



Eric Schmidt Keynote Address at the Center for a New American Security Artificial Intelligence and Global Security Summit

November 1, 2017

Washington, DC

Transcript prepared by Gregory C. Allen and Anthony Cho. This transcript has been edited for clarity. Discussion starts at :30 of video

Paul Scharre: So thanks for joining us. I wanted to start by asking you from where you sit looking at your companies and thinking about the United States as a whole, how big is artificial intelligence? How important this to be a leader in Artificial Intelligence?

Eric Schmidt: If you're in a business that has a lot of data, and in an information business of some kind, it's more than big, it's essential because it will become your competitive advantage. And I think it's true that pretty much all of the knowledge system kinds if companies will become AI companies in that sense. A lot of people believe it'll be more than that, but that I know from the standpoint of Alphabet because it touches everything we're doing. When I was CEO we had mobile first, the new CEO announced AI-first as our core strategy. That's how serious it is.

Paul Scharre: What do you see as the key drivers of competitive advantage? What are the things that allow a company to be ahead of rivals? Is it human capital? Is it data? What are those key components?

Eric Schmidt: It's sort of both. The first observation is that this stuff is still essentially magic. The scientists that are working on it cannot for example explain certain failure modes. We don't really exactly understand how the learning occurs. There's evidence for example that the face recognition and image recognition is using slightly different parameters than what we do as humans. That's an example. So it's a slightly different animal if you will.

So at the moment it's in the province of very highly skilled mathematicians, PhDs, those sorts of people. And we're in a situation where those kinds of people, graduating out of Carnegie Mellon and others, are in the highest demand I've ever seen with huge multimillion dollar packages in their twenties coming out. That's how valuable these people are in the marketplace.

But the other thing that's worth saying is that these algorithms, at least today, require a great deal of training data. And when I say a great deal I mean like, millions of entries in the matrices, billions of pieces of data. So the classic example: people would say "well why can't you figure out terrorism?" Well the good news is terrorism is very rare. Right? So it's much, much harder, if

you will, to apply AI to that problem. Whereas trying to understand traffic, right? As an example something that occurs every day, is far, far easier because you have so much training data.

Another reason to think that AI works so well in biology is that as much as we would like to find differences among ourselves, we are all the same biologically. And the fact that we have so many examples of us that we can measure means enormous gains in our healthcare globally, for everybody.

Paul Scharre: So in your role as the chair of the Defense Innovation Advisory Board, what is your assessment of how is the Defense Department is doing in terms of looking at artificial intelligence and incorporating this?

Eric Schmidt: So a couple of years ago I became the chairman of an external advisory board called the Defense Innovation Board and the idea was to try to get the military up to speed with things which are going outside the military. We made a set of 10 recommendations, which are public and you can read them and I would encourage you to do so. We obviously worked pretty hard on them. And they cover a lot of ground, that looking at this audience you're familiar with: the nature of innovation in the military, the kind, what are the problems and so for the Third Offset. Bob Work is here, etcetera. There are all sorts of examples of all of this. One of the most important point we made is that the military is not leading in AI. That we recommend, for example, the creation of the equivalent of an AI institute.

And I can give you a positive and a negative example. So in peacetime anyway and I think also in wartime what do our men and women do mostly? They mostly watch things. Okay. Well, computer vision is a solved problem compared to human vision. And roughly speaking, people's ability to watch continuous scenes with no change is not 100%. There are studies which show the error rates are one-third, out of a hundred, you know one-third error rates in various tests. Whereas computers can watch a scene which is monotonous for a very, very long time, and they'll alert you for a change. Now that seems like the simplest possible thing, and yet we have this whole, you know, tradition of the military standing watch and so forth and so on as if that's a good use of human beings. So the core narrative about AI is not that AI will be like us, that AI is different than us. And that the best uses of AI will be in human and AI collaboration of one kind or another. The vision example is a classic one. It's also true of self-driving cars.

It turns out that the military has started, it's interesting, it started in the intelligence group, a group called project MAVEN, which they are talking about, which actually combines a lot of these very clever systems under a general named Shanahan [Air Force Lt. Gen. Jack Shanahan, director for defense intelligence for warfighter support], it's very clever.

Paul Scharre: I guess how are those initiatives going? I mean, do you see that DOD is making progress adopting that? Do you see obstacles?



Eric Schmidt: Well, you worked at OSD. I don't need to tell you that the military is very large. It's very cumbersome. The military as a general statement doesn't build things. It uses contractors. So the problem is that military, the contractors build what the military asks for because they are contractors and the military has not been asking for AI systems. And it takes 5-10-20-30 years to go from the spec to the delivery. Am I summarizing the history of 40 years of procurement or did I get it wrong? I could make it worse if you want in my report.

So, so the core problem is how do you get the leaders, who want passionately this stuff done to be able to deliver these solutions quickly. My own view is that they have to write AI and machine learning more technically correctly in every procurement. The military contractor base is largely good at hardware not very good at software. I don't need to tell you that. You use it every day. The software systems aren't interconnected. They don't use cloud computing. There's a long list and these are covered by our recommendations.

So MAVEN is interesting, to give you an example, because it's an intelligence application, which is good because it isn't life safety, you know, in other words it can make a mistake and makes a recommendation. And they train it on open source data, which I think is most interest. So there's another good example of the military coming up with a new strategy that uses open source data to make their systems smarter and one of the tricks of AI is you need large amounts of training data. And there's always more open source data than there is secret data.

Paul Scharre: So one of the things that is different in this space right now, than say earlier major military innovations from thirty years ago, things like stealth, precision-guided weapons, is that these are coming from outside the defense sector, right? How does the defense department begin to import these technologies? And then you talked about the human capital component as well, right? How does DOD compete for talent and then draw in companies in this space?

Eric Schmidt: Well, in the first place, there's lots of money in the DOD. The programs are large. They're multi-year. They have these large procurements. Much of the software work is not of the best quality. There's plenty of room here to bring in some of the kind of people that I'm talking about. And the military has mechanisms to build software contractors, so they don't have to pay military salaries to some of these people, which are way too low for the market. So I think the mechanisms are there, the specific mechanisms to get access. I think it's really about strategy and will.

My overall feeling is that if you start with [computer] vision, right? Which is something that everyone is doing, everyone is watching, everyone is learning, everyone can understand that, right? So why are we taking people, exquisitely trained men and women of our military, and having them watch something boring all day? I know that sounds repetitive, but that's literally what they do right? For hours and hours and hours.



And the other activity, which I find very disturbing, is called, basically it's called "turn and type." They have two different screens. This system uses one dataset, this other one uses another dataset, and they have, again, soldiers taking the data from one screen and typing it into another because your systems aren't interoperable. Again, these are straightforward problems. I know we want to talk about, the stuff that Andrew [Moore] talked about, the incredibly interesting stuff, but a lot of the gains here are making the existing folks just that much more effective. It's all right in front of you.

Paul Scharre: When you are talk about that vision, are you seeing the senior department leaders that you interact with, that they're starting to get it? They're starting to see this is a key area to be ahead on.

Eric Schmidt: They are. Under [former Secretary of Defense] Ash Carter and now under [Secretary of Defense] Mattis, everybody kind of gets this. The problem is, everyone can understand something, but they cannot collectively act is the sort of the core governance problem. So you have to come up with ways for them to be able to get the resources and so forth. Now, if we were in a huge war with a major adversary, I'm sure the rules would be different, but right now the planning procedures and so forth in my view take too long. That's why we in our recommendations we talked about embracing a culture of innovation, trying things, failing quick, all the standard Silicon Valley kind of things that people talk about. It will help if the military moves to cloud computing for much of its infrastructure. The DOD has just announced a major cloud computing initiative, which they're very serious about, and I'm told that compared to other initiatives, the DOD is moving very, very quickly compared to many other things they are doing. My view is that it's not fast enough, but I say that at Google too, so it's okay.

Paul Scharre: I want to turn to China for a second. Some of the top AI companies in the world are in China. What is your sense of how the United States is relative to China in terms of the talent and innovation that's going on on Artificial Intelligence?

Eric Schmidt: I don't know if you all saw, but a month ago China actually announced its AI strategy, and I actually read it. And it's pretty simple. By 2020 they will have caught up. By 2025 they will be better than us, and by 2030 they will dominate the industries of AI.

Just stop for a second. That's, the [Chinese] government said that. Weren't we the ones in charge of AI dominance here in our country? Weren't we the ones who invented this stuff? Weren't we the ones that were going to exploit the benefits of all this technology for betterment and American Exceptionalism in our own arrogant view? Trust me, these Chinese people are good.

We were, in April, we brought our AlphaGo game, which is a remarkable reinforcement learning activity, and brought it to Shanghai, Wuzhen technically, and had a series of demonstration games against the very top Go player in the world, and we won, we won every game. He was a



brilliant Go player. He invented new moves. We invented new moves. It was a historic if you're a Go player. Then we retired the game.

The game that was played learned how to play Go in seven days from zero, right? That's how smart these systems have become – and beat the top human in a game that's 2500 years old. To me the more interesting thing about going there was to see that all the top computer science people from China had shown up. And my conclusion was they are using, going to use this technology. They're going to use it for both commercial as well as military objectives – with all sorts of implications over that period. But, we know what they are doing: they have announced their strategy! You're crazy to treat them as somehow second class citizens.

They have... Let me give another example, we do tests globally, there are these things called top coder contests, where you have a coding program literally for programmers. And Asians as a general statement, and Chinese in particular, tend to win many of the top spots. So if you have any kind of prejudice that or concern that somehow their system and educational system is not going to produce the kind of people that I'm talking about, you're wrong.

Paul Scharre: So what steps should the U.S. be taking to remain competitive, right? Presumably that's not where we would like to be in 2030.

Eric Schmidt: Well for one thing we don't have a national strategy, I just said that. Alright. If you believe that this is important – as I suspect all of us do, but certainly I believe – then we need to get our act together as a country. Alright. So I'll give another example. Look at the history of Sputnik and the highway system: what happened was the Russians, everybody knew that the Russians were going to do this, and then the 'Sputnik Moment' happens, and then NASA happens, right? Eisenhower and the interstate highway system in the 1950s, right? This is that moment. This is the moment where the government collectively with private industry needs to say these technologies are important.

And by the way, whether they really needed the interstate highway system to move all those missiles around, the country benefitted by the interstate highway system enormously. Just think about it. So whether it's from a position of fear, where people are afraid of something, or whether it's a position of leadership, I don't care how we get there. But I do think that a national focus on investing, starting with... So the specifics are straightforward – investing in research. America is the country that leads in these areas. There's every reason to think we can continue that leadership. It's also how we will differentiate our businesses globally. Software is the leading... The leading place that software is done in the world is done in America for many reasons including cultural and so forth. There are all sorts of reasons to do that.

I'll give you another one. Let's talk about immigration. Shockingly, some of the very best people are in countries that we won't let in to America. Would you rather have them building AI somewhere else or having them build it here? I'll give you a specific example: Iran produces



some of the smartest and top computer scientists in the world. I want them here! And to be clear I want them working for Alphabet and Google. I'm very, very clear on this. It's crazy not to let these people in. So I could go on.

Paul Scharre: Okay, good. Well let's open up to questions from the audience. We've got a great selection of folks from across government, the private sector, and academia working on AI. Let's hear some questions. We've got some folks in the back with microphones.

Audience Question: Hi, thank you. Thanks a lot for your talk. Joanna Bryson. I loved most of it, but of course I'm going to pick on the one piece that worried me. You were talking about replacing large numbers of military people with an AI vision system, could that be a single point of failure? Could that be something that could be hacked or manipulated? Whereas, you know, thousands of soldiers would presumably whistle blow if something weird was going on.

Eric Schmidt: So there are many issues around the design of these systems that you would want to get right. So the first thing to do is – the military's position and I agree with this – is that a human should be in the loop in anything involving kinetic activity. That seems obvious to me. So you want to start where the areas, where you're making the soldiers more efficient.

So yes of course, if you build a tool that is then hacked by the Russians or the Chinese or whatever the scenario you have, that's worse than, the tool, not having the tool. So make one that's not hackable. I think these are all straightforward engineering problems and they need to be done right, in the same sense that the ships need not to sink and the airplanes need to fly and so forth and so on. These are known problems. The reason I picked [computer] vision is vision is easily understood if we understand its failure modes. It's also something where people can say "Let me check that. Let me check another source. Let me see. Do I agree that a missile was launched? Can I, the computer says that a missile was launched, let me check over at this other site. Let me see if everything agrees before I do anything. Maybe it's a false positive." We understand how to build reliable systems using that kind of error rates.

Audience Question: I also worked at Google for Eric at one point. [Eric Schmidt: You forget for a decade!] I think that right now the technology that we built to monitor machine learning systems, you probably know the code words internally, the things that we used to make sure that at Halloween our machine learning systems didn't go crazy because they were so overwhelmed with queries for candy. That was very much ahead of what I see happening in academia or in the Defense Department at the moment. And I'm very interested in putting in a large program to help the rest of society catch up. We, in Google, I don't want to sound boastful, we were really good at this and in fact about half of our machine learning teams were writing the systems which are watching and monitoring, constantly testing and diagnosing machine learning. Do you think that we should invest in that? Do you think we should maybe actually start to open some of the commercial techniques to the rest of the community?



Eric Schmidt: I do. I also think that the way this will really work if you go back to [Carnegie Mellon Dean of the School of Computer Science] Andrew's [Moore's] point about how it's very easy to mash up a very powerful system today, you precisely cannot do that in the military. There are a gazillion laws that prevent you from doing "mashups" and having fun, for all sorts of, again, reasons. So I am very much in favor of partnerships between the military and universities, whether they're simply a building next to the university or the military helping fund research in a university, in an unclassified and open source kind of way, so that the military can get the benefit of these kinds of systems. I think that it's good for the country. It's good for the universities. It's good for the research, as long as it's done in the open and not in a classified way.

Audience Question: Thank you very much for those views. Very interesting. My principal question is around how you see the China-U.S. competition. A few observations on that is that China is in a position to spend orders of magnitude greater capital in both creating the demand for these AI products and directly building the AI products. With the One Belt, One Road, they're also integrating and partnering up with many other countries that produce great talent that they will have access to. And in fact, there are some statistics that indicate that because of the immigration climate in the United States, Chinese students and other smart students that would have chosen the U.S. as their number one destination are now not doing so in measurable ways. One can go on and on, but at the same time in the U.S. we've seen this budget proposal, which actually, the most recent Trump budget proposal, which actually cuts down science and technology funding. And there is a whole bunch of things also that sort of get to the heart of what we're trying to do with education by making all of these changes. It's certainly not to promote more science, certainly there's nothing about AI. So all of these decisions, they have a lifetime. I mean, when you make a decision like this and you commit to a decision like this, you don't get to redo it three months and six months down the road. It has a consequence once you go down that path. So given that these decisions now have a window of consequence, how do you rate the probability of, in five years, the US will remain a dominant AI player compared to China?

Eric Schmidt: I'm assuming that our lead will continue over the next five years, and that China will catch up extremely quickly. So, in five years we'll kind of be at the same level, possibly. It's hard to see how China would have passed us in that period, although their rate of improvement is so impressively good.

It feels, as an American, that we're fighting this particular conflict over AI with one hand behind our back, supported by your point. What I would rather do is not adopt Chinese policies but rather "more like us." Do more of the things that we're particularly good at: funding basic research. As you pointed out, the Trump budget does reduce that. It's the wrong direction. Let's find other ways to do it.



There is good news, which is that this explosion in opportunity is being funded by the private sector, and so for the next five years or so, the number one majors in the leading universities are computer science. This is extraordinary. At Princeton, the number one major used to be economics. Now it's actually computer science. Economics is second. My professor friends have told me that the quality of the students they are getting are the top students that are coming into the university, including compared to the physics students, the math students, the other traditional people we would compete with, the fields we would compete with for the smartest students. So there's reasons to believe that this next generation of young scientists and inventors are even better than the people, the old people that are talking about this now. Relative to that, these are 19, 20, 21 year olds. So I'm very optimistic about that.

I'm concerned that the country does not frame the problem correctly. I used immigration and research funding as two examples. There are plenty of others.

Audience Question: Jonathan Goldberger. Really appreciate the talk as well. We've talked a lot about technology, we've talked a lot about software, but this is really a people problem. I'm really interested in your perspective about the behavioral change and culture that embraces this. And how do we do that with organizations that are hundreds of years old, our government even more so, to enable that behavior change, that culture of artificial intelligence?

Eric Schmidt: So I never want to use nuclear as a good example, but if you think about the nuclear emergency in the 40s and 50s, ultimately the government created a set of institutions – these are Sandia, other labs around the country – that were run by physicists, because the physics of nuclear is very, very difficult. These are highly competent organizations full of incredibly smart people working on these very, very difficult science problems. So we do have a model there. But it was understood that it [nuclear physics] was a specialized skill. Part of the reason in our [Defense Innovation Advisory Board] recommendations create the equivalent of an AI institute is that it's not a normal stuff. That's an example.

I think that what I think will happen is that every major public corporation will have as its suppliers some of [AI experts] the people that I'm describing, whether they're separate software companies or whether they work for them directly because it's too fundamental.

I'll give you another example: if you run a normal big business in America, a team of 10 or 20 of these people can probably increase your profits by 5% just by pattern matching, because these algorithms are very, very good at looking at and discovering patterns that you don't see. My favorite example is we [at Google] have these large data centers. We care a great deal about them. We've tuned them. We have the best data center engineers in the world – according to the data center engineers. So we take our group in London, called DeepMind, and we ask them to use their reinforcement learning tools to look at simple energy efficiency of our data centers. And data centers are defined by how much electricity they use, moving power around, cooling,



you know, boring stuff. And you think its more busy on Monday than on Sunday, that kind of thing and they've tuned it as best they can. And they're really good. In two weeks, the [DeepMind-developed] system came up with a 15% improvement over the top engineers in the world. Now what's a 15% improvement in energy efficiency worth to a large distribution company in the oil & gas business or the transportation business of any kind? That's enormous. Those are all of their profits. That's how powerful this technology is.

Audience Question: Geoff Odlum, State Department. The tech sector traditionally is wary of too much regulation with the federal government, oversight, and intervention. On this particular aspect, the intersection of artificial intelligence and national security, what's your sense of the willingness of Alphabet or Facebook or Amazon or Microsoft to work proactively with the federal government especially when it comes to trying to set limits or constraints on the types of technologies you develop or the export of technologies to certain countries?

Eric Schmidt: So why would you want to put limits on the technology that I just described?

Questioner continues: To prevent an AI arms race or to prevent proliferation of weaponized AI.

Eric Schmidt: Okay, there's a lot of presumptions in that question. So the first question is do you think there's going to be an AI arms race?

Questioner continues: Yes. Well you said China will be ahead of us by 2020. Doesn't that concern you?

Eric Schmidt: It does, but it doesn't necessarily mean that its an arms race. Depends on how its done.

Questioner continues: Well you said China is interested in using AI for military purposes

Eric Schmidt: I said for commercial as well as defense purposes, which is what they've said publicly.

Questioner continues: So it doesn't concern you at all?

Eric Schmidt: It actually does, but the reason I'm provoking you is that it's so easy to go to the movie, you know the movie of robots

Questioner continues: Yeah, I'm not talking about The Terminator, but as you mentioned the example of nuclear weapons in the 40s. We might be at that stage in the next few years when it comes to AI.

Eric Schmidt: Hopefully we're not going to be. So let's go through some of the arms race questions. So the first question is cybersecurity. One of the things that is different between cybersecurity and nuclear is that we have the possibility of defending ourselves against it. So



before we get focused on how terrible things are, maybe we should fix our cyber systems to become more resilient, starting with upgrading all of those Windows 95 systems that the government uses. Last time I checked, we're still trying to get to Windows 10. Windows 95 has been thoroughly penetrated by the Chinese and every undergraduate at every university. So if you think your network is secure, trust me, it's not. There's a set of things that we can do there that are well beyond, before an arms race. I can go on.

I think the general answer on regulation is we worry about premature regulation. So there's plenty of things to worry about, but I don't know what you'd regulate now. Under that doctrine you would say "well when you invent an algorithm, you have to give it to the government." Well, at the moment if you give it to the government it would be leaked, okay, because the systems are so insecure. Again, I would be very, very careful about simple answers to this. I would prefer to "be more like us" which is to invest in our core, especially younger talent coming out of universities and allow that we will lead through technological, technology innovation and so forth and stay ahead. That would be my goal.

Paul Scharre: Do you see certain components of this technology that we might want to protect and not allow others, not proliferate, not export to others?

Eric Schmidt: It's unclear to me that would work very well, because the tech industry is moving so quickly, the understanding of the underlying math and so forth is so broad now, and there isn't a fissile material that you can sort of constrain. It's a reasonable debate to have. An example is, that, my own view, is that if anyone is going to invent general intelligence, it will be the DeepMind subsidiary of Alphabet, which is in London. Why is it in London? Because there's a brilliant collection of people who came out of Cambridge and they've been adding people and so forth. And by the way – their complaint? They want to make sure that Europeans will be able to move into Britain because there's plenty of incredibly smart Germans and French people that are worried that in Brexit might not be able to move. Same argument but in a different context.

Audience Question: Eric, thank you for sharing your views. It is really insightful for us. Your peer Elon Musk and silicon valley apparently had a more pessimistic view on AI and he publicly warned that AI is going to be a disaster to human beings, to society. So I have a very simple question maybe in a broad context. So if you look at artificial intelligence, today we call it AI, but essentially what it really does is... Taking what you shared with us: AlphaGo learned this new game in seven days. So that's not really regular intelligence, its creating super intelligence. And so in the future, in the near future, what is going to happen is you're going to have potentially country to country confrontation. Could be, you know, good guys versus bad guys confrontation. So essentially what is going to happen is really superintelligence against superintelligence. I know it is a very hard kind of question to answer, to predict what is going to happen but a simple question I have for you is: so take the DeepMind organization. You guys



have this AlphaGo. Do you have another AlphaGo within the same organization and you are trying to compete between two AI, two superintelligence. When's that going to happen?

Eric Schmidt: I understood your question, and I'm afraid I'm going to have to cast it into the movie department because we're not today building superintelligences. We're building systems that are very good at imprecise pattern matching. It's not, these are not superintelligences in the sense you're describing. There are people who think that we will eventually be able to do that but certainly not in the next few years. We're some time away. It's absolutely true that AlphaGo is a tremendous achievement. It was done against a system of absolutely clear rules, which is why games are so susceptible for it [Artificial Intelligence techniques]. It remains an open question how much farther it will go. I'm not as worried about the scenario you described. You can imagine a movie where there are these superintelligences and they have a big battle and ultimately they have a sort of combination.

And I don't think that Elon is right at all. I think that AI, in fact, is the inverse of what he said. I think it will make people enormously smarter and enormously more productive, live longer, have better lives, in fact lead to less conflict. This is my own personal view and most of the people I work with agree.

Audience Question: Richard Harknett from the University of Cincinnati. Couldn't agree more with everything that you've said and I want to take you back into that strategic space in which we have an absence of strategy. When I think strategy I see a distribution of AI cyber power. Obviously we're talking about China, we're talking about the United States. I wonder if I can draw you out and talk about Alphabet as an actor in this space. We've been caught up in this partnership model for a long time and that has very interesting implications or implies a certain type of relationship. But when we think strategically and we think about Alphabet as an actor in this space, it suggests to me a much better way of thinking about this is an alignment model, right? That Alphabet has interests in this space that may or may not be the same as China's, that may or may not be the same as the United States'. And how do you align those to achieve strategic ends that from the United States' perspective would advance national security. Does that make sense? That we should start to think about your company and others in that space as actors with interests that are going to influence and construct the area that we have to seek security and power in.

Eric Schmidt: I'm sure that we will be a contributor but it's important to remember that we [Alphabet and Google] operate under U.S. laws and U.S. regulations. The three companies that are really leading in this area are ourselves, Facebook, and a Chinese company – Baidu and there are others coming. I think that there will be a set of companies that will amass enough people and data and so forth to do interesting things with their businesses with again, others joining. But, ultimately these issues will be resolved in a higher court if you will. I will tell you that our contribution is we're trying to move this research ahead quickly. We released a library called



TensorFlow, which is one of the mashup systems that people use which has been remarkably popular.

Audience Question: Thank you Eric for the very interesting feedback. My name is Jeff Clune. I'm a professor of Artificial Intelligence and also helping to create Uber's Artificial Intelligence lab, which in many respects is following up on model that Google created. I want to talk about that model really quickly. One of the things that I find really commendable and remarkable about what Google did years ago was paying people to come and just do research within the company walls for pay, which I imagine is a tremendous expense. You have all of Jeff Dean's group in Google Brain, you have Demis' group at DeepMind and many other researchers. And so I wanted to hear a little bit of your thoughts in terms of what you would recommend to other CEOs and government officials, especially in the military, in terms of: is it worth the expense to hire to come and just do research and not focus on products in your company. And I know you mentioned the great example by Mustafa's group of RL of being applied to your own power factories, so what extent have you been justified in this amazing investment in pure researchers that are publishing and kind of doing the typical academic thing but for, on Google's payroll?

Eric Schmidt: So in our case, as an AI first company, you would want to be the inventors of some of the future of it so it makes perfect sense. I worked in research labs in my whole career. I worked at Xerox PARC, you know, 40 years ago. So I can't imagine a tech company without being able to fund either directly or through universities the basic research in the field that they take the most advantage. I make the same argument to the military. It just makes no sense to me that the military would avoid having either [AI] researchers itself or researchers that it funds, collaborations with universities, partly so they can get an internal heads up. It's just so much a better feeling when you can make a local phone call or the equivalent to somebody who is on the payroll, if you will, or in your tribe to try to understand what is really going on.

Think about, let's use the China threat in the military context. So let's say the military is sitting there doing a China analysis, and I'm sure that they do this, wouldn't it be nice if they had some folks that they trusted that they could call that could really talk about where Chinese research is, where American research is, and where differences may emerge.

Audience Question: Dr. Schmidt thanks for being with us today. I'm curious what you think the correct relationship is between the large Silicon Valley technology companies and the Department of Defense and the U.S. Intelligence Community. Please correct me if I am wrong, but my understanding is that the terms of acquisition for DeepMind actually restricted the level of collaboration between DeepMind and the U.S. national security community. I'm curious one whether those restrictions are real and two what is the proper relationship between technology companies, U.S. AI researchers and the national security community?



Eric Schmidt: I don't remember the specifics on the [DeepMind acquisition] terms but there's a general concern in the tech community of somehow the military-industrial complex using their stuff to kill people incorrectly, if you will. I don't know how to say it any other way. It comes from a, it's essentially related to the history of the Vietnam War and the founding of the tech industry. So in reality, I and the other people I am part of the [Defense Innovation Advisory] Board with are at least trying to at least allow for communications to occur.

As to whether I, for example, am conflicted in the sense that I can't suggest Alphabet things inside the military nor would I ever do that. In practice I don't have as a big concern of this because the military procurement processes and the way it operates are so different from the way the tech community operates. My guess is what will happen is that there will be tech, companies founded that are more in alignment with the mission and values of the military that will use this technology to help keep the country safe, etcetera.

A very good example is a group called DIUx which is headed by a guy named Raj Shah. DIUx by statute has the ability to invest in small tech companies for things which are helpful for the military's mission. They are for example investing in drone companies that are doing interesting things. Sensor companies that are doing interesting things. That is a great idea, and there should be lots more of that. That's the more likely scenario.

Paul Scharre: Thank you Eric. We're about out of time.

Eric Schmidt: Thank you all, and thank you for caring so much about this.

Eric Schmidt is the Executive Chairman of Alphabet, Inc. and Chair of the Defense Innovation Advisory Board. His remarks at the CNAS Artificial Intelligence and Global Security Summit were in a personal capacity and not speaking for the Department of Defense.

Paul Scharre is a senior fellow and director of the Technology and National Security Program at CNAS.

Gregory C. Allen is an adjunct fellow at CNAS.

Anthony Cho is a research associate at CNAS.

