



Center for a
New American
Security

JANUARY 10, 2020

TRANSCRIPT FROM CNAS TECHNOLOGY & NATIONAL SECURITY PROGRAM EVENT

U.S. AI Strategy Event: “The American AI Century: A Blueprint for Action”

Transcript from U.S. AI Strategy Event: “The American AI Century: A Blueprint for Action”

I. Welcome Remarks

- Paul Scharre: Thanks everybody for coming. I'm Paul Scharre, I'm a Senior Fellow here at CNAS and Director of the Technology and National Security Program. I want to welcome you all to today's event on America's AI Strategy.
- Paul Scharre: I am very excited for the lineup we have today. We have some great speakers from CNAS, we have outside experts from the White House and the National Security Commission on AI. Before we get to them, I just want to say a few words about where we are as a nation in thinking through how we respond to the AI revolution. Two-and-a-half years ago here at CNAS we launched a new initiative on American AI and Global Security to try and better understand how the AI revolution was changing global peace and security, and what the US needed to do to respond.
- Paul Scharre: Now, it wasn't that long ago where if you were running around Washington talking about artificial intelligence, people would think you'd been watching too many science fiction films. That changed a lot in just the past few years, and now we see a growing awareness across Washington, in the White House, Congress and different agencies about the importance of this technology, the revolution we've seen in the past two years in machine learning, and how this is changing different elements of society and national security as well. And that the US needs to find ways to respond and adapt.
- Paul Scharre: It's important that the US maintains a leadership position in this technology, not just so that we can stay ahead of potential competitors, but also to shape how this technology is used around the world. I'm very excited by the report that we have today that we're talking about. The American AI Century that lays out a blueprint for how to implement an effective national strategy to maintain American leadership in artificial intelligence. It is 25 specific actionable recommendations across a range of areas, R&D, talent, standard setting and other things.
- Paul Scharre: I especially want to acknowledge the authors Martijn Rasser, Megan Lamberth and Ainikki Riikonen for their role in developing this report, and in particular in getting to this level of detailed recommendations. It's all too easy to say things like, "We should do more AI." I think people actually get that. The challenge is figuring out, now what do we do next? What are the specific steps that we need to take as a nation to maintain our leadership position? And we see a lot of action across the space with policymakers, including just this week from the White House with new regulatory action on AI. We see Congress taking action, with not only different legislative proposals in the works, but creating the National Security Commission on AI.
- Paul Scharre: So, we're hopeful that this set of policy recommendations in the report will be a starting point for a conversation among leaders about how we move forward from here. Just a short note on the format of this content. You will have found on your chair these very short handouts. This is not the full report, okay, with all this detail and it's like a few pages. This is a summary. We are now doing this for some reports at CNAS. For this report, this is the only thing that we will print in hard copy, it's just this very short summary. It has a high-level summary of all the recommendations. The full report's available online. There's a link to it

on the back. So you can find the full report online, in web, in PDF format, and it should be reasonably friendly on mobile hopefully across different devices.

Paul Scharre: This is now the sixth report that we've done on artificial intelligence as part of this project. You can find all of the reports online at cnas.org/ai, so very easy to remember, and you'll find hard copies of some of these other reports like this one for example, around outside as well as other CNAS reports on related tech topics. I want to acknowledge the members of the CNAS AI Task Force who helped contribute to this report and our other work on AI. We are very grateful for their support and contributions to our thinking on this topic. We've been able to bring together a diverse array of stakeholders from industry, academia, other nongovernmental organizations and former policymakers to deliberate on many of these topics that are in the report and have helped inform how we think about these things, so I'm very grateful for their insights that have contributed to this and other AI work here at CNAS.

Paul Scharre: I also want to say it's been my personal pleasure in the last two years to have the good fortune to hire some incredibly talented people, many of whom we'll hear from later today, as we've grown our team in response to this work, so I just want to say thanks in advance to all of them for their hard work in putting together this report, today's event, and all of the thought leadership issues relating to national security and technology.

Paul Scharre: I'm now pleased to welcome former Deputy Secretary of Defense, the Honorable Robert O. Work, who will give us some remarks about artificial intelligence. No one in the defense and national security community has been a more vocal advocate of AI than Bob Work. Six years ago, when Bob was actually the CEO here at CNAS, he coauthored the monograph called Preparing For War in the Robotic Age. Also we have some hard copies out there. And it remains a really important and forward thinking call to adapting to this technology and its importance across defense and other aspects of national security.

Paul Scharre: Of course as Deputy Secretary of Defense, Bob launched the third offset strategy and founded the Algorithmic Warfare Cross Functional Team - Project Maven, important first steps in getting the Pentagon to adopt this kind of technology, and since leaving DoD Bob has continued to be a thought leader both formally and informally on this topic, including in his role as the Vice Chair of the National Security Commission on AI. So please join me in welcoming Secretary Work.

Paul Scharre: Thanks, Bob.

II. Opening Remarks

Robert Work: Thanks, Paul. Paul has hired an enormous talent pool here at CNAS, but one of the smartest things I did when I was the CEO here was to hire Paul Scharre to lead the program, and I'm so proud and happy the way it's turned out. Now I read this week that Samsung just debuted project Neon at CES which is the world's largest tech conference. It introduces an advanced new type of AI empowered digital human avatars that "can autonomously create new expressions, new movements, new dialogue, completely different from the originally captured data." And purportedly so realistically that you can't tell the difference between the avatar and a real human. It passes the Turing Test.

- Robert Work: Now these new digital humans, as they are being referred to, can be used for entertainment in business purposes, they're actors' guides, receptionists, and many, many more. I think those of you who know me know that I'm a big movie fan and someday they'll be tied with holograms and humans will develop relationships with their guide, much like Officer K and Joi did in Blade Runner 2049. At least I saw it, I'm not certain anybody else did. No.
- Robert Work: Another article I read this week reports how AI-empowered retina scans will be able to determine signs for Alzheimer, and it's just another indicator of how AI is going to have an enormous impact on the human condition, being able to determine, through micro-expressions, depression, being able to find tumors far faster and more accurately than humans have ever done. And in a new federal rule change published Monday, classified software specifically designed to train deep-learning neural networks on the analysis of geospatial imagery has been designated a dual-use technology and will be subject to many of the same restrictions for exporting arms.
- Robert Work: Now as the three examples suggest, and every week I read articles on AI, and every week there's something new and astounding. It's just moving at such a fast pace. It just indicates how deep of an impact that artificial intelligence is going to have on our lives, our health and national security. And it's hard to imagine a technology, we tried in the Department of Defense, so is there any other technology which will have as broad of an impact? Now we concluded that synthetic biology and genomics will have a very broad impact on the human condition, but certainly no technology that we could identify would have as broad of an impact on national security as artificial intelligence. And the data behind this is just starting to build and build and build. I take this from Stanford's human-centered AI project that gives an annual index. If you haven't seen it, it's quite good. It gives you highlights of reports every year.
- Robert Work: Here are the 2019 results. Between '98 and 2018, over a 20-year period, the volume of peer review data AI papers has grown by more than 300%. China now publishes as many AI journal and conference papers per year as Europe after surpassing the United States in 2006. The Field-Weighted Citation Impact of US publications is still about 50% higher than China's, so it is the quality of the research in the United States that is still leading in this area. Almost 30% of all world AI journals are attributed to East Asia. 40% are attributed to North America. And North America accounts for about 60% of all AI patent activity on the globe, between 2014 and '18.
- Robert Work: So the story is that the US is the leader in research right now, but the world is catching up, especially China. Attendance in AI conferences such as this and many, many that are even larger, continues to grow at an annual growth rate of about 30%. Prior to 2012, AI results tracked very closely the Moore's law with compute doubling every two years. Post 2012, which is where AI and machine learning has really started to grow, compute power has been doubling every 3.4 months, and as a result in just the last year-and-a-half the time required to train a large image classification system on cloud infrastructure has fallen from three hours in October 2017 to 88 seconds in July 2019. So you train your image classifier in 88 seconds.
- Robert Work: During the same period the cost to train such a system has fallen similarly, and the same thing is happening in natural language processing which is going to allow humans to interact with AI empowered systems in a much more convenient and personal way. Now the effects really haven't had a major impact on our economy yet. They're just starting. From 2010 to

2019 the share of jobs and AI-related topics increased five times, with machine learning leading. But there's still a very small percentage of all US jobs, about 1.32%. But this represents just the leading edge of this massive wave.

Robert Work: In 2019 global private AI investment was over \$70 billion with AI-related start-up investments of \$37 billion. Merger and acquisition was \$34 billion, IPOs \$5 billion, and minority stake values around \$2 billion. AI start-ups continue their steady ascent from a total of about \$1.3 billion raised in 2010 to over \$40 billion in 2018, and funding has increased in an average annual growth rate of 48%. Autonomous vehicles are the things that are driving most of these use cases now, with drug, cancer, and therapy right behind it, facial recognition video content and fraud detection and finance.

Robert Work: So, even though things are starting to build up, the impact on our lives still is just a small portion of what is going to happen in the future. Now, for all of you here, 58% of all large companies surveyed reported adopting AI in at least one function or business unit in 2019 which is up 47% from just 2018. So, businesses are starting to pursue AI applications. But we have a long way to go. 19% of all large companies say they're taking steps to mitigate risk associated with explainability and reliability, only 19%, so we still have a long way to go before we can have AI that is reliable, explainable, repeatable. Long way to go. And we have a lot of bias in our algorithms.

Robert Work: And the other thing besides all this activity and investment, what we're seeing in business, education, the graduate level AI has become the most popular specialization among computer science PhDs in North America, with twice as many students as the second most popular which is security and information assurance, which was the big thing last decade. Over 21% of all graduating computer science PhDs specialize in AI and machine learning, and in the US and Canada the number of international PhD students continues to increase, exceeding 60% of all the PhDs produced. We'll talk about that probably later on the panel.

Robert Work: So, it's just amazing what is happening right before our eyes. Now the White House, things are happening and the United States is finally starting to get its collective act together. You have the White House which is really starting to lead the charge on how AI can be used throughout the United States to make us healthier, safer and more productive and more economically viable. The National Commission just released its interim report and by design, because we do not publish, or we do not roll out our final report until March 2021, when we will produce that report for the new administration, whichever party it is. We specifically wanted in our interim report to say, these are the issues that we are wrestling with, and we specifically decided not to have a lot of recommendations.

Robert Work: Meanwhile, the report here from the AI Task Force and the AI group in the technology program really focused on hard-hitting recommendations so that we can start to talk with Congress, and there's a lot of interaction between our work and there's a lot of interaction between the Commission and Paul's team and the White House. But I thought I'd read something that Michael Kratsios said at the Center for Data Innovation event on September 10, 2019, because I think it perfectly captures what the White House is trying to do at the upper level, what the National Commission is trying to do, and what CNAS is trying to do.

Robert Work: "So today our goal is very clear. This uniquely American ecosystem must do everything in its collective power to keep America's lead in the AI race and build on our success. Our future

rests on getting AI right. AI will support the jobs of the future. It is, and will continue to drive our economic growth. It's advancing our national security and it is improving our daily lives. When we lead in AI we will drive our free and prosperous future. Authoritarian nations look at new technologies as another way to control their people, using AI to surveil their population, limit free speech and violate fundamental rights. This is not the America way. Our vision for artificial intelligence is rooted in the rule of law, respect for rights, and spirit of freedom. Our holistic strategy will improve our development of AI, empower the American people, promote innovated uses of new technology, and stay true to our values and preserve our dominance."

Robert Work: And the title of this CNAS report, the American AI Century, is completely in line with the direction that the White House has set, and is also completely in line with the way the Commission is pursuing this because we are unabashed saying, "We want America to win the AI technology race." We believe it is fundamentally important to set the rules, the global norms and standards and rules that Paul talked about, so that AI is used as a force for good in our future.

Robert Work: The other thing I really like about the CNAS report is it truly is a blueprint for action. It is, I think it was 28 pages in the final report, but it's focused on recommendations. There's not a lot of flowery language. It says, "Let's roll up our sleeves and get after it. And here's what we do as the first steps." And so, I really believe that it is an important addition, everything that the White House, the Commission and CNAS is trying to do. I'd encourage you all to read it and embrace and internalize what it's trying to do, and I very much look forward to talking with you as part of the panel this morning. Thank you.

III. Lightning Talks

Martijn Rasser: Thank you very much, Bob. Good morning everyone. My name is Martijn Rasser. I'm a Senior Fellow here at the Technology and National Security Program at CNAS. I want to take just a moment and thank Paul for his vision and leadership, not just for this report but for his overall efforts to ensure there is informed debate on the impact of artificial intelligence on US national security. It's a real privilege to be a part of the AI and Global Security Initiative here at CNAS, which Paul spearheaded, and it's a big reason why I enjoy working here so much.

Martijn Rasser: Another aspect of CNAS that I particularly enjoy is the Center's commitment to nurturing the next generation of national security leaders. It's one of the most important things we do. You'll hear from two of those rising leaders this morning, Megan Lamberth and Ainikki Riikonen. I'm fortunate to have them as colleagues, teammates, fellow authors, and I'm excited that they have the opportunity to share their insight with you. Megan and Ainikki will each do a lightning talk on one of the important topics that we've raised in our report, and they'll offer actionable policy recommendations. Their recommendations are essential to crafting an effective national AI strategy.

Martijn Rasser: Megan, the floor is yours.

Megan Lamberth: Hi, thank you all so much for being here today, and thank you to Martijn for the kind introduction. In this era of global technology competition in artificial intelligence, the US government has shown commitment to developing AI systems that can positively transform

the country's economy and national security. Yet the government is also neglecting one of its key sources of competitiveness: talent. Only with the right talent and expertise will the United States continue making groundbreaking advancements in AI, and while some of the country's AI experts are US born, high-skilled immigrants represent an indispensable component of the country's tech ecosystem.

- Megan Lamberth: In our recent CNAS report, The American AI Century, we argue for why the US must protect and maintain its immigrant talent base for continued AI advancement, and we explain what the US should do to attract and retain the best AI talent in the world. Exact numbers vary, but it's clear a shortage of AI talent exists. There simply are not enough STEM educated Americans to fill the ever growing number of AI-related jobs. And even if the US were to adopt a fully funded, long-term educational initiative to train US students in STEM, this kind of effort would take decades to realize.
- Megan Lamberth: But high school immigrants can be employed today. Yet current immigration pathways for talented foreign nationals are often complicated and expensive. As the need for computer scientists and technologists continues to rise, tech companies have increasingly relied on the H1B visa program to recruit foreign talent. Since 2005, the cap on H1Bs has remained at 85,000 per year, with 20,000 of those visas earmarked for applicants with advanced degrees. But the number of H1B applications have far exceeded this cap for the past 16 years, peaking at 236,000 applications in 2017. And keeping the cap low is designed to protect American workers, but the number of AI-related jobs far outweigh the current number of job seekers and this divide will only widen into the future.
- Megan Lamberth: America's current immigration policies represent a clear example of government policies killing American innovation. If the United States wants to ensure its long-term competitiveness in AI, it must make changes to its current immigration pathways now. First Congress should reform the H1B visa process by raising the overall cap of available H1Bs and removing the cap entirely for applicants with advanced degrees. Second, Congress should simplify the process of applying for an H1B to make it easier for smaller tech companies and start-ups to hire talent. The process of applying for an H1B is expensive and requires extensive documentation from the potential applicant's employer.
- Megan Lamberth: Start-ups simply don't have the personnel nor the resources to be very competitive in the H1B lottery, therefore Congress should also earmark a percentage of available H1Bs for smaller tech companies. And in addition to reforming the H1B process, Congress or the White House should create a new immigration pathway to retain international students studying on an F1 visa in the US. And to apply for this program an international student would first need three things. Acceptance into a pre-approved AI-related graduate program, successful completion of the government's screening and vetting process, and a commitment to work in the United States in an AI relevant field for a minimum of 10 years after graduation.
- Megan Lamberth: And after completing graduate school the program participant would be granted an open market EB-1 green card which essentially means that the individual could work for any US employer. And after 10 years of employment the individual would be granted unconditional permanent residency or citizenship. And a program like this would not only bring international students to the United States, it would attract talent that already have a desire to study and to work and to live in the US long-term. This program would also help smaller

tech companies in their recruitment efforts addressing the so-called "startup visa problem" by eliminating some of the expense and uncertainty these companies experience when extending job offers.

Megan Lamberth: The US cannot afford to close its doors to international talent. If the US wants to maintain its leadership in artificial intelligence it needs human capital. It needs immigrants. Talented foreign nationals want to come to this country. America should afford them every opportunity to do so.

Megan Lamberth: Thank you. Happy to hand things over to my colleague, Ainikki.

Ainikki Riikonen: Thanks, Megan. Good morning, everyone.

Audience: Good morning.

Ainikki Riikonen: Thank you. Thank you for being here today. Martijn, thank you for the introduction. As Megan said, people are the drivers of the AI revolution. We need human talent—and that means the best and the brightest—to build the technology and to implement it. People are the engine of innovation, yet people can also steal technology. They can commit industrial espionage which is theft of secrets on behalf of private industry. They can also commit what we call economic espionage or theft on behalf of nation-states.

Ainikki Riikonen: The US Trade Representative estimates that trade secrets theft can add up to a whopping \$600 billion every year in the US. That's over 100 times the size of DoD's AI research budget, just for scale. So we need to find ways to protect our technology and our competitive edge. So how do we balance competing issues? On the one hand, the openness of our society has been the bedrock of American innovation, and I would say that innovative thinking really requires diversity to work and to reach its maximum.

Ainikki Riikonen: On the other hand, malign actors want to take advantage of our openness. Historically these have been both friends and foes, and they range from France to Japan to Russia. Today it's China. China has a dedicated apparatus for systematically stealing technology and it's doing this on a massive scale. So, how we keep our doors open while keeping bad actors out, and how can we do this without collateral damage to our fellow Americans? Asian-Americans are already suffering impacts from some of our efforts, for instance, scientists being wrongly accused of spying.

Ainikki Riikonen: To combat illicit technology transfer, the US government should stop malign actors before they come in through our doors. It should better equip small firms to defend against cyber espionage. The government should also boost its collaboration with academia. Collaboration includes spreading best practices, mitigating specific methods that are behind technology transfer, and also identifying the technologies most at risk for theft. So, allow me to expand on these three recommendations.

Ainikki Riikonen: First, Congress should authorize consular officials to act on risk indicators for espionage. An example of a risk indicator is whether an individual is funded by China's government on their visit here. According to the Australian Strategic Policy Institute, 500 military scientists from China have come to the US since 2007. The US government can do more to stop these individuals before they arrive, because these exchanges have limited value at best.

- Ainikki Riikonen: Second, the US government should provide more cyber defense support for small firms. Congress should increase funding for this effort. The Department of Homeland Security is building its cyber workforce and it should implement resources specifically for protecting the US innovation base.
- Ainikki Riikonen: Third, the US government should boost collaboration with universities. Some universities already have some good practices to prevent technology transfer such as travel briefings to researchers going overseas, and also conflict of interest reporting. Conflict of interest reporting defangs China's Thousand Talents Program. The program doesn't want you just because you're smart. They want to know actually what you know. Travel briefings could also highlight that while there's no free lunch in Washington—perhaps some breakfast sandwiches courtesy of CNAS—but certainly no free dinner to be found in Beijing. A briefing might also give broader advice like, "Please don't pick up the free flash drives at the conference."
- Ainikki Riikonen: The new White House Joint Committee on Research Environments is helping universities to proliferate some of these best practices like travel briefings and like conflict of interest reporting, and it should continue to do so. Other agencies can also contribute to this collaboration by lending their specific areas of expertise. The Department of Commerce already trains universities on export control regulations, for instance. The State Department can borrow from this model and train researchers to spot dual-use applications for their AI research. Flagging dual-use technologies, which have a heightened risk for theft, can help researchers to protect their work.
- Ainikki Riikonen: A positive agenda for American AI leadership means cultivating human talent and the brightest minds. It means protecting innovation. It means keeping malign actors out and working with the research community to protect both the technology and also the openness of our society. Creating this balance will be key to American leadership in this technological revolution. Thank you.

IV. Panel Discussion

- Martijn Rasser: Thank you very much, Ainikki and Megan. Now, I'd like to invite our panelists up on the stage. Lynne, Olivia, Bob, please come and join me.
- Martijn Rasser: Well, it's great to have such accomplished people here. Welcome. Let me start with some quick introductions. To my immediate left is Lynne Parker. Lynne is the Deputy Chief Technology Officer of the United States. So in this pivotal role, she helps to guide policies and efforts related to the industries of the future with a strong focus on artificial intelligence. To Lynne's left is Olivia Zetter. Olivia serves as a Director of Research and Analysis at the National Security Commission on AI, and she brings to bear a very rich and impactful career in national security. And of course, Bob Work doesn't need any further introduction, so welcome back to the stage, Bob.
- Martijn Rasser: So, Lynne, I'd like to start with you. The American AI Initiative was launched almost a year ago now. What do you consider to be its greatest success so far?
- Lynne Parker: So, yes, as this audience well knows, the President did sign the American AI Initiative through an executive order. It was February 11th of last year. And I think what we're most proud of

is the broad, multi-pronged approach that this strategy has. And, at first, I should say congratulations to CNAS on this report. There's so much synergy there across what the administration is doing and what you're recommending, so thank you all so much for that report.

Lynne Parker: The multi-pronged approach that the American AI Initiative has set forth recognizes that there's not just one thing that we need to do, as CNAS' report also recognizes that. There is so many actions across so many sectors and areas, and so we focused in the American AI Initiative in particular on research and development on regulatory barriers, on education and workforce, on international matters, and those are the main prongs of activity that we've been involved with over the last year. And so we've had a number of important deliverables in all of those areas.

Lynne Parker: And in the R&D space, for instance, we want to know—certainly, the CNAS report is saying we should spend \$25 billion a year—but the question is, we have to start with a baseline. We have to know, what is it we're spending now, and a year ago we didn't know, frankly. There's never been a rollup agency by agency of what we're spending at AI R&D. And so, a few months ago, through the supplement to the President's 2020 budget, we released a rollup, agency by agency, of what the current investments are in AI R&D that now provides us a baseline of—I should say non-defense investments—that now provides us a baseline so that going forward we can say, "Okay, this is where we are now." Now we can say, "Okay, what needs to change? Where does it need to change?" And then we can begin to look at other figures for increasing those investments.

Lynne Parker: We certainly agree that more investments are beneficial and can be impactful, but how many more? How much more? \$200 billion as some have suggested might be a little of an overreach, but it's more than \$2 billion. So, that's something that we're very proud of so that we can now help the conversation going forward. You have to know how to spend those funds, so we released the 2019 update to the National AI R&D Strategic Plan so that we can see which areas of investment are most important. We also, across the interagency, released a progress report on what the agencies are doing in each of those strategic R&D areas so that we can now have a good collection of understanding of existing programs that are happening across the space in AI R&D.

Lynne Parker: If you look at regulatory areas, one of those issues is not regulatory so much as it is having some consistency which has to do with technical standards. And so, part of the executive order called for NIST to develop a plan for federal engagement in the development of AI technical standards, and that now helps us going forward, collectively, to look at these important technical standards that we need in order to be able to do things like measure the performance of AI systems, or create systems that have the level of trustworthiness that we actually want.

Lynne Parker: Just this week, as was mentioned, we are very happy to release the proposed guidelines for the regulation of AI technologies, and so that is now open for public comment for 60 days. That is the first in the world of regulations—a proposed regulatory process—that actually has teeth. It has teeth because all of the regulatory agencies that want to propose any kind of regulation about AI, that process works its way through the Office of Management and Budget, specifically the Office of Information and Regulatory Affairs, which is part of The White House.

- Lynne Parker: And so now those regulations, as they are proposed and they work their way through the system before they become policy, then have to abide by these proposed regulations that have now come out. That's critically important for many reasons such as ensuring consistency across our regulatory approach. It provides predictability to our industries and understanding how to think about these use cases of AI. It also is a messaging to our international partners about what we expect for how AI is used in our nation in contrast to how it's being used in other places. But it's also light touch in the sense that we don't want to squash AI innovation through regulation. So, that's an area we're also very proud of.
- Lynne Parker: And one more point I'll make is, at the international level, certainly these regulatory principles are based on a common foundation that we agree t, internationally, so we're very proud to have joined with the OECD nations back last spring, to sign onto the AI principles. We came to a consensus on what those high-level principles are, not only for what we expect our AI technologies to look like but also how we, as nations, should work together in order to create an ecosystem for us to be able to benefit from AI.
- Lynne Parker: If you look at that second half of the OECD AI principles, you'll see some things that are very synergistic with the American AI Initiative, synergistic with many of the items that are in the CNAS recommendations as well. So all of those areas, again, the multi-pronged approach to our AI strategy is critically important for us to make progress across the nation.
- Martijn Rasser: Okay. Well, since you mentioned the AI regulatory principles, I did have a question for you, but I think first I'll ask Olivia and Bob a question, then return to that point in a bit.
- Martijn Rasser: So, Olivia, the NSCAI released its Interim Report last November. The report identifies five lines of effort, seven principles, consensus principles. Let's start with the lines of effort. What are they, and how do they fit into an overall national AI strategy?
- Olivia Zetter: Sure, thank you, Martijn, and thank you Paul for having me here today. It's been incredibly encouraging to see all of the alignment from the American AI Initiative, from the CNAS report, and from the Commission's initial assessment on the areas that need investment and focus going forward. As Secretary Work mentioned, our interim report sought to set the stage from our first initial eight months of assessment and really drive into the areas that need careful consideration going forward and recommendations to be developed.
- Olivia Zetter: From that we divided into five primary lines of effort. Those lines of effort are investing in AI research and development, applying AI for national security, training and recruiting the next generation of AI talent, protecting US technology advantages, and marshaling international cooperation. These five lines of effort form the core areas that we see needing investment and attention going forward in order to prepare the United States to maintain its competitive advantage in AI for the foreseeable future.
- Olivia Zetter: From those, we came up with our seven principles as well as 27 preliminary judgments that are outlined in the report. In the interest of time, I won't go over all 27 right now, but I will run through our seven guiding principles that the Commissioners came to consensus on. The first is that global leadership in AI technology is a national security priority, and innovation is really the foundation of our national power. We've seen decreasing US investment in our technology base and R&D since World War II, and increase in investment from our strategic competitors. The US really has an imperative to guide R&D and

investments and our technological future in a way that protects Americans and ensures our national security.

Olivia Zetter: The second is that leading in AI application is a national security imperative. As Secretary Work mentioned, AI will have immense impacts on humanity as well as national security. The way that the US applies AI will have an impact both on how we protect our homeland as well as how we set the course for AI leadership around the globe. The third is that private sector and the government must build a shared sense of responsibility for the welfare and security of the American people. Because AI will have such an immense impact on our humanity, and because it is such a complex issue area affecting so many things, there needs to be a whole-of-nation approach to our national security when it comes to AI.

Olivia Zetter: The fourth is that people matter more than ever in AI competition. As Ainikki mentioned and Megan, there is an immense talent shortage and there's a lot of fantastic talent around the globe, so going forward we'll need to invest in our STEM education here at home making sure that we can recruit and retain the best talent from around the globe as well as making sure that we can get some of that talent from both places into government to help us on our national security initiatives. The fifth is that actions taken to protect America's AI leadership from foreign threats must preserve our principles of free inquiry, free enterprise, and the free flow of ideas.

Olivia Zetter: The sixth is that there is a convergence on ethics, and this ties in to the fifth and then the seventh that I'll mention, and that's that everybody wants safe, reliable, robust, trustworthy AI. And the pathway to getting there will require a lot of coordination from experts across the ethics community, across the private sector, and in government. So the Commission's approach to ethics to date has been to engage with as many stakeholders from across those communities as possible to really make sure that we get the ethics question right.

Olivia Zetter: And then our final principle is that the use of AI by the United States must have American values at its core, and this includes the rule of law. So, those three final principles really encapsulate how we see our AI future. It's one that has American values at its center and propels those values into the way that we see our international allies and partners around the globe adopting and applying AI for our future. So I'll leave it there and look forward to the conversation.

Martijn Rasser: Thank you. Bob, I'd like to ask you about China. So, the NSCAI interim report describes China as representing the most complex strategic challenge confronting the United States, and the Sino-American AI entanglement is particularly intricate. Given the Commission's work to date, where do you see the balance between cooperation and decoupling falling? And, as a corollary, what is the greatest obstacle to achieving this balance?

Robert Work: Well, after the end of the Cold War, I think both parties decided on a long-term strategy of entanglement with China, helping its rise, getting it into the World Trade Organization, and the expectation was over time that this would moderate the more aggressive potential tendencies of the Chinese Communist Party and would lead China to become more of a responsible stakeholder. I think both parties now, there's widespread agreement that many of the assumptions of that grand strategy no longer hold, and so the big debate now is between entanglement and disentanglement. What is the best long-term competitive strategy

approach, since we do believe that China is going to be the most daunting strategic competitor in the 21st century.

Robert Work: Now it's clear that China, all of the objective data, it is clear that China used the openness of the US system and took advantage of our own grand strategy. I think it was Ainikki who said \$600 billion a year in IP theft. The Director of National Intelligence said in 2018 alone, one in five US companies reported IP theft originating from China. This is a big problem. Some researchers from China attend US universities, then they establish a parallel lab in China where they commercialize technologies developed through US research. Some researchers from China have used peer review for grant applications and they literally copy parts or entire US research projects in China, and it sends many of its PLA scientists, its People's Liberation Army scientists abroad to be trained in US universities, I think 500 was the number.

Robert Work: So it is clear that the entanglement went too far, that it is working against US interests in a long-term strategic competition. So the debate in the Commission, focused narrowly on AI research, really, is what is the proper balance between entanglement and disentanglement in AI research? And I have to tell you that the Commission has not come to a conclusion yet. We're still in a heated debate. There are some that believe that entanglement is really the only logical way forward because most of the AI research is open source anyway and will proliferate around the globe.

Robert Work: But there are others who believe that we are too entangled and we should take steps to do so. And the difficult part is figuring out how to thread that needle. There are a lot of benefits to both countries with our AI collaboration right now. We went out and we talked with business leaders and to researchers, and they have cautioned the Commission, and I think probably the White House and CNAS, "Look, the human hardware supply chain investment in corporate connections between the United States and China just can't be unwound easily without generating enormous economic cost and unintended consequences for our economy and our research environment."

Robert Work: So, the key thing, I think, Megan and Ainikki pretty much laid it out. You have the good side and the bad side, and what is the balance? The Commission hasn't come to a conclusion yet. We will, by the time we submit our final report in March 2021, but I believe the recommendations in the CNAS report are a very, very good step. It's a very important thing that we need to debate, and we don't feel we should throw the baby out with the bathwater, but we do believe that we have to take steps to protect our IP, et cetera.

Martijn Rasser: So, Lynne, I wanted to go back to the regulatory principles that the White House announced this week. You mentioned in particular there was an international component to the messaging involved with that. So I wanted to ask you, our allies and partners in the G6 are, I think, pretty much on board with the Global Partnership on AI as well as other EU countries, India, New Zealand, I think as well. How do you see that gulf in opinion on how to approach regulations on AI? And I could foresee some negative consequences coming from that, that there's a patchwork of regulations internationally as a result, perhaps making it more difficult for us to effectively counter authoritarian use of emerging technologies such as AI. Do you share that concern and do you see that gap as being widening, or do you think that we'll be able to find some type of middle ground with our partners and allies on issues like this?

- Lynne Parker: Well, I do think it's a great ideal to think that we have a single, agreed, shared approach to how we want to oversee the use of AI but, practically speaking, all of the nations and regions of the world do have different cultural expectations and a different mindset in terms of things like, what level of privacy is enough, what is too much? And I don't think it's realistic to think through the GPAI, the Global Partnership on AI, or any other international mechanism that we'll all come together with a single set of unified principles and approaches from a regulatory perspective that are going to work everywhere.
- Lynne Parker: That's why the work that's happening at the OECD with those AI principles, at a very high level, is what we can agree on at the highest level. The question then is, how do you implement those and have impact and all the nations that are using them? And so now, internationally, the conversation is very much about how do we take those steps to implement these principles that we do agree on, and it's not just the OECD venue. Also in the G20 venue, the same principles with the caveat that I won't get into, were also agreed to. And so, that's something that we now are internationally talking about, at least the like-minded nations. How do we go about implementing those principles that we agreed to?
- Lynne Parker: Now, you can imagine a lot of different mechanisms. You can imagine a partnership of one type or a partnership of another type, and so I don't know how commonly known the Global Partnership on AI is. That is a partnership that was proposed by Canada and France over a year ago in order to look at a number of issues in terms of policy development for AI. Our view, in the administration, is that the OECD's work is already ahead of what they want to propose there. We are certainly very much as supportive of the goals of the GPAI but we want to accomplish those as effectively and as efficiently as possible. So the OECD is setting up the AI Policy Observatory which is bringing together ideas from around the world into a common repository where we can begin learning from each other. They have established a network of AI experts that are going to be digging into that and looking at what some commonalities are, and some level setting in terms of what exists and what some of the challenges are, and so forth.
- Lynne Parker: So there's a lot of activity going on there. Our view in the administration is that given that they have the resources, they have the infrastructure in place, they've already demonstrated the ability to bring people together from around the world to move forward on AI principles and now implementing those principles, that that's a good venue to work in, and so that's what we are supporting at the administration level.
- Lynne Parker: Now when it comes to regulations in general, obviously if you look at... this administration is not big on adding new regulations, in fact it's the other direction. But this administration also has, in the case of AI, has recognized that if you don't have some commonalities in this nation, then you do end up in many cases with a patchwork of regulations at the state and local level. And the challenge of patchworks of regulations at the state and local level is that now companies have to do something different in every area and that is actually hampering innovation. So this is a good area in which adding some regulatory principles for the use of AI is positive. It's pro-innovation rather than working against innovation.
- Lynne Parker: So that's where these principles are coming from, it's to provide that consistency that I mentioned, so that we can all make sure that we're making progress going forward. What the process sets forth is, once the final version of these regulatory principles, the guidance for AI regulation, once that's finalized after the comment period, then the agencies will have 180

days to come up with their own plans for how to be consistent in their particular use cases. So, this approach is recognizing that not all AI is the same. There are lots of different use cases. You don't need to have the same regulatory approach for each use case.

Lynne Parker: For instance, driverless cars. The regulatory challenges there are very different from medical devices. And so, at the level of the memo that we've produced, and that was the real challenge, is how do you come up with regulatory approaches that will work for medical devices, for drones, for mortgage loans. And so, at that level we have a level-setting, and now the 180 days will give agencies the opportunity to dig into the use cases for their regulatory domain that are most critical. Based on a risk benefit analysis, a cost benefit analysis, not all AI needs to have a regulation about it, but some use cases do, and so we're setting out the pathway forward for that process to take place—and that's why we're so proud of really, and really excited about with this approach—is because we can have some consistency across the board, some predictability, but we have a pathway forward so that as more areas are identified that need additional attention, we are able to do that in a consistent and reasoned way.

Lynne Parker: International collaboration is critically important for understanding particular cases that we agree on, but I think overall each nation has to define their regulatory approaches that makes sense for their nation.

Martijn Rasser: Great. Thank you very much for elaborating on that because I think this is going to be a very important issue that we're going to be dealing with for the next few years, because obviously you want to strike that balance between safety and security, but of course not stifle innovation at the same time, so I appreciate the detail on that.

Martijn Rasser: I just want to let you and the audience know, you'll have the chance to ask questions yourself in a few minutes. I've got a few more questions I'd like to ask first, but please start thinking about what you would like to ask the panel, things that are on your mind that you'd like some more insight on.

Martijn Rasser: So, Bob, I'd like to go back to you for another question. So, the Commission's mandate is very wide ranging. But specifically to AI in the DoD and the intelligence community, what are you particularly excited about what you're seeing? And conversely, what's keeping you up at night?

Robert Work: Well, it's a bit frustrating in that, in November 2014, the Department said that it was going to do a broad-base move to AI-enabled autonomy and autonomous operations throughout the Department. And then there was a change in leadership and if you read the National Defense Strategy, it was very, very consistent with the thinking behind what the Obama administration was thinking, this is what we have to do to prepare ourselves for a long-term strategic competition with two great powers.

Robert Work: Now, in one sense, I look back and I say, "Okay, it's been over five years and the amount of movement inside the Department has been modest." That's the most charitable way I can say it. But then on the other side, things really seemed to be starting to move. The Joint AI Center has been established, has been staffed. The most recent NDAA provides I think \$185 million in 2020 and similar amounts across the FYDP, each of the services have started to build what they called component mission initiatives where say, the Navy... how can AI and

machine learning, or how can machine learning help us in acoustic intelligence, for example. Let's have a project on that.

- Robert Work: There are all sorts of things that are happening at the lower level. Now, the National Defense Strategy says, "We want to pursue urgent change at a significant scale." So this is the way I think I would answer it. The intent seems to be there, but we certainly don't see urgency. I don't see the urgency that I would otherwise expect, and we certainly aren't at any scale that would make any appreciable difference in any future clash of arms. So I'm optimistic, I'm a glass half-full type of guy, and I think things are starting to move, but I'm frustrated by just how slow it has been.
- Martijn Rasser: On a related note, and I think this question would be good for both Olivia and for you, Bob, so, what are the obstacles that need to be addressed for the NSCAI's final recommendations to be successfully implemented?
- Olivia Zetter: Sure. So one of our overarching themes that we found is that our national AI strengths that we have in our innovation base, in our R&D leadership, in the incredible bright spots of programs that are happening throughout the interagency, and the strategies that we've set in place, whether it's the American AI Initiative, the DoD AI strategy, the IC's AIM Initiative, are not yet translating into large national security benefits, and that's because of a number of the impediments that you've mentioned.
- Olivia Zetter: Looking at those, in order to move the bottom-up successes that we've had so far into strategic gains, we'll need to have top-down leadership. And so we've seen immense top-down leadership from the White House, and from certain parts of leadership within departments and agencies, but overall, there needs to be recognition by those at the top and those, whether they are across the services or across the IC, that AI needs to be a priority and we need to make the changes that allow us to get AI done.
- Olivia Zetter: Some of those changes fall under the basic requirements to be able to do AI right anywhere but are particularly difficult in government. That leads to acquisition and infrastructure. Our acquisition system was built for large scale linear defense programs—building an aircraft—not for software or even AI, which is an extension of software. Very similar but not exactly the same. So we need to be able to transition those acquisition processes and authorities to allow for more rapid iterative development.
- Olivia Zetter: We also need to be able to make sure that we have the infrastructure in place. If we have the best algorithms and the best talent, if we don't have the cloud compute, the networks, the broad communications capacity within our systems, we won't be able to use them. So those are just a few of the big picture areas that we're exploring, and our recommendations will go into these in much greater detail going forward.
- Martijn Rasser: Right. Thank you.
- Robert Work: Yeah, the Department of Defense I think has approached artificial intelligence and autonomy from a "let a thousand flowers bloom" strategy, where it's more of a bottom-up approach. For the transformation that was envisioned back in 2014, it really requires a combination of a strong top-down approach, and, quite frankly, the Department hasn't had a true champion on why this is important for the last several years. Now, I actually place the

blame, you know I say, "Hey, I'm a little upset that we haven't moved," but there's two things that, in hindsight, I wish I had not, since I was a Deputy Secretary, done.

Robert Work: The first thing that people forget is that, what Secretary Hagel announced in November 2014 was the Defense Innovation Initiative, which was a broadscale, "We've got to relook at the way we have been operating for the last 20 years and we need to change to get ready for a long-term strategic competition." Line of effort three was a competitive strategy to retain military-technical superiority. That was line of effort three. From that line of effort came the Third Offset Strategy. The Third Offset Strategy became shorthand for line of effort three, and it ultimately became shorthand for the whole DII, which was very, very bad because there was a line of effort on strategy which I would say culminated with the 2018 National Defense Strategy, line of effort on operational concepts, line of effort on war gaming and experimentation, a line of effort on increasing stakeholders and Congress in the White House, a line of effort on DoD-IC integration, and a line of effort on information capabilities management—reveal capabilities for deterrents and conceal capabilities for war-fighting deterrents.

Robert Work: That kind of got washed out and everybody had the little bright shiny object was the Third Offset, so many people thought that it was all just about technology, and it wasn't. It was about transforming the Department in a broad way. The second thing is, the DSB study which had such an enormous impact on our thinking, wasn't the DSB study on artificial intelligence. It was the DSB study on autonomy, and in my view, we spent way too much time talking about artificial intelligence and not enough time on how you get after autonomous systems and operations, which would have the greatest impact on the way the Department would operate.

Robert Work: And so I see that changing now, but I think that materially contributed to the fact that we aren't as far along as we are, and I'm very, very happy now that JAIC is in place, Lieutenant General Shanahan seems to be now empowered to speak more forcefully on these issues, and I'm hoping that the Deputy Secretary and the Secretary do also. A very important change as the Secretary of Defense has said, artificial intelligence is his number one research priority. And that should start to flow down and start to affect what is happening in the Department.

Martijn Rasser: That was great. So, Lynne, in that vein, as you survey the AI landscape, what do you see as being America's greatest relative weakness in the global AI competition, and how is the administration preparing to address that?

Lynne Parker: I think if you look at today, there's a lot of concern about our lack of access to data compared to China. And so I think today that would be what I would call, if I had to pick one thing, our greatest weakness. Certainly China is collecting, buying, stealing data at enormous levels and that does give them an advantage for the kinds of AI technologies that are becoming impactful today. That's a very narrow type of AI technology to be quite frank. It's the technologies that are obviously data-driven, they're deep neural nets of some sort, they are doing pattern recognition. There is enormous potential for AI to do so much more that is not purely data-driven artificial intelligence based on deep neural nets and pattern recognition.

- Lynne Parker: So, if you look at it from the US government's perspective, we'll never compete with China on the amount of data. We should own that and be proud of the fact that we do have privacy laws and we're not collecting all of the data on all of our citizens and using it in ways that are not consistent with our values. So that's not something we'll ever apologize for. At the same time, if you look at the state of the field, and as a long-term AI researcher, you look at what people do as a proof of principle of what intelligent systems can do. And if you look at small children—and I tell this story frequently—when you talk about a small child that... their parents are showing them a little picture book and it has a fish in the picture book, and they read the picture book twice, and then they take a field trip and go to the aquarium. The child is able to immediately look at the real fish in the aquarium and say, "Fish," with only looking at a very small number of examples of just drawings of fish.
- Lynne Parker: Clearly, as human beings, we have the ability to be quite intelligent and have lots of reasoning capabilities that have nothing to do with zillions of data examples. So, that's why, at the US level, we're investing significantly in other types of AI at the R&D level that will get beyond just the pure data-driven approaches. DARPA's AI Next campaign, for instance, is investing significantly in this area. NSF is also investing in this area. Broadly speaking, across the federal agencies that are investing in long-term R&D in AI are looking at these other types of AI systems. We will have breakthroughs in those areas. It's not easily predictable, but at some point we will have new AI techniques that don't require all this data, and then suddenly almost overnight the advantage can go away.
- Lynne Parker: That's one thing. It's not going to happen probably tomorrow, and so, in the meantime, what the federal government is doing—and this is one of the actions that's in the President's executive order on AI—is to look at ways that we can make federal data more available for AI R&D and testing. The challenge there, of course, is again the privacy concerns and security concerns. We don't want to just let all federal data become available, so we're looking at ways of making that data increasingly available but perhaps through secure data stores and other kinds of approaches where people who have a need to access the data through federally funded research can access it.
- Lynne Parker: We're also looking at the ripeness of particular technology such as federated machine learning that can allow us to pull data together but—you're able to learn from it, from a machine learning perspective—but you don't actually have to see the data, and so therefore you can protect intellectual property and privacy and security, but still be able to create tools. And so, these are a couple of the approaches that the administration across the board is working on so that we can both invest in those R&D areas that will overcome the need for vast amount of data, and also looking at ways that we can do a better job of making federal data more available for AI R&D and testing.
- Martijn Rasser: I agree with you. The potential for a breakthrough in your being able to train models with small data sets is very exciting. There is some great work going on at the MIT-IBM lab. There's several start-ups that are working on hybrid solutions taking aspects of symbolic AI in deep learning. So, yeah, it's very exciting times, and I think if the United States is able to achieve a breakthrough there, that could be a game changer.
- Martijn Rasser: Before I turn to the audience for questions, I just want to ask you one last thing. So what is the one thing of the administration's work on AI that you would want us to walk away with today? Just like a little soundbite.

- Lynne Parker: Well, I think, to go back to what I started on at the very beginning, I think certainly this is not just a one-action kind of activity. The administration is very proud of the various activities in the multi-pronged approach that we're taking across the... what we call our national AI strategy, which is the American AI Initiative. So I think it's important to recognize that we do, as a nation, have a strategy and the key challenge is for all of us to continue to work in those areas. I mentioned R&D, education in workforce has come up several times, the importance of looking at the regulatory issues and technical standards, looking at the international climate, looking at the national security issues. So I think we want to get beyond the question of, "Do we or do we not have a national strategy?" We clearly do have a national strategy on AI, and now we need to look deeper at how to implement these strategies looking at a lot of continued good ideas. Our work is never done, and so continuing to implement in all of these areas and collecting these good ideas going forward.
- Martijn Rasser: Right. Thank you. Well, with that, I'd love to open it up to the floor. So, if you could please state your name and affiliation, and the gentleman in the back has already raised his-
- Richard Jordan: Do we need a microphone?
- Martijn Rasser: Yes, there's... right behind you, sir.
- Richard Jordan: Right. Thank you, I'm Richard Jordan, Senior Dean of NGO Representatives, non-governmental at the UN for 40 years. Lynne, since February 11 is the International Day for Women and Girls in Science, the executive order to which you have referred, did that include any specific support for women in science and pay parity, equal pay for equal work by men and women doing the same job? Thank you.
- Lynne Parker: The executive order did not mention that explicitly. Certainly that is something that I care a lot about, but what we tried to do in the executive order is make it specific to needs of AI. That, I would argue, is a need that's much broader than AI.
- Michael Hopmeier: Michael Hopmeier, I am a consultant to the Army Futures Command. One of the questions we're grappling with now is, and I appreciate very much your comments on the technology development, but technology is only one part of the problem. What you've been talking about now, I think, has the a priori assumption of, if you build it, they will come. It's a great technology, new capability, everybody will adopt it. What we saw with the air force, for example, with drone technology, there were immense cultural and social issues in being able to accept it. Questions of, can a drone pilot only be a drone pilot if he's a rated combat pilot? Do they even get to wear flight suit? What type of medals? How are you looking at, or what discussion is going on now about the other half of the problem? The cultural and social acceptance of these new technologies on the battlefield? Thank you.
- Lynne Parker: Well, I won't speak specifically to the battlefield, but I could speak to that question from the government use of AI perspective. So, this is a key point. Back, a few months ago, we at the White House hosted the AI in Government Summit, and that's not something I've talked about so much here today, but it's about the government's own use of AI and how do we move forward on that. We brought in over 175 people, multi-stakeholders from government, academia, industry, non-profits, to consider how do we advance the use of AI

in government, recognizing that there are many benefits to improved delivery of mission for the American people and improved efficiency, and so forth.

Lynne Parker: And so, this very issue came up as one of the main talking points. You can go to ai.gov and see our summary that includes some of these points about the challenges of changing the culture, as you say. The whole process in many cases has to change in order to take advantage of these AI technologies. So, rather than being afraid of that, certainly there is a process of embracing it and facilitating it. So what we have done at the administration level, is to work with General Services Administration and they have now spun out, or are spinning up, an AI Center of Excellence, and also a Community of Practice across the federal government.

Lynne Parker: That Center of Excellence and Community of Practice will allow all of the people in the agencies to share ideas and to share approaches that they're looking at in terms of AI pilots, for instance, in certain agencies, that have common applicability to other agencies, and so then they can learn from each other and help foster and facilitate this process. We're very happy that the JAIC in fact, is one of the very first clients of this community of Center of Excellence, and I believe that there'll be learning lessons going on in both directions. I think the Center of Excellence can learn from the JAIC and the JAIC can learn from the multi-agency Center of Excellence about how to really go to the implementation stage of these techniques.

Lynne Parker: So that's something that we're doing at the federal level across all agencies to try to provide abilities for agencies to learn from each other, to help figure out how to address these cultural issues that are so key.

Robert Work: In my view, from the OSD perspective I'm speaking from, the cultural issues are rightly in the military departments and the services. That's where it should reside. The way this, I think, should work virtuously is the Department of Defense tells the US Army, by the end of the FYDP you will have a human machine infantry brigade combat team, and it will be operating. The army then goes about saying, "Okay, what is the concept for operation? Am I going to have mixed units with humans and machines? Or if I'm going to have human platoons and just have a machine weapons platoon?" And then, experiment and say, "Hey, we didn't realize that a commander would have to think differently in this way," and these are the cultural issues we have.

Robert Work: So, if the army's waiting for OSD to outline all the cultural issues, don't wait. The Marines, the Navy, the Air Force and the Army all need to be pressing after this fast. I look to the Air Force and I say, "Gosh, you have the senior acquisition executive in the air force saying, 'I want to make the United States Air Force a software organization'." I mean, the implications of that sounds really easy, but it's pretty profound when you think of all of the different cultural impacts that it might have. And so, what does the air force start doing? They start to do these Kessel Runs where they get really smart coders that are in the Air Force and they put them with operators, and they say, "Here's a problem, let's try to solve it." And they do it really rapidly and innovatively.

Robert Work: Talking with Will Roper, all of the major Air Force leaders now want their own Kessel Run. I mean, you're starting to see this very broad cultural transformation. So, from an OSD perspective, I would say, "Look, don't wait for OSD to tell you what your cultural issues are.

We expect the departments and the military services to figure that out on themselves. We will help you once you identify problems that need to be addressed within the program."

Martijn Rasser: Olivia, do you have any-

Olivia Zetter: Sure. I just quickly... From a training perspective, this is where issues such as human machine teaming and trustworthy AI are so important. So, as we think about how we build the right talent set for... whether it's DoD or other agencies across the government, we need not only the top developers and AI PhDs working inside the government, but we also need general AI literacy across the organization so that there's the right level of understanding about how AI can help specific mission sets, and how much to trust an AI system, so that you're not putting either too much emphasis on it or too little emphasis on it in your battlefield needs.

Olivia Zetter: So this is something the Commission is thinking about as we think through our talent recommendations. It's not only how do we get top AI experts in, but how do we raise the broad literacy of everyone to be able to take advantage of what AI has to offer.

Robert Work: And this whole cultural thing, I'm going to use an Army example because I think it's one of the coolest ones. Recently I was briefed, the Army just did a completely robotic breach of a complex obstacle. You know, ditches, concertina wire, minefields. This is one of the most difficult tactical operations that a small unit has to do. Typically what happens it's an either a company or a battalion training package and you go through all of the different... the mission essential tasks and you practice and you practice and you practice, and then you kind of lead up to this big exercise and you do it. It takes a long time, but the Army recently did the entire thing robotically and General Murray, who's the head of the Futures Command, was saying, he went down to talk with the soldiers because usually it takes a long time to train the soldiers up for this very complex task, and he said, "Hey, how long did it take you guys to get ready for this?" And they said, "Oh, about two hours."

Robert Work: And the reason why is the Army had a very good cultural approach. All of the controllers were Xbox controllers. As soon as the soldiers got the Xbox controllers in their hands and were told, "Okay, these are the functions that you have," within two hours, they're... and they're just moving the stuff all over. So, those are the type of things, those type of innovative approaches. Once they start happening across the Department and across the government, you're going to see this transformation catch like a wildfire. And the more we learn, the better we will be, and the faster we'll go.

Martijn Rasser: Great. Any other questions from the audience? Yes, sir.

Fritz Barth: Thank you. Fritz Barth with the Office of Under Secretary of Defense for Intelligence. At this point, my understanding is that almost all AI is very limited and fragile, right, and I'm wondering if you are looking at... So, if you need to counteract from either adversary or criminal use, now is the time to start working on that before it gets better. Is there any look and examination at that issue?

Lynne Parker: Yes. There's an enormous amount of research that's going on now about, how do you address adversarial attacks on machine learning, and recognizing that there are many types of AI systems that are vulnerable to these like data poisoning attacks and other types of attacks.

So, it's not a challenge that immediately has an answer today. That's why the research and development is so important, but at the same time, there are a lot of efforts as well to see what we can do to understand the systems that we have now so part of that is another area of research which is explainable AI, are helping us to understand how an AI system is actually working so that then we can see what its vulnerabilities are based on a deeper understanding of how the AI system is actually learning.

Lynne Parker: So there's an enormous amount of work happening at that level. Certainly lots of discussions, in general, about... that is a limiting factor to more implementation of AI in practical application. So, there's a lot of ongoing discussion and work in that area.

Olivia Zetter: Sure, and I'll add that as any other connected technology, AI has the opportunity to both protect us against vulnerabilities as well as increase our potential threat surface. And so, this is an area, like Lynne mentioned, that is getting a lot of attention from both the IC, the service labs as well as the FFRDC community to start thinking about it now. An issue such as general cyber security as well as disinformation and how we think about deepfakes as well, are all interconnected in this.

Olivia Zetter: So, through our briefings we've seen some really creative and innovative research areas that people are looking at, and we plan to explore this more going forward. It's an area where we want to be able to invest in potentially automated defenses, but also recognize that our strategic competitors and potential non-state actors as the technology is democratized will be able to do the same, so we have to think really cautiously about it, and then I'll just drop a plug for the Defense Science Board is putting out their counter-autonomy study at some point in the near future, so they've also taken a distinct look at this issue.

Robert Work: The only thing I'd add here is that, most of the advances since 2012 have been in machine learning, and it frustrates me whenever I hear AI/machine learning. Machine learning is just what DARPA would call the second wave of AI, and the first wave was expert systems, if/then logic. And with the compute power we have now, I believe that we are spending far too much attention on the machine learning aspect and not enough on the expert systems, primarily because it uses if/then logic and you don't have to worry about explainability. It is built into the program. If a problem occurs, you can recreate the problem and know exactly what happened.

Robert Work: The autonomous ship that the navy had, it's rules of the roads. You know, you can press the button and the thing will autonomously navigate between Norfolk and Bahrain. All of the rules of the road are all first wave AI. It's just if/then. If the ship is to the port and if it's closing the... here are the sums you do. But it needed to have machine learning in the camera just to say, "Okay, is this a container ship bearing down on me or a sail boat? And how fast can it go?" So it was a righteous combination of machine learning and expert systems would allow that thing to go.

Robert Work: And right now, this is one of the things I think the Department of Defense can really move faster on if they said, "What are some of the expert systems that we could use today to get going forward?" So, then the third wave AI, is what Lynne talked about is what DARPA's looking at now. What about contextual AI, how do you make the machines reason more like a human? But, there is a lot still left to go in first wave. And this is one real critical thing. In the beginning we thought all of the AI applications would be narrow task applications that

were told by a human, "Do this task for me so I don't do it, but do it in furtherance of my objective." There wasn't anything about a fully autonomous independent thing wandering around the battlefield which would require artificial general intelligence. It was take small steps first, learn as you go, get better and experiment.

Martijn Rasser: Ah, yes. Gentleman right there in the gray suit. Yeah.

Frank Hoffman: Frank Hoffman, National Defense University. I have a non-security related question, even though I agree with everything Mr. Work has already said about the defense strategy and our lack of scale and change, but I rarely get to question somebody from the White House, I don't want to lose this opportunity. I read a book that I got for Christmas, The Technology Trap from Frey, divides the history of technological development into human enabling capabilities and human displacing capabilities. Makes a very clear argument that artificial intelligence and the revolution you're talking about is dramatically influential in human displacement and that there are socio-political, cultural, and economic dislocations that come with that, that are severe in social and institutional cultures and context. Is the White House aware that it's advancing a revolution that is going to have a socio-economic and employment impact in the next 10 or 15 years in the United States, and how does that impact our welfare and whole-of-nation approach, because you're, at the strategic level, you're involving I think the welfare and the whole wholeness of the nation to some degree.

Frank Hoffman: And are you conscious about that, because there's an economic office somewhere in your environment that I think has a contrary objective in terms of the manufacturing base, the agro business and some of the other industries that are going to be severely impacted by this technology.

Lynne Parker: Yes, of course. I think if you look at the workforce issues, broadly speaking, we've touched a little bit on the training side of it at the formal, maybe K through 12 or college level, but there's also an important aspect of rescaling. I think if you look at the history of technology and capabilities across the years, the decades, the centuries, there are always these transition times. Most of our grandparents, or great grandparents were farmers, and today we have relatively few farmers, but we have managed to create jobs for people to do other kinds of things. And so that's why the President established early on the National Council for the American Worker in order to make clear that we need to embrace and learn how to adapt in these times of changing jobs due to automation and technology.

Lynne Parker: And so that is a partnership with industry and it has challenged industry to come up with opportunities for people to learn new skills in this new era of AI and automation, recognizing that not everyone is going to be immediately prepared for these jobs. If you look right now in our nation, we have historically low rates of unemployment, but we still have people who are looking for jobs. And so, if you look at the people who are looking for jobs and they have certain sets of skills, and you look at the many open jobs that where industry is looking for people, the scales don't align.

Lynne Parker: And so, the beauty of this National Council for the American Worker is that it challenges industry to provide rescaling and retraining opportunities for people who want to learn new skills for the open jobs of today. And so the way that works is that, there's this pledge to America's workers where over, right now, I think over 350 companies have pledged over 14

million retraining opportunities for Americans. And so, that is an opportunity for people to learn these new skills.

Lynne Parker: We can't ever be afraid of technological change. AI is here. We can't make it go away, and so what we can do though, at the federal level, is to provide people the opportunity to learn and to adapt. Our position, we stated repeatedly, is that as a nation we need to embrace the idea of lifelong learning. No longer is it, you've learned something when you were 22 and now you have some right to do the same thing for the next 50 years. We need to embrace the ability, and Americans are so good at this. We're so good at adapting and taking on the next new challenge and overcoming it.

Lynne Parker: But we do need to provide these rescaling and retraining opportunities. Administration has been very strong in areas of like, there's a federal strategy for STEM education that talks about helping everyone learn more computational literacy. There are initiatives for apprenticeships that are being very helpful for giving people the opportunity to learn new jobs, and so forth. And there are many other kinds of workforce initiatives as well. So that's the approach of this administration. I think it's been very effective in getting the private sector involved as well so that we, as a nation, can move forward and embrace these new opportunities.

Lynne Parker: It's not to say that there won't be some job occupations that will be going away, but there are also new job occupations that are appearing now, and so there's enormous amount of opportunity, and it's enabling to a large extent people to engage in jobs that are, perhaps require more creativity, more interesting kinds of jobs, rather than kind of drudgery sorts of jobs. So there are lots of positives here, it's not to say that everyone is happy about it, but I think there's a lot of opportunity and we want to make sure that people have the opportunity to engage.

Martijn Rasser: That's a good, positive note to end on. We're at the top of the hour, so unfortunately, we're out of time, but thank you all very much for your great questions, your great attention, and please join me in thanking our panelists.