



APRIL 29, 2020

TRANSCRIPT FROM CNAS TECHNOLOGY & NATIONAL SECURITY PROGRAM EVENT

CNAS Event Transcript: Military AI Applications

CNAS Event Transcript: Military AI Applications

I. Opening Remarks

- Paul Scharre: Okay, welcome everyone. Thank you for joining today's virtual event. I'm Paul Scharre. I am a Senior Fellow here at the Center for a New American Security, and Director of our Technology and National Security program. Very excited to have you attending today's event. I am very pleased to welcome the Honorable Robert O. Work, former Deputy Secretary of Defense.
- Paul Scharre: Bob Work has been a champion of artificial intelligence and autonomy, in his various roles in and out of government. He led the Department's Third Offset Strategy. And today we're going to have an opportunity to talk with Bob about military AI applications, about looking at how the Department is doing in adopting military AI, as well as take some questions from all of you, so thank you for joining.
- Paul Scharre: Just a brief note at the outset on format. I'm grateful that we're all able to stay connected through this kind of virtual format. Today's event is going to be a webinar. We'll have a conversation at the outset between me and Bob. I'll ask him some questions. But you have the opportunity, and are highly encouraged, to submit questions throughout, in a variety of formats.

II. Expert Discussion

- Paul Scharre: Bob, one of the things I want to do really is get your sense of, how is the Department doing in adopting AI? What is your sense of how far the Department has come along, and then where we need to go going forward?
- Robert O. Work: Well it's great to be here, Paul. I think you can look at it in one of two ways. If it's a glass half full, if you're a glass half full kind of person, there's an awful lot of activity going on. There are many, many experiments. Many, many prototyping activities. You would say, "Wow. In the last three or four years, things have really picked up."
- Robert O. Work: If you're a glass half empty kind of guy, you'll say, "Well you're saying that autonomy and AI is so important. Why isn't the Department being more aggressive?" I think that's fair. On balance, I would say right now that we haven't gotten to the point where the Department has fully committed to a future where AI and autonomy is central to their operations and thinking and concepts. I think we're moving in that direction, but I don't think we're there yet.
- Paul Scharre: It feels like we're at a place where there are a lot of true believers. You've got people in different pockets of the Department that are very excited about AI. But there's yet to be this killer app. This application that people across the Department can look at and say, "You know, we need that. We need to get that for our program, or our service." What's your take on where we are, in terms of showing the return on investment for AI?
- Robert O. Work: Yeah, that's a great point. When we started pursuing guided munitions for example, it was really easy to explain to the warfighter, "This is what guided munitions are going to do for you. Unguided munitions, almost everything you fired missed the target. We're going to give

you a new weapon, that can home in on its target, and have extremely high probability of kill. Not only that, but the accuracy is going to be independent of range.”

- Robert O. Work: You would ask the warfighter, “What do you think about that?” The warfighter would say, “I’ve got to have it. Give it to me as fast as you can.” You had Army artillery officers asking for guided artillery shells. Of course, the air-to-air guys always wanted a missile that could out-stick their opponent. You had ground-to-ground missile guys, anti-submarine guys. Everybody wanted that.
- Robert O. Work: It’s been very difficult to explain to the warfighter, “This is exactly how AI and autonomy is going to make you a better operator on the battlefield.” So, you’re right, we have to have some type of app that demonstrates to the warfighter, “This is what we mean.”
- Robert O. Work: That’s what we had hoped Project Maven would do, which was all about computer vision. It was about demonstrating to the force what computer vision could do for them. But that unfortunately has been stove piped primarily on the intelligence side. It hasn’t had a broader application to the operators.
- Robert O. Work: You’re right. I think when this thing really catches fire, is when operators of all types say, “This is the thing I have to have to make me better.”
- Paul Scharre: Yeah. I mean, it feels like on one level, the fact that AI is such a general-purpose technology, makes it harder in some respects than if it was something very discrete—a precision guided weapon—you can see what that is. It’s a more accurate bomb. That’s great. It’s going to be less likely to miss the target.
- Paul Scharre: But because AI is a much more general-purpose technology, like electricity or computers, and in so many potential applications, what does it do for you? If we sprinkle some magic AI dust on a program, on some kind of operation, what is that going to do? It has tremendous potential, in terms of improving operational efficiencies, and effectiveness, and accelerating decision cycles.
- Paul Scharre: But it feels like we’re yet at the point where we have a publicly available example of that, that you can point to people and say, “Here we had a process before AI. We added some AI. Now here’s the new process. We’ve increased operational availability. We’ve reduced cost. We’ve accelerated the decision time. We’ve done something that’s a very tangible and concrete way.” It feels like those are right around the corner, perhaps. But at least those aren’t public yet.
- Robert O. Work: Yeah, I think you’re right. As I said, there’s a lot of experimentation and prototyping going. The people who are working on those programs have seen the light, and they’re ready to go.
- Robert O. Work: The Army recently did, for example, a breach of a complex obstacle. As a former infantryman, you know that this is one of the most difficult battlefield tasks that an infantry platoon or an infantry company trains for. You have a complex obstacle with anti-tank ditches, concertina wire, mines, covered by fire. Affecting a breach against that complex obstacle requires a symphony of different actions, all coordinated together.

- Robert O. Work: The Army recently did a completely robotic breach of a complex obstacle. The robots were kludgy. The AI systems weren't all that advanced. But the Army demonstrated that you could actually do something, without putting a single human in harm's way.
- Robert O. Work: Now everybody who participated in that particular experiment is going, "Well yeah. In the future, 10 years from now, I don't want to be the human going through that giant thing. I want a robot going through. How do I get there?" Well, you need to have new, advanced AI-enabled control systems, AI-enabled communication systems, AI-enabled counter-jamming systems, to make the whole thing work.
- Robert O. Work: As I said, if you're a glass half full type of guy you're saying, "Wow. This is just really cool. There's all sorts of things that I can see in the future."
- Robert O. Work: But the other thing, Paul, I think you and I have talked about this a long time. This is all about operator trust. Because in the end, autonomy is about delegating to an entity, the ability to create their own courses of action, and choose among them. When you're doing that with a machine, you have got to establish a concrete trust between the human operator, who is tasking the machine to do something, so that the human has trust that the machine will respond the way the human is expecting...perform the task the way that they train.
- Robert O. Work: That's one of the other things that's holding us back. The operators haven't pushed the, "I believe," button yet. When they have trust in the algorithms to do what they were trained to do, since we train our algorithms just like we train our soldiers, that's when you'll really start to see forward movement quickly.
- Paul Scharre: Let's talk about this issue of trust. Because this is I think so foundational to adopting any new technology. But particularly this one. Where you can't see it, right? It's not like a new widget, that you can take a look at it, and you can say, "Okay" and kind of poke at it. It's much more of a black box.
- Paul Scharre: How do we get to both trustworthy AI systems, ones that we ought to trust? But also getting to operator trust? Getting to a place where operators are going to feel comfortable, putting their faith in this kind of system?
- Robert O. Work: Well the first thing you have to do is, start with the validation and verification of the algorithms. We're still learning how to do this with AI algorithms. You've got to be able to have a process that says, "Look you have extreme confidence that the algorithm will perform the way that you want to."
- Robert O. Work: Now in a lot of these deep learning systems, they are, like you said, they're like a black box. The algorithm may recommend an action, and it's unable to explain to you why it came to that decision. We're working very, very hard. DARPA is working very, very hard on explainable AI. We have to have a system by which we can validate and verify that the algorithms are trustworthy.
- Robert O. Work: Then you have to do a series of experimentation with operators, so the operators can actually see how the algorithms perform in a complex environment. We'll never be able to duplicate the chaos of battle completely. But we will be able to throw a lot of things at the algorithm, that it might not have been trained on, to see how it reacts.

- Robert O. Work: First, you have to verify and validate the algorithm itself. Then you have to experiment with it, to make sure that it performs the way you expect. Then you have to train with the operator, the human, so that the human has a lot of trust.
- Robert O. Work: There's a story of the 1973 Yom Kippur War, where the Israelis were surprised by both the Egyptian and Syrian attacks and were really being pressed. The U.S. shipped a whole bunch of TOW missiles to them. These are tube launched optically wire guided missiles. The Army had proven them, and the Army operators trusted them. The Army would have said, "Yeah, give me as many TOWs as I can, and I'll kill as many tanks as you put in front of me."
- Robert O. Work: But the Israelis had never trained with the TOW, and they were more inclined not to use them in battle. Because the operator didn't have the trust in the system. He didn't understand how the system would operate, and therefore they didn't use them as much.
- Robert O. Work: This is going to go doubly in spades for AI systems. Because the human is going to have to trust the machine. That's a whole, not only trust the machine to just do things like guide a torpedo, which we're all comfortable with. But trust the machine to make decisions, that the operator has delegated to the machine. That's going to take some time.
- Paul Scharre: I mean, I think that's a great point. It's one that's critical to really any military adoption of new technology. It's a process of getting comfortable with it, learning how it works. Figuring out, what can you trust it with, and what can you not trust it with? Because we're not looking for blind faith here. We're looking for understanding its capabilities and limitations. It's particularly challenging for some of these technologies.

III. Audience Q&A

- Paul Scharre: We have some great questions coming in, so I want to start turning to them. We have a question here from [Questioner]. One of the exciting things about this webinar format is, we're able to tap into folks from all around the world. Glad you were able to join. [Questioner] says, "What is your view on cooperating with allies? Whether they're NATO, Five Eyes, others, on military AI? What are your preferred partners for working with allies in military AI, and what are preferred areas for cooperating with allies on that?"
- Robert O. Work: Wow, that's a good question. Right now, the thinking is that different countries will be able to bring a lot of different capabilities and advances in AI. Even the smallest country, with a really high tech base, would be able to say, "I've got a new way to explain AI." Or, "I have a new way to validate and verify." Or, "I have this new application."
- Robert O. Work: The U.S. is very anxious to work with our allies. One of the things you're looking for also is interoperability, as much as possible. How far can we push in the interoperability of AI-enabled systems?
- Robert O. Work: Now I think, and this is just an opinion. I'm not speaking for anyone here. I think what will happen is, we'll start with our Five Eye partners. That's Canada, the UK, Australia. Paul, help me out. I'm having a brain lock here.
- Paul Scharre: Sorry, who did we get? Canada, UK, Australia -

- Robert O. Work: Canada, UK, Australia -
- Paul Scharre: New Zealand, and us.
- Robert O. Work: New Zealand, and us. Here are Five Eyes. We'll start with them, because we have all sorts of procedures. We have all sorts of technical sharing agreements already. We'll start with the Five Eyes. That's E-Y-E-S. It's just, it's called the Five Eyes, who are authorized to see a certain type of intelligence.
- Robert O. Work: But what we want to have is 50 eyes. AIs. We want to have as many as 50 partners. Because the competition in AI as it plays out, it's really going to be a competition of values also. Because we already know that authoritarian regimes will utilize AI in ways that a democratic regime would never contemplate or consider.
- Robert O. Work: We believe it is important that we have like-minded AI partners, who want to make sure that AI maintains personal privacy; is not used for state surveillance and suppression. Things like that. The more allies that we can get, the more likely that the future of AI will reflect the values of democratic nations.
- Paul Scharre: Thank you. That's great. We've got another question coming in. This one from [Questioner]. "I agree that the performance of the Department, in terms of adopting AI, is mixed. If you were to provide Secretary Esper with two recommendations, to enhance the adoption of AI into the Department, what would they be? How would these recommendations change if there were no constraints in authorities, but budget's constant?"
- Robert O. Work: Okay, there's a lot in that question. But the first thing I would do is, I would take the JAIC, the Joint AI Center, and I would put it directly under the Deputy Secretary of Defense, rather than the CIO. This assumes that the Deputy Secretary of Defense would push AI as one of the top two research priorities of the Secretary. Secretary Esper has said this several times, that AI is right at the very, very top of what he's doing.
- Robert O. Work: You have to have a top down push. You have to have someone at the upper echelons of the leadership, giving a clear signal that AI-enabled autonomous operation systems and platforms is the general direction that the Department wants to go.
- Robert O. Work: Then you have to do it within the programming and budgeting process inside the Department. One thing that I would do, if I was the Deputy Secretary, is I would say, "Okay. I'm going to withhold five percent of the entire defense budget, while you are developing your programs—Department of the Army, Department of the Air Force, and Department of the Navy. I want you to come forward with the absolute best ideas for autonomy to create a military advantage.
- Robert O. Work: I'm going to assume that you're going to go after AI for all the back-office stuff, through the efficiency side of the house with the Chief Management Officer. But I want you to come forward and give me the absolute best ideas, for how we use AI-enabled autonomy for competitive military advantage."
- Robert O. Work: After the three departments have created their programs, then they would come in and pitch the ideas to OSD. The best ideas would be given money in the program to pursue these

things. You have a top down push, and you have a programmatic sweetener, that says, "Look. We will help you get there. But we need you to develop the concept of operations, and the ideas for the platforms."

Robert O. Work: But don't just, as you said earlier Paul, don't just come to us and say, "I'm sprinkling AI on all sorts of stuff." I want to see a specific program that you believe will provide us with competitive military advantage. That's what I want to hear. We'll give the best ideas the money to pursue them.

Paul Scharre: I love the idea of an incentive structure like that. I think that that's really important for driving the right incentives inside the Department. I want to press you on one thing though. Because clearly while there are huge advantages for AI in combat applications, you said that you would assume that services would already use this for back office things. Do you think that's a fair assumption?

Paul Scharre: I mean, at least in the various battles that I was in, at a very low level, in the Pentagon, duking it out with people about programs. I think one of the most ferocious fights we ever had was about personnel codes. I won't repeat them here, but I was called some pretty nasty words in OSD, trying to talk to services about how we code personnel.

Paul Scharre: Do you think that's a fair assumption? That people would adopt those management efficiencies when they're available?

Robert O. Work: Well I think you have to split these things apart. Because the Chief Management Officer, which was created by the Congress, was specifically designed to look at the business operations of the Department and to make them more efficient. Make them more akin to modern business practices.

Robert O. Work: The Secretary has said, he wants to free up, or the Department of the Navy alone, wants to free up \$40 billion over the five-year defense plan, in efficiencies. The only way that you're going to get to that kind of money is to start automating many of the processes. I don't know how many transactions occur in the Department of Defense each day. But these transactions should all be automated. If you automate them, you'll be not only more accurate and more efficient, but you'll be a lot more cost effective.

Robert O. Work: Andrew Ng, who used to be the CEO of Baidu, and is now back in the United States, said, "If it takes a human one second to think about what to do, you should automate it." That was his rule. Many of these transactions, if it's a no brainer once the transaction is set up, you just have to say, "Yeah. Let's do the transaction."

Robert O. Work: I would say, "CMO, you've got that. You have got to push that as hard as you can." But the Deputy Secretary, in my view, should really focus on the things in the program that provide this competitive military advantage. That I think having those two parallel tracks is the way that you can have AI and autonomy have the broadest impact in the Department.

Paul Scharre: Wonderful, thanks. We have a pile of questions coming in, so I'm going to bundle a couple of them that are related here. The question from [Questioner], asked, "Thank you for your remarks. You talked about the need to overcome the black box problem, for humans to trust

the AI system to make decisions. But explainability is not the same as accountability.” That’s very true and is a key concern for a lot of people; making sure we have accountability.

- Paul Scharre: “In the event of an unfortunate accident, where do you think accountability should fall?” Now I’m going to pair that with, so hold out our mind. A related question here from [Questioner]. “What about the issue of automation bias? How do we prevent problems that may arise, when an operator places too much trust in an AI system?”
- Robert O. Work: Okay, these are both good questions. The Department of Defense has already made a determination on accountability. First, it notes that international humanitarian law does not, I repeat, does not preclude autonomous systems on the battlefield. But it makes a statement very clearly, that because a machine is not a moral agent, it cannot be held responsible for mistakes it makes. The Department’s Law of War Manual states clearly that it is the commander’s responsibility who employs the AI system, and they are accountable for the mistakes that these systems make.
- Robert O. Work: This is a key point that you have when the Campaign to Stop Killer Robots and the Department of Defense argue over these very tough, legal, ethical, and moral questions. The Department of Defense says, “Look. We’re going to have narrow AI systems. We’re going to have systems that are made for a specific battlefield task, that we can actually test. We can actually see; will it perform the task we expect? We are not pursuing general AI systems that have independent ability to make decisions on the battlefield. Everything will be in a human command and control system.”
- Robert O. Work: The commander who employs these AI systems has to understand what they were designed for. Has to understand how to use them. If the commander uses them inappropriately, and there is some type of legal or moral or ethical issue that evolves on the battlefield, it is the commander, the human, who will be held accountable.
- Robert O. Work: This goes under, in fact, Paul was one of the primary authors of DODD, Department of Defense Directive 3000.09. It says that everything we do, from the time we design the algorithm, to the time we test it, to the time we use it, there must be appropriate human judgment over the use of force. Period. All stop. We do not delegate to a machine a life or death decision on its own.
- Robert O. Work: The human might say, “I want you to attack that target.” Then the machine may do it all on its own. But the human makes the decision, what target or type of target or group of target to attack. Accountability in my view, is clearly stated in the Department of Defense. It is the human.
- Robert O. Work: Now automation biases goes to the V&V discussion we had. The validation and verification. It is a problem that everyone is aware of. We are still learning how to do V&V processes for AI enabled systems.
- Robert O. Work: But I’m quite optimistic about this. Because as I said, generally what the Department of Defense is doing, is asking automation and AI to do very narrow tasks. “This is what we want you to do, machine. We don’t want you to do anything else. This is what we have designed you to do. This is how we’re going to design it.”

- Robert O. Work: Then in the V&V process, we try to determine, “Do we have any bias? Is the algorithm operating in a way that is unexpected? How brittle is it? How easy could it be spoofed?” In the V&V process, you'll try to identify all of these things. Then you'll go through a demonstration phase, to determine whether or not the algorithm performs in real life, the way it performed in the lab.
- Robert O. Work: Then you go through a phase of trust building with the operator. Automation bias, at that point. I mean, you have three different processes, that hopefully would identify this type of bias, and would allow us to correct it.
- Paul Scharre: Okay. We have a couple more questions I'm going to bundle up here together. These are both about organizations inside the Department — “How the DOD should organize itself?” [Questioner] asks, “With so many different initiatives and prototypes being explored and tested across the DOD, how can we make sure, how could DOD improve upon translating success stories and lessons learned to other parts of the Department?” For example, there's things like Maven that are stove piped, that the lessons are limited to that particular department.
- Paul Scharre: A related question. This is a short one, but I think a good one, from [Questioner], “Who owns AI in DOD?”
- Robert O. Work: Well, let me answer the second one first. This is a personal opinion. AI is a means to an end, in my view. This was the view of the Defense Science Board, and this was the underlying premise of the Third Offset Strategy. Where you would get competitive military advantage is in autonomy. AI is the way that you get to autonomy.
- Robert O. Work: To answer [Questioner]'s question, if you're saying, “Who's responsible for the R&D?” Each of the three military departments is pursuing AI applications in their individual domains of warfare. The Navy is going after unmanned surface vessels, and unmanned underwater vessels. The Air Force is going after command and control systems, and unmanned loyal wingmen that would accompany a manned fighter. The Army is pursuing a wide variety of robotics and planning systems.
- Robert O. Work: Of course, the IC, which DOD works very closely with, is using AI and machine learning to autonomously go through mountains of heterogeneous data to identify the thing that an analyst really needs to focus on. You have responsibilities in each of the military departments, and the IC, pursuing AI-enabled applications for their own domains of warfare.
- Robert O. Work: Within the Department of Defense, it's a good question. The Department decided to put the JAIC under the CIO, Chief Information Officer. I think that made a lot of sense in the beginning, because they were worried about what they refer to as joint foundational activities. Like the DOD cloud strategy, which is absolutely critical for an AI future. DOD data strategy, absolutely critical for AI and machine learning. Algorithmic libraries, where you could store all these different libraries, so each of the military departments could avail themselves of them.

- Robert O. Work: I think in the beginning, it made sense to put it under the CIO. But if you're looking for a broad base transformation to gain military advantage, I believe it has to fall under the Deputy Secretary of Defense.
- Robert O. Work: You could make a strong case that R&D would be underneath the Undersecretary of Defense for Research and Engineering. But because of the split that the Department made, you hear R&E talking very little about AI, and CIO doing most of the talking. Now that's AI, the technology behind AI, the algorithm development, explainable AI, etc.
- Robert O. Work: But for military applications, I think the JAIC is the right place to be responsible for pushing those type of things. The JAIC just started a new program it calls Joint Warfighting. We're just working with the services to try to identify the killer apps, literally and figuratively, that would provide us with military advantage. That's the way I would say, or how I would ask [Questioner]'s question.
- Robert O. Work: I've blabbered on so long, Paul, I forget the first question. If you could repeat [Questioner]'s-
- Paul Scharre: "How do we ensure the Department is able to share lessons learned about successful AI programs, from one office or program to another?"
- Robert O. Work: Well, JAIC, the Joint AI Center, came about, because the Defense Innovation Board, under Eric Schmidt, recommended that because AI technology and autonomous operation systems and platforms were so central to the future of the Department, that there should be a center of excellence for the Department because the amount of AI literacy in the Department was low.
- Robert O. Work: That's how the JAIC came about. In hindsight, I wish they would have named it the Joint Autonomy and AI Center, to make sure that we're not focused so much on the technology per se. We're focused on the applications that would provide us with advantage.
- Robert O. Work: JAIC 2.0, they now have completely restructured. JAIC 2.0 has brought in a new AI dev-sec ops process, which mirrors the development security ops development of software and algorithms in the high-tech sector in the valley. They are bringing that in, because in JAIC's view, that is the only way to scale these applications across the Department.
- Robert O. Work: I'm extremely excited about the new JAIC business model, JAIC 2.0. We don't have enough time to go through the whole thing, and I'm not the right one to talk about it. But it strikes me as exactly the right thing now, to get to what you're saying—knowing these are the best algorithms and getting the algorithms into a common algorithmic library that then anybody could use.
- Robert O. Work: The hope is that, the Navy uses an algorithm that is really good, or excuse me, the Air Force does an algorithm that's really good in doing swarm logic, and swarm control, of a whole bunch of unmanned aerial systems. They create this algorithm, and they put it in the algorithmic library. Then JAIC tells the Navy, "You know what? With some modification, this is going to be able to control swarms of unmanned underwater vehicles. Rather than developing your own in a long process, let's start with the Air Force algorithm, and modify it to your needs."

- Robert O. Work: That's kind of the thinking of how we'll get this started. I believe, again, I am a glass half full type of person. I wish we were going faster. But things are starting to fall into place. You have a Secretary who said, "AI and autonomy is among the very, very top things that I am personally worried about." You've got JAIC 2.0 going on. You've got all sorts of activity going on in the Department, at all different levels.
- Robert O. Work: I don't think it's going to be too much longer before things start to knit together, and you start to see an acceleration of activity.
- Paul Scharre: On a similar vein, I've got a question here about the ability to scale up, in a big way, AI applications. Got a question from [Questioner]. "Given that traditional vendors are ill suited to develop critical AI applications, how should the Department go about awarding significant program dollars to new entrants, with better access to talent? It seems like this has been a real struggle in the past. For example, with companies having to sue the government."
- Paul Scharre: I'll just add onto this, we've seen DOD is able to tap into smaller companies that can do some really innovative things with AI. We've seen this with Maven, and other projects. But once you start scaling up the dollars, that's when it starts to get challenging. What's your take on this issue of getting to AI at scale?
- Robert O. Work: Yeah, this is another thing where the words are all there. The Department says, "We have to tap into the commercial sector talent." We want to have a vibrant AI autonomous ecosystem, of smart, aggressive, innovative, cutting edge companies. But I've worked with a lot of these companies, and we have not broken the code.
- Robert O. Work: In every case, there may be some ... Not in every case. Never say never, never say always. In most cases of the companies that I work with, what happens is, they're saying, "You've got a great idea. Come in and brief me." They come in and brief the guy, or the gal. Then they say, "This is great. I've got you to brief my boss. You have to brief my boss. Come back in two weeks, because it's going to take us two weeks to get on his schedule, and we'll brief the boss."
- Robert O. Work: They go back, and they come back two weeks later. They brief the boss. The boss says, "Wow. This is totally cool. I've got to get you to my resource sponsor. Come back in three weeks. Because this is a three-star, and it's really hard to get on his schedule." These are small companies. Every one of these delays is a killer and redoing all of these briefs are a killer. So, the consensus of the companies that I advise, or I work with, is that the Department still really isn't serious about breaking open to get there.
- Robert O. Work: There are pockets of things like DIU, the Defense Innovation Unit, SCO, the Strategic Capabilities Office, the Rapid Capability Offices in all the services—they are oasis of innovative contracting, and innovative programming, in what I consider to be a desert right now. I really feel for a lot of these small companies, who really want to help the Department of Defense, but they haven't been able to do so, for a wide variety of reasons.
- Paul Scharre: Yeah, this is a pain point I hear continuously as well. Let's turn to a question from [Questioner] from [Organization]. "As we look to operationalize AI capabilities in the

military, what areas do you see as producing results in the near term? Say in the next five years?”

- Robert O. Work: Well, there's a lot of stuff going on in autonomous weapon development. I think right now, the focus is on cooperative swarms of weapons. Say you fire seven weapons at an enemy target, and all of the weapons start to talk with each other. One of the weapons notifies the others, after looking at what's happening, “I am going to go high.” I'm speaking of these as if they're people, so forgive me.
- Robert O. Work: But one of the missiles says, “Hey guys, I'm going to go high, and I'm going to turn on my active radar. I will pass all of my radar data to you. But you all keep silent. Because as soon as I activate my radar, I'm going to be a big fat target.” Then another of the missiles says, “I will come in from the west.” Another missile says, “I'll come in from the east.” Another missile says, “I'll go high and dive.” Another missile says, “I'll stay on the deck.” All of it designed to completely flummox the defense.
- Robert O. Work: Within five years, we're already testing this kind of stuff. Within five years, I would think that we would have a lot more of it. What that means is, you have to fire fewer weapons at a specific target to achieve effects. That helps you in many, many ways.
- Robert O. Work: There's all sorts of unmanned control logic that is going very, very fast. The Navy has unmanned underwater vehicles, unmanned surface vehicles. The Air Force has loyal wingmen, and all sorts of unmanned attritable systems, systems that are relatively low cost that they launch just once. I think you'll see a lot more of that activity going.
- Robert O. Work: Interestingly enough, one of the things that's holding us back is Congress. For example, the Navy would like to have what they call medium unmanned surface vessels. An unmanned ship, that's pretty big, maybe 150 feet, and says, “We'd like to count this as part of our Navy, because it's going to do all sorts of stuff for us.” But Congress says, “No, you can't count an unmanned ship as part of the battle force yet. We're not convinced.”
- Robert O. Work: What is the Navy to do? This is really frustrating. Congress says, “I'm not going to give you money for this unmanned ship until you give me the concept.” The Navy says, “Well I have to have some of these ships so I can develop the concept.” We're just caught in this endless duo.
- Robert O. Work: I also think you're going to see a lot, well, you're going to see a lot of computer vision systems. I think what will happen is, they'll go on the sensor themselves. You'll put AI machine learning logic on the sensors, so they will be able to do some of the processing right there and will only have to send back the information that the human operator asks for. This helps you a lot, because you don't have to send back so much information through your pipes. It helps you in the communications electronic warfare fight.
- Robert O. Work: I think we're getting close to planning systems, that will help commanders develop courses of action. That might be probably to the latter part of the five-year period you asked about. I think you're going to see robotics, in all sorts of different things. As I said, there's a lot of activity going on.

- Robert O. Work: To Paul's point, you haven't seen, right now I think in '21 if I remember it right, it was about \$750 million in AI research—explainable AI, V&V, stuff like that. But there was another \$1.7 billion, \$1.8 billion in autonomy—autonomous airplanes, autonomous ships, autonomous underwater vehicles. That seems to me to be kind of the right ratio, in the Department of Defense. You want the Department of Defense focused on applications.
- Robert O. Work: But that's not a lot of money in the big scheme of things. When you have a \$740 billion budget, and you've got about \$2.5 billion in a given year on these technologies, I'd like to see that go up by at least an order of magnitude.
- Paul Scharre: Well, and this is such an interesting thing about these priorities. Esper has said AI is his number one priority. Setting aside the fact that that's not what R&E is saying, and they have an ever-changing list of priorities. When you look at the dollars, it's just not there. In no way, shape, or form do the dollars add up to AI being the number one priority.
- Robert O. Work: I agree.
- Paul Scharre: We've got another question, from [Questioner]. “How do you think about adversarial examples and operational risk?” For folks who aren't familiar, these are some of the spoofing attacks that can be done against deep learning systems. We can feed them tailored, very specifically tailored data inputs, to then manipulate their output, for things like, say image classifiers.
- Paul Scharre: Specifically, [Questioner] asks, “Separate from legal or ethical concerns, is it wise for DOD to start to build up and rely on these AI-based systems, in their concepts of operation? What will happen when deep pocketed adversaries deceive these AI systems, or break them?”
- Robert O. Work: [Questioner], this is a question that is being debated within the Department daily. Just how brittle are these systems? How susceptible are they to spoofing, etc.? Again, I'll answer it in two different ways.
- Robert O. Work: First, I equate what's happening in AI, in the technology of AI, and in the applications of the technology, a lot in the way that naval aviation was pursued in the interwar period. Everyone knew that aviation was going to have some type of big impact of naval warfare at sea. You had some people who thought, “All they're going to do is, they're going to scout for the battle line, for the battleships. They will spot the shot, and help the battleships put rounds on targets.”
- Robert O. Work: You had a group that said, “Well, you're going to have airplanes coming after the battleships to attack them, etc. You're going to have to have some way to provide them with air cover.” Then you had this one group of people that said, “We're going to use naval aviation as an independent attack arm of the U.S. Navy.”
- Robert O. Work: There were all sorts of arguments. People were saying, “No way, aviation, it's always going to be the battleship, blah blah blah.” But the U.S. Navy allowed an insurgency to thrive. This was a battleship navy, and the battleship admiral said, “Look. I don't know if they're right. But I'm not going to stop them.” They encouraged an insurgency, and they started to test.

- Robert O. Work: The technology for aviation was changing almost weekly. You'd build an airplane, and the airplane would be obsolete within a month. Because you'd have a more powerful engine, or a better plan form that would allow you to have much longer range, and a much higher payload. I'll call this—technology approach.
- Robert O. Work: Meanwhile, you have the technology. Then you give it to the testers. The testers say, “I think an application would be a dive bomber. I think another application would be a torpedo bomber. Another application would be a fighter plane.” Then you have this experimentation going on in the fleet, and they're testing it all. They're learning as they go.
- Robert O. Work: They didn't really have a perfect sight picture on exactly where they would end up. They just rode the wave. On December 7th, 1941, when the battle line was on the bottom of Pearl Harbor, everybody could say, “Well we have a ready solution.” I think AI is the same way. The technology is moving so fast, that it is outpacing any conceptual work that we can build. It's hard for anybody to explain to an operator how AI is going to help them in their immediate battlefield problem.
- Robert O. Work: It's much harder to say 10 years from now, when you put all this together, “This is how the force will operate.” It's just hard. We've got to learn. To me, this is like ride the wave. The Department says, “Look. We're not exactly certain how AI and autonomy is going to end up. But we know it's going to be a transformational effect. We've hit that, 'I believe,' button, and we're going to pursue it aggressively. We're going to ride the wave. We're going to look at how we protect our algorithms, and make sure that one way to do this is start with narrow AI type of approaches.”
- Robert O. Work: Specific approaches that you can test, and you would know immediately, if you trained an AI to say, “I want you to pick out this type of target in the scene. This is the target I'm designing you for.” You test it, and you say, “The algorithm is picking up all sorts of funny things. It's not picking up the targets that we expected it to.” Then you fix it in the V&V process.
- Robert O. Work: Then you take it into testing, and you try to figure out how brittle. You try to test. You stress test it in cyber. You stress test it in different types of data, and how you corrupt the data. You try to solve all of that. This is exactly what the Navy did in the interwar period. They were learning the technology, and how you would apply the technology, as they go. Then you would experiment, and then you would say, “Here is our first concept.” You would continually refine that concept.
- Robert O. Work: So, I don't have a specific answer. Everyone is concerned about the brittleness and the vulnerability of AI. At the same token, they say, “If we can solve these technological issues,” and places like DARPA, the Defense Advanced Research Project Agency, says, “Look. These issues are solvable, we think.” There will always be vulnerabilities that we have to think about. But DARPA believes that you can get to explainable AI, and DARPA believes that you can have protected AI systems. We have to prove that.
- Robert O. Work: Am I worried about it? Yes. But I'm more worried about us not pursuing this technology for competitive military advantage. Because it is clear that our competitors are doing it. You don't want to be on the back end of a military technical revolution. All you have to do is ask the Iraqi army how that felt.

- Robert O. Work: It's important that we continue to move, we continue to experiment. We continue to take risks. In my view, you're going to wind up in a future where the joint force will operate in a much more efficient and effective manner.
- Paul Scharre: That's a good transition to the next question here. We have yet to talk about China. But obviously China's developments, across their military really, in modernization, their militarization of the South China Sea, their military buildup. But of course, in AI as well, are a huge motivating factor here.
- Paul Scharre: [Questioner] asks, "China has plans to use AI to drive its military technology. It has signed on to international efforts to ban autonomous weapons, but with a rider," with a number of caveats, really, "saying they have promised not to use autonomous weapons, but they will continue the development. Of course, they're developing military AI in a variety of areas. It sounds a lot like their commitment not to militarize the South China Sea, until the day they did. What's your take on China's military AI development?"
- Robert O. Work: Well it's very difficult. In the Cold War, we could fly satellites over the Soviet Union. We could count the number of bombers on their airfields. We could count the number of silos in their missile fields. We had all sorts of technical intelligence about the types of missiles they were creating. It's very hard to fly over China and say, "I wonder what their algorithms are doing." Because all of these will be tested of course in the lab, and they'll be tested in exercises that we might not be able to see.
- Robert O. Work: Unfortunately, we won't know for certain who is in the lead in the algorithmic competition, until you're in a fight. You're in a fight in which algorithms are actually playing a big part.
- Robert O. Work: I am worried about China. Because China, I just want to read something here. "China believes AI-enabled operations is the way they will leapfrog the United States, to become the leading world military power." They refer to this as "Intelligentized Combat," or intelligent combat operations. This is what they described it as.
- Robert O. Work: "Intelligent combat operations refers to combat operations conducted with intelligent weapons and equipment platforms, using artificial intelligence as the core and with the technical support from information networks, big data, cloud computing, the internet of things, and intelligent control." That's the way they feel. That is the way they are approaching.
- Robert O. Work: They believe that the future, superiority in war, they actually say this – "the mechanism for securing victory in intelligent combat operations is more represented by intelligent autonomy." The first thing, that superior algorithms will generate superiority in war. That's the way the Chinese are thinking about it.
- Robert O. Work: Now, that's what they're writing. We believe, at the R&D level, that the U.S. still enjoys a lead in AI R&D. But that lead continues to shrink. The Chinese are really catching up on the basic R&D fast. There is a sense in the commercial world that the Chinese are ahead in applications of these technologies. That's what worries us. I mean, worries the Department of Defense. Are they ahead in the application of these capabilities?

- Robert O. Work: Their entire thinking about future war puts AI and autonomy at the central core of how they will achieve military technical superiority. We're in a competition, without a doubt. It's a competition that we can't really grade with any type of accuracy. So we look for hints in the intelligence record, and we look at their exercises, and see if their weapons are operating in ways that we can't explain.
- Robert O. Work: But yeah, this is a competition. There's an old saying for politicians—"run like you're losing." In my view, the Department of Defense has to run like they're losing, in this competition for an AI-enabled and autonomous future.
- Paul Scharre: Well thanks, Bob. We are unfortunately out of time. We have a whole pile of really great questions that we're just not going to be able to get to today. Thank you everyone who submitted questions. If I butchered your name, I apologize. I can assure you, with my last name, I get that. But thank you, Bob, for joining us. Thank you everyone for dialing in. This has been a great discussion, Bob. A lot of great work the DOD is doing on AI, and a really important challenge ahead. Thank you for coming and talking with us.
- Robert O. Work: It was my pleasure. It was a great conversation, and I, like you, would like to compliment the questioners. I thought the questions were great.
- Paul Scharre: Great questions, great questions. Well wonderful. Thanks, Bob. Thanks everyone for joining. Take care and stay safe.