories where people have this economic insecurity, and insecurity is in fact a big motivator or de-motivator and has a substantial political influence, and I think that's what we're seeing.

Ming: I've got to say though, 100 years ago there was a fundamentally different problem, which is if I didn't need you on the farm anymore, in six weeks I could train you to work at factory. There's nothing I can do with someone in six weeks that's going to train them for the kind of future jobs that we need. And in fact, what I might train them to do, say something all of us might well believe in is, hey, let's get programming into our school curriculums. That's a pretty shaky promise.

Because I honestly don't think by the time those kids graduate that there will be a job for programmers.

Kurzweil: Well, if you talk to young people today, they're not working or planning to work in farms or factories. They are learning illustration with tools and creating apps for mobile devices and creating websites. And the reality is that the number of jobs is actually moving up, and there's also a lot of new types of economic activity that are not exactly jobs, where people make money with Airbnb or with selling things on eBay or doing work for websites and making money that doesn't register with the economic statistics, and still the statistics move in the right direction.

So the people didn't move from the farms to the factories. When the textile machines in 1800 in England which started the Industrial Revolution occurred, they started the Luddite movement because they felt their jobs would go away, which they did. Employment went up but not necessarily the same people. There were new industries created with whole different types of people and then education increased. That trend is continuing, even though we're now automating mental work.

Bratton: Look, I think the context in which these transformations we're talking about are happening are ones that in and of themselves AI can have an enormously positive role to play at an infrastructural level, not just the augmentation of an individual's intelligence, but the augmentation of systemic intelligence and the ability of infrastructural systems to automate what we call political decision or economic decision. And it's taking place in the context of an increasing accelerating, and what will clearly probably be an even more accelerating ecological precarity. The planetary substrate on top of which this emergent intelligence may in fact appear is one that's—its ongoing-ness is in particular question.

I happen to think that AI has a big role to play in understanding something like what forms of ecological governance may be necessary to sustain the kinds of systems that we want. I think what we saw yesterday was an example of the fact that AI and automation more generally—and I mean not just the automation of labor but the automation of the movement of matter through logistical systems, supply chain systems, and so forth has already destabilized to a

THE 16

certain extent the sense of what it means to be human in ways in which we need to think quite seriously about.

You know, the term anthropocene is one that we hear a lot, and it refers to this notion of a geologic era that is defined by the agency of humans, of a particular species. But it also can be understood—the anthropos of anthropocene can be understood as the agency of a notion of humanism, of the idea that the human experience of human experience is of paramount and central conceptual importance in how it is that we organize our industries and these systems as well. And I think it's something that humanity has a difficult time dealing with. But I think we need to be really—I really would want to see a shift in the discussion around AI precisely to the level of systemic intelligence that may allow for a kind of longer term durability in this way.

I happen to think AI will be what I call a Copernican trauma. Copernican traumas are these sort of moments in history where some sort of way in which we thought we were the central special case, species, whether it's the planet that was the center of the universe, or Darwinian biology was a Copernican trauma, neuroscience is a Copernican trauma of demystification of mind. Queer theory is a Copernican trauma. AI will prove to be a Copernican trauma.

We don't deal with Copernican traumas very well. There's enormous pushback. And I think the humanist pushback against AI, which will be—

Kurzweil: We survived them, though.

Bratton: We have, to date. I don't know whether this is a guarantee. I certainly hope that we do and we will. But you mentioned, quickly, just on the question of design—and it's true, my interest really around AI is what are the implications for design and design disciplines and design thinking in the ways in which we build intelligence into tools at a large scale shift what is designed and designated in this way. And I think it's an important conversation to have around the indirect effects of AI and automation. And driverless cars is an example. You know, if you go to an architecture school—and I have spent a lot of time in architecture schools—the way they deal with the question of driverless cars is not about how do we optimize the sensors on the cars or the decision making systems on the cars or the pathfinding algorithms on the cars. What are the implications for the rest of the city when you don't need to have 20% of your surface area of the urban core just paved asphalt for the storage of transportation units that you don't need to own anymore, or garages or the rest of this as well. Urban planners have been trying to get rid of parking lots as long as there's been parking lots because they're horrible. Turns out the way you do it is you put sensors and intelligence on the cars.

Kurzweil: There's some popular songs against parking lots.

Bratton: There's popular songs against parking lots, yes. And so I think there's ways in which we can think about this sort of systematically. And the other thing we don't talk about so much with driverless cars is how important they would be—coming from Los Angeles, where I spend

THE 16

a lot of time—is how important it would be to provide an access to the city for people for whom it's expensive. You live in certain parts of the city, you can make it to Santa Monica and still pick up your kids in the particular part of the day, but if you don't live near there, it's hopeless. And so one of the ways in which the automation of transportation will allow for an important kind of social shift is that it moves the responsibility for capital ownership from the individual user and putting it back into the system. Transportation becomes infrastructural, and it's the opposite of the cellphone that you talked about, where something that used to be essentially owned by the system then got pushed to the culture of the end user.

Kurzweil: It'll save two million lives a year, which will be useful.

Ming: Yes.

Bratton: At the very least. So anyway, my point is not so much that the questions of augmentation and intelligence aren't significant, not at all. But we can't wait until some threshold point where it somehow appears in order to understand what the terms of that absorption should be and the ways in which we want to design their implications.

Ming: And I want to be clear so that there's no misunderstanding. Five years ago, my son was diagnosed with type I diabetes. I hacked all of his medical devices, broke several federal laws, and built an AI that predicts about an hour into the future whether his blood glucose levels will go high or low. I recently visited Eli Lilly in Indianapolis, an interesting experience for a transgender person, and what was fascinating was their head of research and the CIO took me aside and said, "Hey, we'd like you to talk to this researcher." There was a guy there extending my model on an artificial pancreas, purely in simulations, because the data's hard to come by. But it worked better than a real biological pancreas. It may not be very far in the future before the Paralympics are a lot more interesting than the regular Olympics. [LAUGHTER]

I am a believer in what can be done with these technologies. All I'm calling for, somewhat in complement to you talking about the system, is let's make certain that the human institutions, the cultural institutions are keeping pace with technology. Because unlike a hundred years ago, technology is vastly outpacing the human side, and it takes me 20 years, if you'll forgive the metaphor, 20 years to build a human being, but it only takes us a couple of years to iterate on a new machine learning tool.

Kirkpatrick: And just to be clear—you've used the phrase, and I like it, adaptive creative problem solver. You essentially think we need an entire society filled with people about whom we could use that description.

Ming: I think this is a rational investment, no matter what we're looking at. I mean, I fundamentally believe, when you model this stuff out economically, the impact it would have on the US economy, the South African, the Indian, the three places where I've most detailed, are increased GDP by 10, 60, and 110% if we could really focus on the things that were actually

THE 16

predictive of life outcomes. But at the same time, I think that there is an implication of a pretty profound divide that could evolve over the next 10 to 15 years of people whose labor is worth more than what AI can currently deliver and the increasing number of people who slip below that. And we'll still have jobs for them. I mean, I don't think work is going to disappear. If nothing else, there's always subsistence farming.

Bratton: And this goes to the question you asked me at the beginning about yesterday and there's sort of AI in this as well. I wanted to just sort of first say—because I thought the panel this morning on this subject was very interesting. I think we absolutely need to understand what happened yesterday first of all as a global phenomenon, that the rise of ethnonationalist populism is a global phenomenon. It has global causes, that we see a common psychodemographic divide, whether it's in France or India or Iran, between an urban cosmopolitan constituency and a rural monocultural constituency, and that it's a global position. It's going to have to have in a certain sense a global solution. And I don't think that that global solution necessarily is one that can be solved one eighteenth century jurisdiction at a time.

And so just very quickly, the question of how it is that AI may be a participation in the re-conception and redesign of political infrastructures that are able to in fact intervene in this way, these to me are the conversations that we need to be having and prototyping and understanding in this as well.

Kurzweil: Just to comment on the speed of change, that's certainly happening and that's a corollary, I believe, with the law of accelerating returns. We're greatly speeding up the paradigm shift rate. However, it's remarkable how quickly people adapt to change. You know, phones that you can talk to emerged and people say, "Isn't that amazing, you can actually talk to your phone?" And people say, "Yeah, but it doesn't really work very well." And then a few years later it starts working better and better and then people say, "Isn't that amazing, you can talk to your phone?" and people say, "Yeah, but it's been around for a long time." And then we can't remember when it wasn't the case. An automated pancreas or any kind of solution to type I diabetes or any other condition is adopted very quickly once it's shown that it can work and people then would not even want to think about a few years earlier when it didn't work. We adapt to these technologies very quickly. Think about what life was like three or four or five years ago. If you were to describe the nature of the technology and how we lived our lives, it was very primitive. There weren't social networks just six or seven years ago, and Wikipedia isn't that old, and just imagine life without search engines. That's only about 15 years.

We very quickly adapt to these changes. And it's not going to be a matter of, jeez, there's going to be some people who can't compete with the AIs. The AIs is not a civilization apart. It's not an invasion from Mars of intelligent machines. We create, as you said, these are extensions of ourselves and they are already—I mean, who here, or anywhere in the world, can do their work or get their education without these brain extenders? So we're very intimately merging with them, and these dystopian AI futurist movies where it's the AI versus the humans for control of

THE 16

humankind is not the way it's unfolding. We don't have one or two AIs in the world today. We have two to three billion, depending on how you count, and it's very deeply and intimately integrated with humanity.

Kirkpatrick: Well, this is pretty heavy stuff for after dinner, I realize. [LAUGHTER] I do want to hear from you all.

A: Robert Klitzman from Columbia. That was great. I'm wondering, are there any cognitive functions that you think AI won't be able to cover or do as well as human beings at some point in the near future?

Kurzweil: Well, I think the whole point, when I say by 2029 computers will master human intelligence, that will encompass the different things that humans do. In fact, the cutting edge of being human today is getting the joke, being funny, creating beautiful music. No other species does that. Every human culture ever discovered has done music. Expressing a loving sentiment. Think from our primitive emotions, which come from the old brain. But these things are the cutting edge of human intelligence. Being able to read a novel and have a reaction to it that's coherent and summarize it. These are the types of things I believe computers will do by 2029. But we will integrate with them, as we do now already, and it'll be more intimate. They'll ultimately go inside our bodies and brains and make us smarter, and as I say, funnier. But no, I don't think there's any cognitive abilities that computers are inherently unable to do.

Breckenridge: Ross Breckenridge, Silver Creek Pharmaceuticals. One way of interpreting the last 24 hours is as sort of a revolution against technologies that happened 20 years ago. And, you know, I come from a country where we had a similar revolution in June. What do you think the effect of what you're talking about now is going to be on society in a few years' time? And if you're thinking about—I mean, especially Ray, as part of a commercial company trying to leverage this, what efforts would you have to make to change society before you can introduce these things without causing some huge revolution?

Kurzweil: Well, I agree with the panel that we need to actually think about these things and create social structures and economic structures where people don't feel left out. I think people are in fact not left out. I mean, I can cite lots of statistics that things are getting better. There was a poll taken recently of 24,000 people in 26 countries and they were asked is poverty getting better or worse? And 87% said, incorrectly, that it's getting worse. Only 1% said correctly that it's been reduced by 50% or more, which is what the World Bank reports. Our information about what's wrong with the world is getting exponentially better. So people really think things are getting worse. The world's getting more violent, that's the opinion. You know, if you read Steven Pinker's book, "The Better Angels of Our Nature," this is the most peaceful time in human history, and people say, "What are you crazy, didn't you read about the event yesterday and the day before?"

THE 16

So I do think it's really insecurity about what's happening. By every statistic, there are more people employed, they're being paid more, the jobs are more interesting. But there's this deep economic insecurity because people—you know, a hundred years ago actually, people weren't that aware of, "Gee, there's these new machines coming that are going to possibly displace me." Now people are very aware of that and I think we have to provide some means by which people have the right economic security, and that brings in lots of social and economic and political issues, which we could debate. But I agree with the panel that this is a key issue.

Kirkpatrick: I still can't help resonating with the word 'dignity' that came up in the opening session. I think that has to be somehow reinserted into the debate about what people need in terms of all of these issues.

Rotenberg: Marc Rotenberg. I've spent quite a bit of time looking at ethical and social issues related to new technology. My question for you, Ray, assuming that we do have techniques to significantly augment human intelligence and to enhance cognition and emotion, it's a practical question I guess, and that is what is the business model that sustains this, that gives people access to this opportunity? And I ask the question in part because what I've observed with the digital economy over the past 25 years is that increasingly we no longer possess digital artifacts. I mean, as you say, they're held in the cloud. We don't necessarily own them. Oftentimes, we lease them or we subscribe to them, and every incentive seems to be to extract from us more value. So to the extent that I have a bit of a dystopic view of this, it's not that it can't be made to work. It's that in the simple instrumentation of the business model, I am concerned about what happens when someone can't pay for the next six months of their enhanced cognition.

Kurzweil: Well, another one of these I think incorrect perceptions is the growing digital divide. I mean, every statistic shows that's absolutely not the case. The Internet and access to it is spreading like wildfire in Africa and other developing countries and we now have two to three billion smart devices that are connected to the cloud and it's going to be six billion in a few years. What's driving this, one of the implications of the law of accelerating returns is a 50% deflation rate in the capability of information technologies. So I can get the same computation, communication, genetic sequencing, and lots of other information technologies that I could get a year ago for half the price today. And we put some of that price performance improvement into price so prices come down, and some of it into performance so performance goes up. That's why you can buy an iPhone or an Android phone that's twice as good as the one two years ago for half the price, which is a fourfold increase in price performance, which is 50% per year. And ultimately, these are very inexpensive, and we have lots of resources that are free, and I've argued this with Christine Lagarde, that we don't count that in the productivity statistics. So when I was a teenager and spent thousands of dollars to buy an Encyclopedia Britannica, I was thrilled with that, but that was a few thousand dollars of economic activity. Today, I've got one that's better and free and doesn't count for any economic activity. And these resources are in fact spreading wildly and becoming ultimately very inexpensive. So she

THE 16

said, well, you can't eat, you can't live in, you wear information technology, and that's going to change. 3-D printing, we'll be able to print our clothing for pennies per pound by 2020. We'll be able to, in the 2020s, print out modules that you can snap together and build a house. That was already demonstrated in Asia as an experiment, but that'll be mainstream in the 2020s. There'll be a new agricultural revolution with AI-controlled food production that'll be very inexpensive. And it's subject to this 50% deflation rate. So ultimately, it's very widespread.

Ming: There's an implicit assumption behind that that the simple connectivity—

Bratton: Or two or three—

Ming: Probably many. But simple connectivity is the thing that will change it because everyone owns a smartphone and suddenly humanity is transformed. But I don't see any strong evidence to support that. I mean, Andrew Ng, I've done some work with him in the computational neuroscience world. He and Daphne Koller, amazing people, found Coursera, wonderful solution, frankly, to *their* problem, which is if I, amazing genius at age ten had been trapped alone on an island, what would I have needed to transform my life? Well, access to all the information in the world, the best classes the world has to offer. But the people that show up at Coursera are already educated. There's this idea that simply passive connectivity will fundamentally change the world. I simply have never seen evidence of that being born out. If we don't go actively out and change people—

Kurzweil: I mean, I was on the MIT board for eight years and we put out all of the MIT courseware, including videos and so on, everything online, and there were 10,000 schools sprung up in Africa—this was at an early stage, where they would sort of gather around a computer with kind of a shaky connection to the Internet and take courses from the best minds in the world for free. And this is much richer now. You know, from TED videos to every other kind of technology. I meet people who are real experts now on things like deep learning and many other fields who learned it online, and this is just beginning, this revolution. So it's not correct to say that people don't have access to this.

Ming: I didn't say—of course they have—that's what I just said is they have access.

Bratton: It's just that it doesn't matter—

Ming: It's just it doesn't matter if they don't have a reason to believe that the hard work of investing their life into that access will pay off.

Bratton: I think one of the things that we're going to see over the next few years—and I would say that I also think that the futures that are going to emerge are far more conditional than they generally get described. I'm uncomfortable about talking about things in terms of will and so forth, as it's a matter of a kind of disclosure of a revelation. But one thing I would say is that I think over the next few years, one of the things that we're going to see in terms of AI and its application at these levels, at least in the initial phases, is a certain kind of AI

THE 16

skeuomorphism, where we are applying AI to systems or processes that are sort of recognizable and have a particular kind of social or economic value at a point in time, and we feel that if we can accelerate and automate those things, for example, all the things that we sort of described, that the actual effects of those things would somehow remain the same. And they won't. By applying AI to a 1980s or 1990s or even early 21st century economic and sociocultural system, that all of the advantages and effects of those would simply be exponentially more efficacious is, I think in terms of the history of any kind of transformative technologies, extremely unlikely.

And I don't think there's a dystopian sort of thing. I just wanted to say that. It just means that the challenges, the work of imagination are in front of us, not behind us.

Kirkpatrick: Well, what I was going to say is I don't think anyone else in this room could say that they've been moderating since 8:30 this morning, but I have and I am pretty much ready to stop. [LAUGHTER] And I've got to say thank you all. I thought this was incredibly rich and incredibly interesting conversation.

[END]