

Rapid Response Transcript – Eric Schmidt

“You should be running toward AI,’ w/Eric Schmidt (frmr Google CEO)”

[Click here to listen to the full Masters of Scale: Rapid Response episode with Eric Schmidt.](#)

ERIC SCHMIDT: When most people think about AI, they think about killer robots, and that is precisely not what we're talking about. So the business person should run as fast as they can to this.

Everyone else is, if you don't get there first, somebody else will, and then you're going to be in trouble.

These systems are imprecise, they're dynamic, they're emergent, and they're still learning. They can't explain what they do and they may make mistakes.

Think of them as nuclear weapons of a different kind. They're so dangerous that they need to be protected.

This is a new epic. It's an epic of coexistence with an intelligence that's not human. We've never been through that before.

We're playing with human lives. It's important that our systems not enable the worst of us, but instead promote the best of us.

BOB SAFIAN: That's Eric Schmidt, former CEO of Google and co-founder of Schmidt Futures.

Eric has been outspoken about how business leaders and all of us need to plan and adjust based on where tech is going.

I'm Bob Safian, former editor of Fast Company, founder of the Flux Group, and host of Masters of Scale: Rapid Response.

I wanted to talk with Eric because his insights, particularly around artificial intelligence, are both encouraging and disturbing.

In a new book called *The Age of AI* that Eric co-authored with Dr. Henry Kissinger and MIT's Daniel Huttonlocher, he argues that AI will change business, society, and potentially humanity itself.

Eric is not an alarmist, but he is alarmed about potential shifts and unintended consequences.

As the chairman of the National Security Commission on Artificial Intelligence, and before that chair of the Department of Defense's Innovation Board, Eric has in-depth knowledge of the military and national security risks posed by AI.

But his insights run even deeper, from the impact on startups to the responsibilities of tech giants like Google and Facebook. AI is going to make us smarter and more productive, he says, but that's only the beginning.

We're entering a new and unknown era, as he describes it. One that requires extra work, extra vigilance and extra care to make sure we amplify the positives and de-amplify the negatives.

Don't expect the government or someone else to protect us against concerns, Eric says. Look in the mirror, and bring the human element to bear.

[THEME MUSIC]

SAFIAN: I'm Bob Safian, and I'm here with Eric. Schmidt, the former CEO of Google, and co-author of several books, including *The Age of AI*, recently released with Henry Kissinger and MIT's Daniel Huttenlocher as co-authors. Eric, thanks for joining us.

SCHMIDT: I'm so glad to be back on *Masters of Scale*.

SAFIAN: Yeah, this is your second time on *Rapid Response*, your third time on *Masters of Scale*. You were a guest on this *Rapid Response* show last fall, talking about the challenges of COVID-19 and the obstacles facing business leaders. More recently, you've talked a bit about Facebook and the state of big tech, and I hope we can get into some of those topics a little bit, but I want to, of course, start with the book.

Artificial intelligence is a topic, a technology that engenders some excitement in some quarters, fear in others, increasingly prevalent in all kinds of organizations, plans, operations. Can you start by briefly defining what AI is in the way that you think about it, and then why you took on this book?

SCHMIDT: Well, thank you, and it's really great to do this. When most people think about AI, they think about killer robots, and that is precisely not what we're talking about. What we're talking about is the fact that the world around us will be, over time, full of intelligence systems that are human-like, but not human intelligence, and the book is a statement that we have to get ahead of the implications of this on the human race.

If you look back 20 years ago, 15 years ago, when I was doing social media stuff, we just were naive, and I want us to avoid that with this even much more powerful technology.

Dr. Kissinger, when he first came to Google, which was about 12 years ago, he gave a speech where he started by saying, "I believe Google is a threat to civilization," which of course the Googlers loved, because they love the attention.

And the reason he said that, I now understand, is that he doesn't believe that the power of information should be in the hands of any single individual or firm. So our book is an exploration of what the issues are when these extraordinarily powerful AI systems will be part of our daily lives, guiding us, helping us, misleading us.

SAFIAN: When you say that AI is human-like, but not human, what makes it like a human, because it's in some ways faster than human, right? And what are the ways that are not human?

SCHMIDT: In the book, we use a couple of examples, one of which I think is quite interesting, it's called GPT-3. So GPT-3 is a transformer that is generative, and what it does is it sucks in all the text information around the world and, remarkably, appears to understand it. And if you ask it questions, it can generate answers. You know, design me a website, that sort of thing.

And what's interesting to me about GPT-3 is that it was arrived at in a way very different from how human brains work. And now there are six or seven companies that I'm aware of that are trying to build trillion-parameter models, which is five or 10 times larger than GPT-3, and so far, the people who are doing them tell me that they're increasing returns to scale, that the larger the model, the more information, the more understanding.

The important point is that many people think that within a decade, we will have systems that have human-like intelligence, but remember, it's not the same. I'll give you the rule. These systems are imprecise, they're dynamic, they're emergent, and they're still learning. They're imprecise because they can't explain what they do and they may make mistakes. They're dynamic because they're constantly changing. They're emergent because when you combine them, you get behaviors that you didn't expect.

So imagine a situation where your kid's toy is talking to the kid, but what's it learning? Don't worry about the kid. The kid has a teacher and a parent—but who's teaching the toy, right? And what if it gets the wrong idea? Maybe it becomes racist temporarily. How do we control that? How do we deal with that? None of these questions are answered.

SAFIAN: And so the way these systems learn and grow is maybe not predictable, or at least need to be monitored?

SCHMIDT: Yeah. The question is, can you thoroughly monitor such a system? And the problem is that they're so complicated. You don't know what they know, and therefore you don't know what they're learning. There are many, many researchers that are working on the explainability problem, and also on the bias problem. So it's possible that

these systems will ultimately be able to fully explain themselves, and guarantee, if you will, that they're not biased in some horrific way, but don't bet on it.

SAFIAN: When we talked a year ago, I'm just going to quote something you said to me, you said, "the reality is that AI is going to make us smarter, more productive. It'll help in education. It'll help reach people who haven't been reached. It'll make our business more efficient and scale and so forth." It's a remarkable story. So there's a lot of positives around that you see around AI, also. What does that balance look like?

SCHMIDT: I'll give you the punchline. And the punchline is that this stuff is incredibly, incredibly important and powerful and positive. It's also incredibly, incredibly worrisome, and we need to get ahead of the downside and promote the upside, obviously.

SAFIAN: Because it's not going to go away?

SCHMIDT: No, no. In our industry, this is a wave that is going to take over everything. Thousands of thousands and thousands of students are busy learning how to do complicated machine learning algorithms. It's not a narrow thing. Its applicability is extraordinarily broad across pretty much every human endeavor.

In the book, we use an example of a drug called halicin. We are collectively developing a broad-scale resistance to antibiotics. So this particular team set out, and they said, let's figure out if we can find a drug that is very different, but also is an antibiotic.

So they took the principles of antibiotic resistance, and they taught the system that, and then they applied it to a hundred million compounds and they said, sort all this stuff out. Then they built another network that scored the ones that they came up with, and out of that came a single candidate whose name is halicin, who appears to work.

Now, that's not something that a human could do. It's also not something that a computer on its own could do. Now, what's the value of a broad-spectrum antibiotic for which we've not developed resistance? It's enormous, and it's something that we couldn't do without our computer friends.

Another example is if you look at AlphaChess and AlphaGo. They were given the rules of the game, and then they learned how to play it, and they learned how to win, and they could beat any human in history. But more importantly, it developed a new set of moves, a new set of strategies. It plays differently from humans. Now, the fact that these systems discovered a new set of moves, does that indicate that there are possibilities in the virtual world that we as humans cannot see and cannot comprehend?

It's always possible that there are principles of the world that humans as a species cannot comprehend. What if the AI system comprehends them in some form, but it cannot explain it to us because we cannot understand? If that's true, Dr. Kissinger would

say from the lesson of history that one of two things happens. Either people revolt, they become fearful. Or, a new religion forms around something that we cannot understand and we cannot master. I don't think we know, but these are examples of why we wrote in the book that this is a new epoch of human industry. It's not just a new generation. The direct analogy is the transition from what was called the age of faith to the age of reason, which occurred a few hundred years ago.

Today we have the age of AI emerging. Does it replace the age of reason? Is it based on the age of reason? We don't know.

SAFIAN: To understand those two examples, the new antibiotic is sort of high-level number crunching that maybe a human couldn't do, or couldn't do in that time, but it's identified something. Whereas with the game advance, it's like there's a whole different way of playing, that these are two different ways that AI shows itself.

SCHMIDT: That's correct. Try to imagine what the next five to 10 years look like. And we have a section in the book called "From Turing to Today — and Beyond," there will be improvements in algorithms, which will come to the point where it begins to look like these systems have volition. That is that they can begin to set their own objective functions.

In theory, they can also begin to write code. That is the beginning of a real transformation, because now the system is imprecise, dynamic, emergent and learning, but it's also changing itself.

SAFIAN: It's creating.

SCHMIDT: It's creating itself. Now this is speculation. One of the things I've learned in my now 50 years in this industry is that the consensus in terms of the technology direction tends to be correct. It's very hard to pick the businesses and some of the applications, but the underlying technology platform is fairly predictable. The way to think about these is to think of them as nuclear weapons of a different kind. They're so dangerous that they need to be protected. Dr. Kissinger in the 1950s, served in a series of task forces at MIT and Harvard and other places, Caltech.

At the Dawn of the nuclear age, where they worked out the principles of nonproliferation. They were informed by the physics of the time and the strategists of the time, but they were also informed by the humanists and the philosophers of the time. And I think it will happen. We can debate when. These systems are going to be very carefully protected and they're not going to be very many of them, and we're going to have to come up with some limitations, agreed upon for the whole world as to how they can be used. And hopefully once we've invented these things, they won't be misused to destroy the world.

SAFIAN: As I'm listening to you, there is this sort of edge case, advanced case of AI. You want it not to proliferate. At the same time, there is more basic AI that's underway right now and proliferating everywhere. And you're not necessarily saying that that type of AI should be restrained and not shared.

SCHMIDT: Yeah. So it's clear. I'm not calling for regulation at this time. Partly because it's so new they don't know what to regulate. In other words, you can regulate killer robots, but we're not building them. So that's fine. You can ban the movie scenario all you want, but that's not what I'm worried about. What I am worried about is even with existing technology, you're going to see systems, which at least in conflict are going to skirt human decision making.

SAFIAN: So a lot of the listeners on this show are business people running their own startups or parts of businesses, and they are being asked or offered an opportunity to start using AI in different parts of their business. Should they be anxious about that, or should they be running to embrace it.

SCHMIDT: So the business person should run as fast as they can to this. And I'll give you a simple way of thinking about it. If I had 10 or 20 Google-level engineers in essentially any business, I could improve its revenue, reach, profitability, scale. And the way I would do it is I would take those engineers. And I would say, "I'm not going to tell you what to do. I just want you to study our business and how we deal with our customers and our technology platform. And I just want you to make it smarter. And the definition of smarter is I want you to use machine learning in order to produce whatever the business is that you're in."-These work best, which is why they're so successful in tech businesses, when you have a lot of training data.

If you just took an AI system, and indeed Google has done this, and you monitored what was going on, you can not only predict the correct answer, but you can also improve it relative to the existing traditional algorithms. So we talk a lot about how businesses are going to be just in time to the person and so forth. It's perfectly possible that every business will have to learn what the preferences and interests are of their customers. And the businesses that don't talk to their end users are going to be in trouble. So if you look at Netflix versus the movie studios, they're both proud organizations, but Netflix had one thing that the movie studios didn't have, they didn't talk to their customers. They talked to the theater distributors, but Netflix of course, developed recommendation engines and so forth.

If you're starting a company, your company's going to look like this. You're going to have an iPhone and an Android, platform, you're going to have a fast network, and you're going to have a backend server, that backend server's going to have a lot of data coming in that you can learn from, and you can improve and improve and improve and improve.

[AD BREAK]

SAFIAN: Before the break, we heard former Google CEO Eric Schmidt talk about how artificial intelligence and algorithms are unleashing new possibilities, some quite troubling but also powerful enough that any business person should run to take advantage of them.

Now he talks about the practicalities of implementing algorithms, the missteps that Facebook has fallen into, and the limitations of regulation. He also shares lessons from the pandemic and what he calls the golden age of biology, as well as the looming uncertainties as we enter what he calls “an epic of co-existence with an intelligence that’s not human.” Eric’s clear that he doesn’t have all the answers, but he’s a firm believer in the value of asking the right questions.

If you're a business that has not been as technologically savvy as a Google or a Netflix, is AI the kind of technology that you can use to like leapfrog to get sort of in the game, or do you need a foundation? Like if you're a smaller business and you don't have the volume of data.

SCHMIDT: So today the data discrepancy is very real. The more data you have, telemetry, scheduling, all of that, the more like you are to be able to use this. But my overall comment to businesses is you should be running toward this. Everyone else is, if you don't get there first, somebody else will, and then you're going to be in trouble, because they'll do a better job of serving their customers. And frankly, what happens is people who are well-meaning, but not very technical, listen to this and they say, " Okay, well I'll hire a consultant."

No, that's not how it works. You actually have to do of the hard work. You have to actually find the engineers. Unfortunately, it's not like email, which you can just turn it on and it works.

We're not there yet. Google had a series of initiatives called AutoML and others, which are attempts to make this easy. If you look at where the research is, it's still in this incredibly complicated algorithm design. This is one of the reasons why the brand new PhDs in computer science are making millions of dollars a year, which is sort of shocking. The reason is they're worth it.

SAFIAN: If I'm hearing you about AI, it's a tech that right now sort of favors bigger businesses that have more data and can invest and build capacity, but somewhere down the line, it may be able to help level the playing field a little for smaller businesses, or am I misinterpreting you here?

SCHMIDT: No, I think that's roughly correct. Remember that any reasonably well run business is going to do this on the cloud. So the typical transition is they have some legacy system and they basically lift it out of those computers and they put it onto the cloud and they don't change anything. They have the good idea of putting all the data on one place, kind of combine it all and standardize the identities.

But once you have that, you have a database of all of your customers, all of your activities, all your SKUs, all your products, then you could begin the use of these algorithms.

SAFIAN: One of the big companies that's making use of algorithms is Facebook or now Meta. You've talked recently about Facebook's missteps, you've been critical, but you also aren't supportive of regulation, can you explain that?

SCHMIDT: Well, let's understand what social media is really about. The system is organized around engagement, watching, watching, interacting. So if you're trying to maximize your revenue. The best way to increase engagement is outrage. They're not truth tellers, they just try to keep us outraged on any side of the issue. Once you understand that the system is built, they all go for this, then you understand you have to have either intelligent leadership or public pressure or employee revolt or something for the extreme cases. And I think that's what you're seeing with Facebook.

SAFIAN: And none of those examples you cite of ways to respond are about government getting involved.

SCHMIDT: Well, again, whenever you see something you don't like, you wish the government would fix it. There's lots of things in the world I see every day that I'd like the government to fix, but the government's not very good at it. I would argue to you that the odds of getting the regulation around attention and outrage right is pretty low. Sridhar Ramaswamy, who ran all the ad systems at Google, started a competitor which is where you subscribe to it as opposed to using advertising. And he said, "Engagement is the problem not the solution. Engagement based ecosystem rewards the worst actors on any platform because their awful behavior gets the most attention."

The systems were designed to produce this outcome, we have to have a proper conversation about how to fix it. At Google when we face these issues, we would have a complicated trade off between quality and revenue, and I think we made roughly the right trade offs. If you look at YouTube, which is not particularly a big problem in this area, YouTube has a series of non-amplifiers. So for example, if you have a questionable video, like a radicalization video, which is permitted under the rules of the platform, we won't serve you the next one and the next one and the next one, we'll serve you something different. That's a good example of a company making a decision to limit the extreme rewarding of radical ideas again, on any side. I'm not making a political comment.

SAFIAN: Sridhar Ramaswamy, who was on the show a few weeks ago. I mean, his new business is sort of based on the idea that the business model that underscores a lot of media companies, social media, as well as Google that it makes those good choices impossible in the long run, that you're going to be driven to bad places if that's the business model.

SCHMIDT: And this is Sridhar's view. If you look at the development of the printing press, you had essentially broadsheets. They were regulated. You had advertising that was regulated. So society has had technology points before in history where they were ultimately regulated because the system's optimization was incompatible with how humans behaved. We're going to have to figure out some kind of de-amplifier. I'm very much in favor of free speech. You're welcome to have any opinion that you want, especially if it's really, really stupid, but I'm not in favor of you being amplified by computers. I'm not in favor of robot speech. And there's always been crackpots in society. The difference now is that not only can they find each other, but then they can organize, and finding each other and organizing against things which are factually false is not helpful to society moving forward. We also know that even if I tell you a video is false and I show you the video, it will change your behavior even if you know it's false.

SAFIAN: That watching a video that's untrue even if you're told it's untrue, you still start to believe some of it.

SCHMIDT: It changes your behavior. Inside in your head, the images, especially bad images, difficult images are seared in your brain. I think that the purpose of the book is to say, we need to have a conversation about these more than just the tech people. We need to have groups of people that include economists, psychologists, business people, political people. If you're upset about the Facebook stuff, it's going to get a lot worse. These open source libraries will allow anyone to produce misinformation at any level. I'll be able to download the software and make fake videos, falsify what politicians say, make all sorts of false statements.

There's plenty of people willing to actually produce lies that hurt people, and they don't seem to have a moral compass. The majority of the anti-vax stuff comes from a relatively small number of people who can be identified and can be stopped. Now, why should we stop anti-vaxxing? Because they're killing people, they're repeating falsehoods, and the normal person's just trying to get through the day and is susceptible to these arguments. Now, how do we solve that? The answer is that the platforms have to identify who they are and then take them down.

SAFIAN: And when you say the platforms have to identify, I mean, if I understand what Mark Zuckerberg says they're trying to do at Facebook, there's so much volume of content across the platform that it has to be the algorithms that are identifying and parsing it as opposed to humans are humanists at each stage of the way doing it.

SCHMIDT: Well, there's no way for Facebook or Google or whatever, to police all the content that's coming into them. So the best scenario is probably to come to an agreement as to what is permissible to be amplified and then to build networks, neural networks that can identify that. In other words, the person who's just screaming all day and screaming falsehoods can probably be identified and then build an agreement that that's not okay.

Now, there's all sorts of issues. Let's say you do this with the legitimate companies, Facebook, Google, and so forth and so on, what do you do about the startups? What do you do about the nation state that's trying to weaponize this such as Russia with kompromat, and trying to actually pollute our systems. So it's hard to be a social network now, you have to police all of the bad behaviors, police all the people trying to manipulate all the good people that you have, and we do not have a theory of how to regulate this.

SAFIAN: I was looking back at our talk a year ago, and in some ways, it felt like a different era. We were talking mostly about pandemic. There were no vaccines available. There was no promising oral antivirals like those recently announced by Merck and Pfizer. There was no discussion of Delta variant either, although you warned us to expect some ongoing flare ups. I'm just curious on that front, with all the progress that we've made, we're still not back to normal, whatever that means. Where are we on that for business leaders? Where should they be focused in those areas?

SCHMIDT: Could I lose my cool for a sec?

SAFIAN: Please!

SCHMIDT: Look, this is horrific. We have normalized 1,200 people a day dying from a disease many of which could be kept away. My friend Colin Powell died. He was ill with a disease and COVID took him out. Someone gave him that COVID. The core problem here is that people don't seem to understand the basic fact about this disease, which is that your primary contagiousness is before you yourself have symptoms. So you're highly likely the day you have the symptoms, to have spread it to people in the proceeding two days. As a moral matter, I don't want to be that person. I don't want to infect anyone ever.

Is there some reason why you can't, as a matter of politeness, wear a mask when you're with other people, especially since you may at least in some probability be contagious? So today it looks like the belief among the doctors, and again, this is not medical advice, please, is that the best path is the two shots of the mRNA vaccines followed by a booster, for example, eight months later. And it's thought that will give you protection for a fairly long time. So we're in a situation where a combination of mask wearing and vaccination will allow us to get back to work, get us back to the office.

Now, there's all sorts of implications from the pandemic. The fact that people are not coming back to work. The acceleration of digital change. The tech industry has benefited enormously from this. But as we return, I mourn the dead, and in particular I mourn the dead that didn't have to die. We're going to end up with something close to a million people dying from this disease in the United States. That ties the record to the number of people who died in the Civil War. How is that okay? How are we normalizing this? Why

are we not having a day of shame on ourselves for our callousness to our fellow Americans?

And these were incremental and individual decisions which collectively produced hundreds of thousands of additional deaths that did not have to happen. So, that's what I really think. Now, having said that, the current consensus is that the winter will have some outbreaks. The primary round of transmission appears to be at home. But hopefully this is the last such wave. This happened because people wore masks and they got vaccinated. Thank God for the people who figured out that masks worked. Thank God for the people who put Operation Warp Speed together.

Remember, Warp Speed. It's industrial policy. It was the government guaranteed the products whether they worked or not. Wow, that's pretty big. That's a violation of our doctrine. We're celebrating Pfizer and Moderna and to some degree J&J, but there were 20 others that didn't make it. Those were huge risks that failed. And then the university community and the acceleration in bio has been profound.

SAFIAN: Did the pandemic make AI more prevalent?

SCHMIDT: I think that as a general statement the AI revolution has been occurring independent of any outside fact. It's self generated. What is true is that in the biology world things have changed dramatically because the research is now occurring before peer reviewed publication. And that's collectively how that community responded to accelerate its own rate of discovery. And that's a wonderful thing.

SAFIAN: So, in this environment then, Eric, what does it mean to be human?

SCHMIDT: Well, today, what it means is that we have these AI systems and we don't seem to mind them, recommendation engines, Google translate, that sort of thing. We foresee, and we discussed in the book, that over time you'll be surrounded by these systems and the way they work will have a lot of impact on your daily life. It's obvious, for example, that the broad applicability of deep fake and other manipulative software will allow people to try to sell you fake things. There'll be an explosion in misinformation, sometimes evil, sometimes misinformed. So it's pretty likely that to get through the day people are going to have to have their own assistant and it sorts all this out and says, "That's false. That's stupid. You don't like him anyway," and so forth. And you can imagine over time, we will become very codependent upon this digital thing, we'll call it a machine for lack of a better term, that we need to get through our day.

Now, these systems will educate us in ways that are very different. These systems are going to have so much data that they're going to appear to be another kind of intelligence. It won't be human-like, but it will be intelligence. And what we don't know is when we're coincident, when we're coexisting with, how will we react? Will we become more patient, less patient, more stressed, more neurotic, less stressed? How will we

master it? How will our identity come out? These are questions that are not for me. These are questions for the great philosophers. But the important thing is this is coming, and this is a new epic. It's an epic of coexistence with an intelligence that's not human. We've never been through that before.

SAFIAN: And as we move through the progress, we just have to be more intentional.

SCHMIDT: Well, sometimes these things just sneak up on you, and I want people to understand the implications. We're playing with human lives. Humans are not computers. Humans are not rational all the time. It's important to understand that we're building for an audience of all humans and all humans include an awful lot of behaviors that we either don't like or will not be supporting. It's important that our systems not enable the worst of us, but instead promote the best of us.

SAFIAN: Well, Eric, this has been great once again and thank you for taking the time and opening our eyes and our minds.

SCHMIDT: Thank you so much. I'm really glad and I'll see you again.