

Can We Handle the Coming Convergence of Man and Machine?

Speakers:

Benjamin H. Bratton, University of California, San Diego

Mary Lou Jepsen, Openwater

Tessa Lau, Savioke

Rohit Prasad, Amazon

Moderator:

David Kirkpatrick, Chief Technomist

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

Kirkpatrick: Let me quickly introduce you to our four panelists here. From the far end, Benjamin Bratton is a professor in the art department in the University of California at San Diego. But he's also trained in both sociology and computer science. And what was the other thing you do that's—

Bratton: Philosophy.

Kirkpatrick: Philosophy, you teach that? He does a lot of things; he's quite a Renaissance person.

Mary Lou Jepsen is also a Renaissance person, who was at one point named one of the 100 most influential people in the world by *Time* magazine. A longtime friend of mine who I first met when she was the chief technology officer and co-founder of the One Laptop per Child initiative with Nicholas Negroponte, which really had a catalytic effect on the devices that we have in the world today. She's got some interesting things to say about her current project, which is called Openwater.

Tessa Lau is the chief robot whisper and chief technology officer at a company called Savioke, which builds robots for delivery.

Finally, Rohit Prasad is the head scientist for Alexa Machine Learning at Amazon and a vice president of Amazon. Obviously Alexa has had a catalytic effect on everything. In fact, two

years ago, those of you who were here will recall, we had an Alexa on stage, just because we thought it was so cool; we just wanted to see how we could use it on stage. We didn't think of too many good ways at the time, but at least we were telling the world, "This is a fundamentally new change in technology."

Maybe we should start with you, Rohit. When you think about what you're doing with Alexa, talk about it from the biggest picture standpoint, what is going on with this new interface and connection to the cloud and AI deployment?

Prasad: So fundamentally, what's going on is we're trying to improve daily convenience. With everything you think about—like, we all grew up, a lot of us here, grew up with *Star Trek*—when that was airing. And you'll remember the convenience of being able to talk to a computer without any touch or looking at it. It was just so liberating as a concept. So we brought that to fruition with Echo. And the goal is, indeed, to make everything be easy by the most natural means of communicating, which is voice, which is what we are born with when we start off in our lives. It is the most easy, convenient way. And if you had a computer that was built into everything and was available everywhere: you can play your favorite tune, you can reorder batteries for instance, you can set up a meeting—all of this just by voice, is extremely liberating. What's also happening with the deluge of devices and services that are connected now, that if you are picking up your smartphone and trying to find how to turn on the light or check on your garage light, that has too many steps. Alexa removes that from you. So ultimately, it's all about daily convenience so that you can do more with your time.

Kirkpatrick: Do you think of this as being sort of the interface of the future? Is more and more of our intersection with computing going to happen via voice, in your view and in the view of Amazon?

Prasad: Definitely a lot will happen with voice. But it's not going to replace everything we know as it exists, as well. I mean, there's a place for touch, there's a place for screen—and there will be many other haptic feedback mechanisms in terms of interfaces. But voice will find more and more actions that we want to take in our daily lives. [Many things] will be done by voice and that's what we are after right now.

Kirkpatrick: When we talk, you know, in the session title about the convergence of man and machine, people and machines, we are accused by some of sexism, but we didn't intend that. Is that where you see us going? Is the world kind of reaching a point where software and systems and human intelligence are getting a little conflated? You know, long term?

Prasad: I wouldn't say they are getting conflated. They all have different strengths. If you think about machines, they are better at computing. If you asked a machine to compute a factorial, it's way more easy for a machine to do it than humans. If you asked a machine to search across a massive catalogue of music or video, again, for machines it's easier. I think what we are going to find more—and this is beyond Alexa—we'll find more and more

collaborative tools with AI and machine learning in general, where you will be able to do more with machines and humans working together.

Kirkpatrick: Okay. We'll get back to some implications of that. But thank you so much. So let's just go down the row here. Tessa, how do you whisper to robots?

Lau: I think we all grew up in a world where robots are the future, right? We all grew up with the science fiction and the stories and the movies, and we all want to make robots happen. And the way I'm helping to do that is by birthing a new generation of service-delivery robots. And I happen to whisper to them in order to get them to do their job. Robots are fairly complex things, and so one of the things that I am passionate about is how do we make them function in today's society and do things that are productive and help people, in order to make their way in the world?

Kirkpatrick: So talk a little bit more exactly about what your robots are doing right now.

Lau: We started Savioke with a vision of getting robots out into the world. You know, we all have this dream about a future in which robot assistants are everywhere. But it turns out that robotics is pretty hard, and so one step at a time. The first step that we're taking is to get robots out into hotels doing room-service delivery. So a Savioke Relay will park at the front desk and he'll bring your food, your dinner, your toothbrush, directly up to your room. He'll take the elevator; he'll push the buttons using Wi-Fi, walk down to your room; he'll tell you when he's there by calling your phone, and then he'll bring you your item. And this is changing the way hotels are providing guest services inside their properties.

Kirkpatrick: He doesn't speak, right?

Lau: So Relay deliberately does not speak and it was because four years ago, when we started, Alexa didn't exist, and, you know, that kind of technology was not mature enough.

Kirkpatrick: Let me ask you this: why is Relay male? And why is Alexa female? And should they get married?

[LAUGHTER]

I mean, I'm just curious. Is it connected to the fact that he doesn't speak? It's very interesting that you're describing him as a male. I'm just curious.

[LAUGHTER]

Lau: Every one of our hotels actually gives their Relay their own name. One of the hotels that we just brought up in Las Vegas has an Elvis and a Priscilla.

Kirkpatrick: Okay, so he's, you know, can be—

Lau: Androgynous.

Kirkpatrick: Yeah, androgynous robot.

Lau: Hermaphrodite.

Kirkpatrick: We'll get back to the gender of Alexa in a little while. I know you have interesting things to say about some of the things you've learned, but let's get to that, maybe, as we go forward in the discussion. Because I think it's very interesting the challenges you've encountered.

So Mary Lou, talk about what you're doing now. Just start out with that, because it's so kind of mind-blowing.

Jepsen: I was at Facebook and realized as I was working on VR that there was something much bigger I could do. I figured out a way, a discontinuity in Moore's law in the optics, that enables a wearable with MRI resolution that can transform healthcare costs. If you think of it, 25 million people a year die of clogged arteries and heart disease and cancer, combined. It also enables telepathy communication with human thought alone, to transcend language.

Kirkpatrick: Just wait, she said, "Communication with human thought alone." I just want to make sure that came through.

Jepsen: Basically, a system that can see with very high resolution inside of our body. If I threw you in an MRI machine for an hour, I can tell if you if you've got a tumor. If I throw you in for 10 to 100 hours, I can tell you what words you're about to say, what images are in your head, whether you're in love or not, and so forth.

Kirkpatrick: This is in a traditional MRI machine, you mean?

Jepsen: Yes. Yes. And so actually, this summer we were just trying to match the resolution of MRI. We just got the funding, we got the team up and going, and just bought ourselves a year to work on the basic physics. We're doing this with camera chips and LCDs, highly modified, I know a lot about making LCDs and camera chips in consumer electronics. So I've shipped billions of dollars' worth of stuff on the hairy edge of physics. But this discontinuity in Moore's law enables us to actually get a billion times higher resolution.

Kirkpatrick: You're using infrared light, not magnetic resonance?

Jepsen: Right, infrared light. Your body's translucent to near infrared light. Like if we were all sitting here bathed in infrared light and you guys could see infrared light, we would look like kind of the consistency of glasses of lemonade with blood in it. Because we're 3 percent blood. And everybody thinks scattering is random; it's not. The physics show it's deterministic and reversible; if you can make pixel sizes on the order of the wavelength of light and modulate the waves in the wavelength of light, using all kinds of four- and five-syllable words in physics, like—

Kirkpatrick: In other words, you can tell what is being shown by the light with incredible exactness?

Jepsen: Yeah. We can focus rather than to a cubic millimeter, to a cubic micron. Which sounds like 1 thousand times resolution higher, but it's an X, Y, and Z, so that's 1 thousand times 1 thousand times 1 thousand which is—

Kirkpatrick: More than MRI—

Jepsen: —two commas, it's a billion times higher, is what we're doing in the lab today. Not product, but we're trying now to decide what products to develop. Like do we try to cure cancer and catch it at Stage 1? Do we go straight to telepathy? Or do you take the largest healthcare expense in the world by far, brain disease, and attack—try to get a handle on that? Two billion people globally live with brain disease; one billion of them can't hold jobs. And so we're working on all of those areas right now as we building up the devices.

Kirkpatrick: So a billion people on the planet cannot hold a job because of the nature of their brain disease?

Jepsen: Between like Alzheimer's, Parkinson's, clinical depression, schizophrenia; those are sort of the top ones. But there's ones below that.

Kirkpatrick: Wow. That's a big project to solve.

Jepsen: Yeah.

Kirkpatrick: Okay, bigger picture question—sort of the same thing I asked Rohit: do you see us headed towards a convergence of people and computing?

Jepsen: Yeah, I think time is going to seem so slow, as we're able—I mean, it's so cool to do the voice, now. And like, can we get the thoughts out there? Like can you imagine a director being able to dump a rough cut for a movie or an architect who figures out a new building, just to be able to like skip the CAD and just go straight out? As we're able to jack in to the complete thoughts, rather than how we type them out or—we basically get stuff out of our brain now by moving our mouths or typing. Or using Echo.

Kirkpatrick: So how soon would we get to that Matrix-like world that you just suggested?

Jepsen: I do consumer electronics, so our first products are in a few—like two, three years.

Kirkpatrick: But we won't be able to dump our brain with an idea for a fully formed movie in two or three years, right? I mean—or will we?

Jepsen: Well, okay, so before we—so there's this deal of responsibility that I think you talked of in the start. And one thing is that the misuse of that is tremendous. And so we've agreed to make a product that will only work consensually. And we'll only ship it when we can define what

it means to be responsible shipping it, because of the neuroethics, the ethical and legal implications. We make the ski hat with LCDs that can read your thoughts. Can the police make you wear it? Can the military make you wear it? Can your parents make you wear it when you come home at 3:00 in the morning and they want to know where you were? Can your spouse make you wear it? You know, like these are—like we think it can only work when the individual wants it to. So there's a bunch of issues like that to be responsible to—

Kirkpatrick: Well there's so many things we can follow up on, but one thing, before we move on to Benjamin, you were talking, when we talked on the phone, about the challenge of keeping these technologies kind of gated. And we were talking about China in particular, that you currently are letting no technology out of your lab. But presumably other people are working on the same thing all over the world, in some form, right? There must be a few.

Jepsen: Ah—[SIGHS]—

Kirkpatrick: No? You don't think so? Good, okay.

Jepsen: No, I think the talent, like this—I mean, I usually do things everybody thinks are stark raving and get them to ship in a few years. But if somebody else was doing it, I assure you, I wouldn't—

Kirkpatrick: So you don't think anybody's doing it.

Jepsen: No.

Kirkpatrick: Okay, but do you think if you ship the product, that somebody else would start doing it very fast?

Jepsen: Reverse engineering; I mean, the sincerest form of flattery, when we have our first developer kits out. And so we need to be ready when we have the developer kits out with our security and—

Kirkpatrick: I also think you said you wouldn't make it in China? Is that a fair statement?

Jepsen: Well, I mean, most all consumer electronics are assembled in China and the software is flashed in China, on every single cellphone, every single cellphone in the world. And so there's already security issues. But we haven't decided where our manufacturing is going to be done yet. We're doing a lot of development in San Francisco right now, in our labs, and mostly with partners, actually, at various places mostly in this country. We're not sure where we're going to ship or where it's going to be so-called 'made' in.

Kirkpatrick: Okay.

Jepsen: But it's certainly a concern, for sure.

Kirkpatrick: Okay, that's sort of a side question, but I'm glad you mentioned that.

Jepsen: But again, we have to decide what it means to be responsible in that. And I think there's a big issue right now, even with cellphones being flashed. All cellphones are flashed—

Kirkpatrick: Flashed? What do you mean by flashed? What does that word mean?

Jepsen: The software is put on in China. As you assemble the motherboard, and the screen, and the, you know, the components, the titanium back or whatever you've got, the software is flashed there with code that can send it through a debug cycle that tends to last—it depends which phone—24 or 48 hours, but—I mean, just think about that: all software—there's software put on every single cellphone in the world in China.

Kirkpatrick: Okay. We'll get to that topic in separate sessions later.

Benjamin, okay. Any response to anything you've heard? Or just talk about how you look at this issue of AI and machine versus human intelligence and what constitutes intelligence? I know these things are things you really think about and have written about—

Bratton: Maybe I'll try to answer it on a couple of different registers, because—in response to what's been said. I think—one of the things I work on a lot is the development of what I think are perhaps better models of what we mean by AI, what AI is and what it actually isn't. Because I think, in many respects, the more difficult problems for the integration of AI into society and culture and politics as a whole, are ones that have to do with the mismatch of the presumptions and expectations of what AI is and what it can do, than they are core technological problems in and of themselves.

Secondly, perhaps as we were discussing before, some of the geopolitical issues involved with this as well. And so for the first one, I would say a lot of my work is looking at non-anthropomorphic and non-anthropocentric models of AI. I think going back to the Turing test, the idea that we will think of AI as being intelligent if it seems to perform thinking the way that we think that we think. We will grant that this is, in fact, intelligent. And I think that in truth, the types of AI that we are the most successful, deep learning systems, machine learning systems, are ones that are much more dependent on how they sense the world, the input that they have around the world. And the form of thinking is one in which sensing and cognition, for example, are much more blurring together. And I think over time, we will want to see AI as part of a much broader spectrum of forms of material intelligence, of which our—and Mary Lou's work I think, shows quite clearly—our own thinking is really deeply material at its basis. Rather than something different.

The question of the geopolitics of this is also quite interesting. I think, as we were talking, I think one of the things we're going to be seeing over the next coming years, is an increasing fortification of what I call hemispherical stacks; that we have a Chinese stack, a North Atlantic stack, a Russian stack—

Kirkpatrick: You mean a stack like the internet, Facebook—

Bratton: All of that, right.

Kirkpatrick: —NSA security, kind of thing?

Bratton: All of this as well. And that the distinction we may have between the geopolitical domain that is structured and essentially enforces itself, its own sovereignty, through its ability to design and monitor and to cohere those technologies, will be more structured. And the AIs that we'll be developing will be ones that evolve inside those encapsulated domains—

Kirkpatrick: Oh, I see! The point is that these geographically discrete stacks of software and connectivity—

Bratton: Right.

Kirkpatrick: —and infrastructure—

Bratton: Right.

Kirkpatrick: —will engender divergent forms of AI, in effect. Is that—

Bratton: That's right. They will be essentially localized within those kinds of domains—

Kirkpatrick: Wow.

Bratton: And that the forms of evolution that they're able to bring forth will be ones that will be specified within this as well.

Kirkpatrick: You know, Rohit, one of the interesting things about the way you defined what you're doing is you didn't mention AI at all, I don't think.

Prasad: Yeah, it's pervasive across everything Alexa does. In some ways it's the ultimate intelligence, because when you're having a conversation, you're trying to accomplish conveying a concept and having an actual discussion on it, it's a super hard problem. It is the way, what makes us as humans be intelligent, and Alexa's one of those AIs where you can't even, in this kind of a setup, you can't tell what the next goal the customer or the user has. So the goal is always evolving. You don't really know what state you are in, so it's a very hard problem. It is an AI as well, but the kind of relationship people have with Alexa also makes it like tell you that it's now making AI a reality in a consumer domain, which is always great to see.

Kirkpatrick: Yeah, you sure have done a good job with that. But tell us about that experiment you have underway with some universities. I forget what you call it.

Prasad: Yeah, it's called Alexa Prize. It is an ultimate challenge; it's like a grand challenge where the universities, academia, is being challenged to come and build what we call social bots that can converse with you on Alexa or an Echo device, for 20 minutes. It's a pretty daunting target—

Kirkpatrick: Like have a real conversation?

Prasad: Have a real conversation on societal topics that you read about in newspapers. So it's meant for that. We're in the first year of that competition. It will take several years to beat that barrier. If you imagine a TED Talk is like 18 minutes, when you're really having a one-way conversation. So a two-way conversation to last 20 minutes is an incredibly hard target to beat. And we're working together with universities, the two things we are doing unique, which is very hard to do with academia, is academia usually doesn't have access to this kind of data. Millions and millions of customers are interacting with these social bots. And second is compute. We gave eight of these credits for these universities to try their best algorithms on scale on the data and the feedback from the customers. So these are the two unique dimensions. But the grand challenge here is to actually have a 20-minute conversation.

Kirkpatrick: And people can already see how that's going by doing what with their Alexa?

Prasad: Yeah, you just say, "Alexa, let's chat." And you will get one of these social bots that we are right now down to the finalists, where there are only three finalists remaining. When we sent out the application process, hundreds of universities across the world applied. We selected 18 teams; 12 were funded, six participated unfunded. And now we are down to three teams that are in the final stages—

Kirkpatrick: And if you say, "Let's chat," you'll get a random selection of the three?

Prasad: Random selection of the bot, yeah.

Kirkpatrick: And then that data will all be factored into the competition in effect.

Prasad: Yes. So we've already chosen, through customer feedback and one wild card through Amazon, the three finalists. And we'll be hosting the finalists shortly. And then at [AWS] re:Invent, which is at the end of this month, we'll be announcing the finalists.

Kirkpatrick: So in other words, how long the conversation lasts is part of the result of the competition?

Prasad: Exactly, the length is a metric, barometer. But also how coherent and engaging these AIs are is also important. So it's very hard to continue a conversation if it's not coherent and engaging.

Kirkpatrick: Okay, but is that another way of saying how it feels to the homeowner, the person, like another person?

Prasad: Or at least an intelligent entity, right? I wouldn't go as far as a person, given where we are in terms of the state of the artificial intelligence. But it is, as you could imagine, is it getting human-like? Definitely.

Kirkpatrick: I want to go back to Tessa. The thing that's interesting about Alexa's extraordinary market success is that maybe you've even stumbled into this opportunity to really advance AI more rapidly, because suddenly, you've got this population that's really there just conducting full-time, day-long experiments all the time, right? So you're just moving forward faster than we've ever been able to, in such a system, in a way.

Prasad: Absolutely. I think having an AI that is affecting millions of customers, the amount of feedback we are able to get, and all this in a subtle fashion where if Alexa plays a song that you like, you will continue playing it. If it plays a song that wasn't the one you wanted, you're likely to barge in and say, "Alexa, no, no. I don't want that one." These are all signals that make us improve Alexa very fast.

Kirkpatrick: Okay, wow. So Tessa, talk about some of the things you've learned about your robots? And first of all, is AI specifically part of how they work, also?

Lau: Oh absolutely. I mean, there's a long history of AI robotics and so, you know, you wouldn't have robotics without AI and vice versa, because I think their synergistic and they interplay well with each other. A lot of what we do in robotics started out in the AI community where people didn't know, is it possible to make a machine know how to get from point A to point B? And now we have GPS and now we have robots, right? So there's a long tradition of AI and robotics living together. What we're doing right now, I think, is sort of the first step. And there's a lot more complexity down the road that you can get to once robotics is maturing and is getting more widely spread. What we're doing right now is really just scratching the surface of what really can be done with robotics and AI.

Kirkpatrick: But talk about what you're learning about the interaction between people and robots. Because there's some interesting situations that you've been experimenting with.

Lau: Yeah, well, similar to how we have a lot of experience with Alexa and talking to people, we're getting a lot of experience with robots interacting with people. We've done almost 200,000 deliveries to guests in hotels. That is many robot hours of time in human-facing environments. And we're learning from all of that interaction. And one of the interesting things that I've realized is that as we're collecting all of this data, we're designing the systems as we go. And we're designing the systems to cope with the kinds of situations that come up in day-to-day usage. And so, you know, the way that you treat that robot in that hotel, actually will influence how it behaves in the future. Because we're programming it to deal with whatever happens to it. If you're nice to your robot and are enjoying its company and we see that. That's something that guests like in a hotel; we'll program in more interactions like that. We'll make it more social, more engaging. If people start abusing our robots, then we're going to program in self-defense into our robots.

Kirkpatrick: And that has happened, right?

Lau: And that *has* happened. I would love to say that everyone who has ever stayed in one of our hotels is sober and well-behaved. That is not the case.

[LAUGHTER]

Kirkpatrick: So what have you done to give the robots self-defense capabilities?

Lau: I would love to say that we put a Taser in them, but they wouldn't let us.

[LAUGHTER]

So you know what we do right now is, since robotics is so young, as a field, and getting robots out there into the world is so new, our first step has been to make our robots vulnerable, and polite, and well-meaning. But mostly, we want people to think of our robots as needing protection and needing help. And that's its first line of self-defense, honestly. Because if it's got those big blinky eyes and it makes you think, "Aww," then you're going to help it. You're going to not try to abuse it and if you see someone else abusing it, then you're more likely to empathize with the robot. And honestly, that's how we're trying to give our robots an ability to defend themselves.

Kirkpatrick: Now, Benjamin, I have to just note that a number of the things that we've just heard at this end of the panel could be said to be anthropomorphizing kinds of things, which you were arguing against before. So any thoughts on that?

Bratton: Well I would say that the difference—what I'm talking about in terms of anthropomorphizing and embodiment of artificial intelligence are—I want to make this distinction—two different things. I mean, to me, Alexa is a robot. It's just a robot that's optimized for one kind of sensory capacity, which is, you know, hearing. And it's optimized for understanding how bipedal hominids vocalize at it and to understand that hearing. But it could be used for lots of different kinds of hearing. It's not ambulatory. It's not embodied in the world in the same sort of way. And so the way in which it can learn about its world through its sensory apparatus, and then learn how to think about its world, through its sensory apparatus, is differentiated. So I see these as two kind of species that are sort of evolving for different niches and different ecologies and will be relatively successful or unsuccessful accordingly in this way or another.

The question, though, this is a different point I would make, between how we interact with it. And to me, I would also say that in addition to that, some of the things that you've been talking about, Alexa, these are interfaces, right? This is a form H/AI ID, human/artificial intelligence interaction design, as much as it of the actual AI itself. And I would be wondering if you guys have done—answered this question of sort of—as I brought up—sort of the folk ontologies about AI that we have, that we wrestle with. If you guys have done any work about how your users think Alexa is working? How they imagine it to be working?

Kirkpatrick: Ah-ha! Interesting question.

Bratton: What happens when they say, “What’s going on in there when I say ‘reorder something?’” What do they imagine is going on? And how far is that from the reality? And whether or not you think that gap matters? And if so, why? Maybe a way of thinking. I mean, it’s a way of phrasing this question, like: what some of the dangers of anthropomorphic illusions might be? Or what their power might be.

Kirkpatrick: Do you want to quickly take that?

Prasad: It’s still an early problem. There are also a lot of things, which are different, between transactional activities versus when people are saying, “I’m depressed. I want to commit suicide.” These are different thoughts where we have gone into looking more deeply into what Alexa’s response should be. But it’s still early, I would say, in terms of what—the capabilities are expanding and we definitely have the opportunity of learning more in that dimension. But we have given a lot of thought there in terms of also, how Alexa responds. And in design, by making sure there are no misconceptions there.

Kirkpatrick: I want to get your thoughts, Mary Lou, on any of this, but before we go off of you, Rohit, talk quickly about what you’ve learned about how [those with autism] have responded. Because that was so interesting.

Prasad: Yeah, there’s two—this has been a surprise with Alexa—the two places where it’s been really humbling in seeing how Alexa has been adopted. One has been in the mobility community, where people are disabled or can’t move or aren’t mobile enough and it has really helped in that setting where you don’t have to move your arms, hands, and you can just speak to Alexa and get entertained or get information, turn on your smart home lights, where we know customers who have outfitted in their parents’ house all the lights with the connectivity of the hub so that now you can just have Alexa issue commands—fulfill commands that you are issuing.

And the second has been—this has been on my friends who have autistic kids, and they have felt like how the kids have been interacting with Alexa has built their own confidence. And that’s something that I didn’t think of when we were building Alexa.

Kirkpatrick: So really unexpected, even medical, implications. Unexpected.

But Mary Lou, any thoughts? And among many things, I want to know whether you are going to be building an emergent AI with your system as well? But keep going.

Jepsen: I think AI is pretty limiting. I mean it’s religious math; it’s a program. It’s over-blown right now. I say this as a former student of Marvin Minsky, one of the founders of AI. But at the time AI was starting, when it was coined in the ‘50s and the ‘60s, Doug Engelbart and a bunch of people said, “Well I don’t want AI, I want,”—and they thought it was cute, they inverted it, “IA.” So what if we make our brains smarter instead? We’re more creative. We’ve got 100 billion neurons, each with 100,000 different connections. Can we swirl our minds around with each other? We’re going to deal with the trust. We’re going to have to deal with—apparently we

hear 200 different lies a day. I'm not saying anybody's lying, but like, you know, maybe I look fat in this dress or whatever. You know, like the truth is something—then that's not the worst of it. I think the person in the White House probably is unfiltered but otherwise, we tend to filter—part of social lubricant.

So what happens when we can—you know, you can learn all I know about optics and electronics, and I can know all you know about AI. What would happen if we would collaborate and transcend language with these parts of our brain? We should be able to do a lot more, a lot faster. Are there all kinds of privacy and legal issues? Could you imagine Human Resources in that meeting where we're sharing our minds? I mean—

[LAUGHTER]

Kirkpatrick: When you really fully believe that kind of problem is going to be before us in a small number of years?

Jepsen: Yeah, in less than ten years. I really firmly—it's inevitable. The physics are inevitable, whether we execute or somebody—and right now, when you have the national academies of science of most every developed country in the world saying, "Of the top five things you can do as a technologist, reverse engineering the brain is on that list," so there are many approaches. Not like ours, so a lot of different approaches doing brain-computer interfaces and so forth, trying to crack this problem. A *ton* of funding going into it. We can make much more progress using the consumer electronics manufacture and supply chain of Asia as is, with the discontinuity of Moore's law. But, you know, I guess the fundamental thing is how do we do this responsibly? Which is why I decided to spin out and do a start-up—

Kirkpatrick: Because you were at Facebook before doing some similar work?

Jepsen: Well no. I wasn't doing this at Facebook. And in fact, you know—

Kirkpatrick: But Facebook has talked about doing this kind of thing themselves.

Jepsen: Facebook has talked about it, but Mark wouldn't let me do it. Which is one of the reasons I spun out. Mark Zuckerberg wouldn't let me do it. But like actually, I went to MIT to do this in 2005. And I got distracted with Nicholas and we started One Laptop Per Child because I thought I could get it done faster. Then I pitched this to Sergey Brin and he acquired my company in—I don't know, 2012. So I went and helped start Google X. And I thought, you know, "Oh I come into Google X, I want to do this brain-computer interface," and Sergey said, "No! You have to do this other stuff first." Because it was more valuable to Google. So I've been wanting to do this for a long time. I think a lot of people have been wanting to do it for a long time.

So I forget the question. Oh, AI realized. So what happened is that we can amplify human excellence. Like, how do we get better at what we do? How do we see how we get better? How do we monitor that? I mean you've got like programs like "Couch Potato to 10k in Three

Months.” What does that mean for our ability to learn? A lot of us think that we can’t, particularly in science and technology, where if you miss one key thing in math, you feel like—sort of like you’re blocked, almost, from it. So we can fix these things.

Kirkpatrick: Okay, I definitely want to hear from people in the audience, and comments are allowed, as well as questions. But Benjamin’s been making some interesting facial expressions, so I want to hear what you have to say.

Bratton: As is my wont.

Kirkpatrick: Go ahead.

Bratton: I tend to do. I wanted to defend language for a moment in this process, if I could, and do it without language.

I would want to suggest that perhaps the question of augmenting intelligence, and the augmentation of intelligence, should be something, if we think about this as the function of AI, rather than AI in a Petri dish that’s sort of all by itself, intelligence exists in the world around us at many different scales and many different tempos, or across the phyla in many different ways. And if we’re thinking about AI as something that would amplify and augment those, it should be something more broadly conceived.

I would also insert into the discussion—this goes, again, to maybe the question of what do we mean by intelligence? I’m suspicious of the notion that—and perhaps it’s because I haven’t heard it—that our thoughts really are so fully formed before we articulate them, through language, through drawing, through inscription, through typing, through writing—I have a sense that, from my own work, that it’s in fact the actual act of sort of collapsing the field condition into—

Kirkpatrick: Expressing them is what creates them. Yeah.

Bratton: —into the particle of semantic syntax, that is part of the process. And that even if we were able to sense a thought at absolutely, super-low resolution that the thought that would be captured and expressed in the form of a CAD program or something, would be quite noisy and fuzzy, because thoughts are noisy and fuzzy and they require this process. And so perhaps the convergences that we’re talking about are ones that are external as well as internal in this—

Kirkpatrick: Interesting. What do you say? What do you say to that?

Bratton: I just wanted to ask—

Jepsen: Oh it’s a function of speech, too, you have to—you know, the near-sighted Mesopotamians were the special people that etched the cuneiforms to like invent writing back thousands of years ago, and they needed to be near-sighted because it was hard to etch. So you know, that was super key for the development of writing, and then we got away from tablets and got ballpoint pens from NASA and so forth. Sure, so writing got easier. But what if

you had a visual idea and could just dump? What if you could get the layers of music out of your head? It's very hard to do that by repetitive hitting something like at the right exact tempo. What if you could just get the rough cut out onto your computer? What if you've got writer's block, why? What are the converging—or diverging ideas going on? And then, you know, a final thought: We're not the only thing with brains on this planet—which is just to follow up on your thread. Octopuses have neurons all over their bodies. They might be much better able to solve certain kinds of neural nets or problems or whatever kinds of problems, because they're wired—

Kirkpatrick: Oh and we'll be able to figure that out, down the road, using your same technology in some possible world.

Jepsen: Or just like the ones with the great smell that can smell land mines or cancer on you. What if they could tell you? You know, your dog could tell you they thought you had cancer or you have a disease.

Kirkpatrick: Yeah, so communicating with animals is also one of the implications, potentially down the road of such technology?

Jepsen: Right. We might stop eating them and start collaborating with them.

Kirkpatrick: That's interesting, but let me just ask you. Tell me if I'm channeling you correctly in response to what he said about the expression and the thought.

Jepsen: Oh sure.

Kirkpatrick: Because isn't it reasonable to think—and maybe Benjamin, also—that if we had the capability that you're describing, Mary Lou, that in itself could become a new form of expression itself—

Jepsen: Yeah, that's the point. Yes.

Kirkpatrick: —therefore it would just become yet another one of those processes by which our thought is expressed.

Jepsen: Right.

Kirkpatrick: Now it might be noisy. Who knows?

Jepsen: If, when—I mean, we get better at it.

Bratton: It's a kind of language.

Jepsen: It's not a language. It's okay, if you say visual and music are languages, sure. But right now all I can get out of my head is by moving my mouth, perhaps dancing, and typing my fingers. If I draw it's pretty slow. And if I could just get the image out of my head, it would be *so* much easier. Even a rough cut for me to develop.

Bratton: I'm all for the AI cuttlefish. The whole thing's worth it for that.

Prasad: I think there will be some language needed to even convey that concept. I'm with Mary in terms of how it's going to get manifested, but the expression requires conveying your thought in a certain form. And how that will be communicated across different species, a human machine, there will have need to be a certain media. I think language will play its part. Many other modalities will play its part as well, because you do need to convey that concept in one form or another.

Kirkpatrick: But as a computer scientist working in AI, when you hear Mary Lou say what she thinks is possible, do you pretty much think that she's in the right ballpark?

Prasad: Yes. I think that in my former life, in terms of in a different company, where we were doing a lot of DARPA work and we are looking at brain signals for ways for communicating. So it will happen. And all of the other things Mary mentioned is also going to become supremely important when something like that does materialize.

Kirkpatrick: Thank you for saying that. There's a hand back there. Let's quickly identify yourself. We'll try to get as many voices on the table as we can.

Reynolds: Name is Josh Reynolds, Rob Roy Consulting. Question is this: I appreciate, Benjamin, you bringing up the importance of making sure that the human side of this is considered very much as well, and Mary Lou, I very much appreciate you being the first panelist to talk about intelligence augmentation.

My question is this: I've heard an awful lot of discussion of about how the human brain and human language can trigger computing. I would love to hear to what extent are each of you concerned about the impact this has on the development of human cognition? To what extent should we be concerned about atrophy of neural processes? And to what extent should we be worried about being so lazy?

Kirkpatrick: Okay, hold those responses. I think that's an interesting comment in itself, if you take the question mark off it. So who else has anything else they want to say before we—Dimitri, please get the mike over here.

Audience Member: So to me the most obvious sort of concern here deals with the intermediation and the agency and autonomy and authentication, right? Because what all of these technologies are doing is they're disintermediating people, increasingly, right? So when you're talking about thinking, for example, there needs to be some type of technology that's translating those thoughts to each person. So my concern with all these technologies is—and this is my question—are you thinking about how this—what the intermediation process is and how that safeguards our agency and autonomy in this future that we're engineering?

Kirkpatrick: Okay, I want to hear one or two more voices before we go to the audience, unless—okay, back there? Right straight back. And then we'll go there and then we'll hear—I

am the one who sees the clock, so I know what we're up against here. Okay. Identify yourself, please.

Kibbey: Lowinn Kibbey. This may be more for Ben and Mary Lou, but are we skipping past the concern and fear of super-artificial intelligence in saying that we can use this type of technology to imbue ourselves with greater intelligence? And therefore, it's not a threat but an enabler?

Kirkpatrick: Hm, really interesting question. Okay, one last over there.

Herman: Hi, David Herman, co-founder of PsyML. I think when like the general premise an issue here is the idea of the objective function. What is the robot, what is the system being trained to do? What is the objective? I hear a lot of talk and, you know, a lot of startups I interact with it's this idea of utility, you know, what is this useful for? What can this do better for me? Outside of the concept of the human interaction. And, you know, you see a lot of people from backgrounds in computer science and machine learning and AI, with a real kind of gap from people from human psychology. And for me, really successful innovation is things that bring more humanity to a system, and so how are you thinking about that? Where you're building things that aren't just useful, but trusted?

Kirkpatrick: Okay, great. Let's go to Robin and one last person and then—sorry to pile them all up, but I think we got a bunch of interesting ideas out here anyway.

Kim: Sure, Robin Kim from Ruder Finn. This probably builds on some of the other questions. There is often an assumption and a crisis of identity. Okay, as human beings interact machines and we're seeing a lot of studies that show that early childhood development. Kids are getting more curious when it comes to interaction but they're also getting less polite. How are you tracking the etiquette and the emotional intelligence and that impact on kids, early childhood development? And are you building that into how you're channeling AI and robotics to deal with it, beyond self-defense.

Kirkpatrick: That's a great question.

Jepsen: For the one about disintermediation through communicating with thought. I would actually say the opposite. Sharing your brain with somebody is probably going to be more intimate than sex. In terms of the concern for children, I don't think we're going to let children use our systems. I can't remember your question, big AI? No, is AI going to be more powerful? The question and the reason we're talking about all of this—

Kirkpatrick: Computers, like, it's the HAL problem.

Jepsen: Right. The reason we're talking about all this is so that we can define what it means to be responsible. That's why we need to have these open discussions about this.

Kirkpatrick: Okay, yeah.

Lau: Really quickly, I heard a lot of different themes there, but I'll try to summarize. I think humans are pretty darn smart and I think that all of the technology that we're talking about here today is really about creating tools that make us even smarter. Whether those tools eventually become intelligent in their own right, we figured out how we can define that, maybe someday. But I think all of the technology that we're talking about here, the way I look at it, is that we're creating stuff that actually we're going to be able to amplify our own powers, to be able to do a lot more things, maybe faster, maybe better, maybe in a different way. And that's what I'm excited about.

Kirkpatrick: Very cool. And either of you want to say something quick?

Bratton: Yeah, just to put a correction—I'll try to be very brief. In many ways actually my work, because I think it's much more critical of human-centered AI design, not the opposite of this at all. I think that human-centered design has many good things about it in terms of making things usable. It's also, in other respects the peak of human-centered design, by one perspective, is the casino slot machine. It's human-centered design, right? So human-centered design is how we get the Anthropocene writ large. It's how you get Pacific trash gyres. When you design entire complex systems, complex natural and synthetic ecologies around the incremental reward of individual human desire. I think we need to solve for other kinds of problems and to go through the rough Copernican trauma of understanding that we're not the only intelligent species. That our interests are not the only ones that matter. And enter into systems that actually will last a bit longer and our lives will probably be better.

Kirkpatrick: Whoa, some big ideas here. We're not going to eat meat because we're going to figure out what they think.

[LAUGHTER]

Prasad: I would just say on the atrophy and politeness, both are important topics. And politeness we, as Alexa, we definitely take a lot of care on how Alexa can encourage more of the polite behavior. Early days for that, as well. On atrophy, I think there's a lot of research needed to be done in terms of what these AIs are making, acquiring skills very fast in different realms, if you think about it. How are we forgetting something? Are we losing out certain aspects? I think there's more research needed to be done in that area.

Kirkpatrick: What a great conversation. Thank you to all four of you.