DEFENSE INNOVATION BOARD
PUBLIC LISTENING SESSION

- - -

CARNEGIE MELLON UNIVERSITY
THURSDAY, MARCH, 2019

- - -

The Defense Innovation Board Public Listening Session, taken before me, the undersigned, Eileen L. Drake-Ober, a Notary Public in and for the Commonwealth of Pennsylvania, at Carnegie Mellon University Studio Theater, 5000 Forbes Avenue, Cohon Center, First Floor, Pittsburgh, Pennsylvania 15213, commencing at 1:36 o'clock p.m., the day and date above set forth.

- - -

APPEARANCES:

Panel Members:

Dr. Richard Murray
Dr. Missy Cummings
Dr. Michael McQuade
Milo Medin

- - -

ALSO PRESENT:

Joshua Mancuse
Christopher Brunett
Provost James Garrett

- - -

P-R-O-C-E-E-D-I-N-G-S

MR. BRUNETT: Good afternoon. My name is Christopher Brunett. I serve as the designated federal officer for the Defense Innovation Board. It is my role to open this public listening session of the science and technology subcommittee of the Defense Innovation Board. Thank you to Carnegie Mellon University for hosting us today -- hosting today's sessions.

If you've not done so already, please silence your electronic devices. This session is part of the Defense Innovation Board initiative called artificial intelligence principles project. Today's session is being recorded and live streamed to allow members of the public to attend virtually. It will also be accessible on the board's website, innovation.defense.gov. Thank you to the defense media activities for providing their expert support to this event.

Welcome to all of our in-person and virtual attendees. As we begin this public meeting, allow me to share a few procedural remarks. This board is a discretionary

independent advisory board operated under the Federal Advisory Committee Act and the government Sunshine Act. Today's meeting was announced in a federal register notice posted Friday, February 15th, 2019. There have been no significant changes to the meeting's agenda as posted in the federal register notice. The public was invited to submit written comments for board members to consider. 16 written comments were received in advance of today's session. These comments will be posted on-line with the minutes of this meeting. We welcome additional written comments on a rolling basis, which can be submitted via our website. The primary purposes of this session is to provide an opportunity for members of the public to provide verbal comments to the board subcommittee today.

As a reminder, these are comments to the board, not question and answer session. Board members may ask clarifying questions.

With that, I now turn the meeting over the board's executive director, Joshua Marcuse, for his opening remarks and

introduction of our board members.

MR. MARCUSE: Thanks, Bruno, and welcome, everyone, and thank you to Carnegie Mellon University for hosting us. This is our first official public listens session for the AI principles project, and we cannot be more excited to start this off in the country's premier institutions for AI research development, thought leadership, and its application.

More than a half century ago, Carnegie Mellon's own Herb Simon and Alan Newell wrote the first artificial intelligence program. Since then CMU has forged a culture around using technology to solve real problems, which is why we believe CMU is the perfect host for this sort of engagement.

I would like to introduce our board members present today: Dr. Missy Cummings, who is the professor of engineering and director of human autonomy lab at Duke University; Dr. Michael McQuade, the vice president of research right here at CMU; Mr. Milo Medin, vice president of wireless services, Equitable; and Dr. Richard Murray,

professor of control and dynamical systems and bioengineering at Caltech. Dr. McQuade and Dr. Murray are the co-chairs of the Defense Innovation Board's science and technology subcommittee, and as such, they are leading this initiative.

I would like to give a brief overview of this project and sort of how we all came to be in this room together to kind of set the context for this conversation. Last July, the department asked the board to undertake an effort to help establish a set of artificial intelligence principles for defense. After a few months of planning and some internal discussions, the first weeks of 2019 saw the board begin convening a mix of academics, researchers, ethicists, lawyers, business executives, non-profit leaders, venture capitalists, policy experts, and a wide variety of other people in the AI field to provide input to this process through a series of roundtable conversations.

We are taking care to include not only experts, who often work with the department, but also AI skeptics, DOD critics, and

leading AI engineers who have never worked with the department before. There may be differences of opinion among this diverse group, since these matters are controversial. We will not shy away from disagreements, as respectful and forthright dialogue should lead to meaningful understanding on all sides, and a robust context of ideas should generate new insights into this question of how to set up an ethical framework for the Department of Defense.

DOD recognizes the need to view AI differently than other technologies, especially on ethics and the imperative to get this right on how DOD employs these technologies. Artificial intelligence not only affects the men and women who serve in uniform in our country, but societies around the world, and that's why the board is a different process than we typically do, to ensure that our process of developing these principles is robust, inclusive, and transparent. We want everyone to take part in this dialogue because these issues touch everyone, and that's why it's so important

that all of you are here today to join us, and to those who are watching on-line, or to come to watch this video in the future, we appreciate your taking this role as a citizen very seriously.

Today's public listening session is one element of this board's initiative, and we hope that this helps the board -- helps advance the board's dialogue. I want to outline the flow of how this remaining few hours will unfold. In a moment, I'll ask Michael McQuade to say a few words on behalf of the Defense Innovation Board, and then we'll hear from the Carnegie Mellon provost, Dr. Jim Garrett, someone we're very glad to involved in this initiative who we've had the pleasure of interacting with over the last couple of days.

After Jim's remarks, we will move to the public comments for the bulk of the time, when audience members can address the board, and I know this is the part that you are all here to see, so when we get to that point, I'll explain how the public comments will proceed in more detail.

If you haven't already submitted a comment on-line -- and, again, Bruno said we have received 16 of those -- when you RSVP, I'll ask you to take one of these comment cards and give them one of my colleagues who will be in the aisles. This is the way in which you essentially raise your hand, and so, we really hope that all of you -- if you haven't taken a comment card, will do so, because we would love to hear from as many of you as we have time for.

Just before the public comments begin, we'll collect those cards -- and here are the important part, we really are going to limit every comment to the five minutes. It's perfectly acceptable for you to be 20 seconds, but it is not acceptable to be 5 minutes and 20 seconds. We do this to be fair so we can hear from everyone, so I'm going to tap on the microphone to let you know when you have one minute left, and when you have no minutes left I'm going to be as polite as possible in making you sit down.

We also have written comments, and so, for those of you that submitted written

comments, but aren't physically here, I have asked Bruno to read those comments, because I think it's very important that those public comments be heard by everyone, as well as the board members, so luckily Bruno has agreed to read them so you don't have to listen to my voice that much longer.

And now, I will turn things over to Dr. McQuade for a few words.

DR. McQUADE: Thank you, Josh. Let me also welcome everybody here with a -- two hats on; one as a CMU employee and one as a member of the Defense Innovation Board.

The Defense Innovation Board was started under Secretary Carter with an objective to bring an outside view around topics of innovation, culture, technology to the department to find ways to help the department, if possible, improve the way it operates and achieves its mission for the country.

We have, over the last two years, had multiple engagements with the department. They have ranged from examining the way the department does software, does including both

acquisition and execution of software. We have also substantially had issues around workforce and workforce deployment, and we have had a number of issues, a number of examinations, around technical capacities of the department, so a fairly broad remit, and through all of those discussions, the overarching implications of AI as a broad-based technology has been front and center for quite some time, and for those of you who follow the department, you will see that there are a large number of initiatives over the last year or 18 months relative to the way AI is being examined and rolled out, including some activity here at CMU, substantial activity here at CMU.

In the process of that discussion around AI, both on the Defense Innovation Board with my colleagues and with the broader community, we have constantly had in the middle of that conversation that it is not just the technology, it is not just the fundamental science or mathematics or however you want to characterize AI, it is also a social element of the way -- the way the

department and society operate, and, therefore, carries with it both obligations and responsibilities around ethics and the ethical use in an environment where the department is trying to prosecute its mission.

Josh gave you a little bit of background about why we're here and how that got set up. I would simply end by saying the following: That in this day and age, and with the technology as broadly implicative as AI, it is absolutely necessary to encourage broad public dialogue on the subject. I would also say it is, in this day and age, even more important to encourage respectful public dialogue in -- in that comment. So everybody who is willing to speak today, we want to hear what you have to say. I would just ask us all to be respectful in the way we communicate that message, and please be as open and evocative for us as possible, because at the end of the day we are trying to represent the broad set of issues and consequences and discussion around ethics of AI.

So thank you, very much, for being here. We really do appreciate hearing from everybody. And thanks to all of my colleagues for coming to CMU.

Josh?

MR. MARCUSE: Many of you know Jim as a pillar of the CMU community, a three-time graduate of the university and a longtime professor of civil and environmental engineering, having also served in a variety of leadership roles at CMU before becoming provost, has a strong background in sensors, data analytics, and a unique blend of technical research and leadership experience. It's my great pleasure to welcome him to the podium.

MR. GARRETT: Thank you, Josh, and good afternoon to all of you. I'm Provost Jim Garrett. It's my pleasure to welcome you, those of you here in person, as well as those of you on-line and on the live stream, to this afternoon's public listening session with the Defense -- Defense Innovation Board.

I hope that prior to today's session you had the opportunity to learn about the

amazing work produced by Carnegie Mellon University for artificial intelligence strategies for a host of important industries. In fact, our university has held a place at the forefront of innovation and emerging technologies for decade, globally and locally right here in Pittsburgh as an essential component of our city's own tech revolution. We're proud that this recognition is one of CMU's signature strengths, alongside our deep tradition of infusing collaboration and interdisciplinary academics that expand the horizons of knowledge. Every day our students, faculty, and researchers work collaboratively across disciplines, and share their curiosity, knowledge, and discoveries with each other to innovate and to inspire each other.

Even more broadly, I believe that as a university, our role is to focus not just internally on the impactful technologies and scientific research produced by the Carnegie Mellon University community, it's equally vital that we bring outside perspectives and voices to our campus. By providing platforms

such as today's public listening session, we live up to the respect and world-class reputation that we've earned as a university. We help to lead conversations that explore new ways of thinking, and ponder questions that make us evolve in the work that we do.

Today's listening session surrounding the ethical and responsible use of artificial intelligence is a prime example of how our university also values freedom of expression, especially as it relates to the accountable ways that we're using the dynamic technology -- this dynamic technology. The sense of accountability, as well as the societal implications and artificial intelligence's impact on the workforce, are also topics that we teach our students in a number of exciting and relevant academic programs in this space.

We're proud that you chose Carnegie Mellon University to facilitate these important discussions. Our university, as was pointed out by Joshua, is the birthplace for artificial intelligence, so I can't think of a more fitting place to examine the

guiding principles for the technology that is revolutionizing the world, and what solutions we're trying to explore.

Together, all of us have a tremendous responsibility to do the right thing. I hope that the DIB members and all of our distinguished guests visiting today will benefit greatly from different perspectives that will be shared, and are inspired through Carnegie Mellon's leadership towards discovering common objectives as global technology leaders.

Now, please join me in welcoming back to the podium Joshua Marcuse, who will kick off today's public listening session. Thank you.

MR. MARCUSE: Excellent. Thank you, very much, Jim, and thank you again for hosting us here. We're very happy to be here, and had great, great sessions earlier today with faculty.

So I thought a good way to get this conversation started was to offer a very brief primmer of, I think, three key ideas that undergird this conversation about

ethics, so let me just run through a few things that I think will help inform the audience and the public about what the department's current approach to this issue is.

In a memo to all DOD personnel last month, on leading with an ethics mind set, the Acting Secretary of Defense Shanahan observed a key component of leadership is reinforcing ethical behavior across the full spectrum of our work and recognizing ethics principles as the foundation upon which we make sound informed decisions. Ethics are, indeed, foundational to the responsible use of artificial intelligence, and we believe that the department can show technical leadership and moral leadership in the world on this issue.

Secretary Shanahan's insight has been -- the approach that DOD has also sought in all of its other adaptations to emerging technology can be applied to artificial intelligence, and that's why DOD leaders asked the Defense Innovation Board to develop an ethical framework for the adoption of AI

on the very same day it announced the establishment of the joint AI center. That request was the basis for the board launching the AI principles project about which this meeting is a crucial part. As a foundation for this discussion, I want to cover three ideas that I will only very briefly introduce.

First, I'm going to discuss how the department believes the law of war applies to DOD's use of AI, particularly in military operations, so I would like to point out military operations are probably a small fraction of the situations which the department might choose to use artificial intelligence; second, what the department's current policy is on autonomy and weapons in studies because of the important intersection between autonomy and artificial intelligence; and, third, what the department's newly released artificial intelligence strategy says about AI in ethics.

Those of you who don't know, the department really just launched only about a month ago its -- its first AI strategy.

These are exciting new developments in the field of AI, like many emerging technologies that preceded it, but the department's commitment to our values and the rule of law is enduring. To put it another way, the addition of AI to any existing system process or product doesn't diminish the department's commitment to abide by the law of war.

So what really is the law of war? The law of war is a body of international law, specifically adapted to the conduct of warfare. For the United States, this body of law includes treaties the United States has accepted, such as the 1949 Geneva Conventions and customary international law which results from the general and consistent practice of states done out of the sense of legal obligation.

There are five fundamental principles that the foundation for a law of war, which is explained in DOD's official law of war manual, and I'm just going to very briefly run through what these five fundamentals are.

First, the military necessity

justifies the use of all measures needed to defeat the enemy as quickly and efficiently as possible that are not prohibited by the law of war. Humanity forbids the infliction of suffering, injury, or destruction unnecessary to accomplish a legitimate military purpose. Proportionality means that even where one is justified in acting, one must not act in a way that is unreasonable or excessive. Distinction obliges parties to a conflict to distinguish principally between the armed forces and the civilian population, and between unprotected and protected objects. Honor demands a certain amount of fairness in offense and defense, and a certain mutual respect between opposing military forces.

The key points to understand about the intersection of the law of war and AI I would say are the following: First, international law of war provides a well-established body of law to address the legality of conduct in the context of armed conflict.

Second, existing law of war rules can apply when new technologies such as new types

of artificial intelligence are used in armed conflict.

Third, the fundamental principles of the law of war provide a general guide for the conduct of war where no more specific rule applies; thus, provides a framework to consider novel, legal, and ethical issues posed by emerging technologies such as AI.

DOD has a robust process to implement the law of way, including training regulations and procedures, the reporting of incidents involving alleged violations, investigations and reviews of incidents, and appropriate corrective actions.

DOD lawyers are engaged in efforts to articulate how existing law of war principles apply to emerging technologies such as artificial intelligence.

And, last, and conversely, experts are engaged in understanding how or if the potential of something like an emerging technology such as artificial intelligence precipitates strategic or tactical questions about which current law of war is not yet developed, and that is the reason that we're

here, because we are asking that question, and, hence, our listening session is designed to aid the department in undertaking that process.

So with that, we move to the second area, which is what is the current policy on autonomy and weapons systems. The department has issued specific guidance on autonomy in weapons systems, which is not synonymous with artificial intelligence, but is clearly related to it, and that document is known as DOD Directive 3000.09, which was signed on November 21st, 2012, and then was reissued with minor administrative revisions in 2017, and I would just say this is a very public document. You can all find it if you Google it, and it's very short and very clear, and I encourage you to read it.

The purpose of this document is that it establishes three key ideas. The first is -- and I'm going to quote directly from the policy -- it says it establishes DOD policy and assigns responsibilities for the development and use of autonomous and semi-autonomous functions in a weapons

system, including manned and unmanned platforms, and this document also establishes guidelines designed to minimize the probability and consequences of failures in autonomous and semi-autonomous weapons systems that can lead to unintended engagements.

So basically there are four things this document says, and, again, I'm just going to give four short excerpts here that basically give you the general idea of what the policy says to the department.

First, it is DOD policy that an autonomous and semi-autonomous weapons system shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force.

Second, systems will go through rigorous hardware and software verification and validation, and realistic system and development and operational test and evaluation.

Third, persons who authorize the use of, direct the use of, or operate autonomous or semi-autonomous weapons systems must do so

with appropriate care, and in accordance with the law of war, applicable treaties, weapons system safety rules, and applicable rules of engagement.

And, last, autonomous or semi-autonomous weapons systems intended to be used in a manner that falls outside of these policies must be approved by the undersecretary of defense for policy, undersecretary of defense for acquisition technology and logistics, and the chairman of the joint chiefs of staff before formal development and again before fielding.

Just for a little added context, those three individuals names are three of the most senior officials in the department.

So the third area I want to cover -- and this is almost unique, it's breaking news here -- is what the AI strategy has to say about this topic. So this recently released and unclassified summary of the strategy -- and also it's public, so please feel free to look it up -- and there's a section in it, one of the five pillars of the strategy, which is leading military ethics in AI

safety, so I just want to read a -- a short quote directly from the strategy on that subject and that pillar about leading ethics and safety.

It reads, "The department will articulate its vision and guiding principles for using AI in a lawful and ethical manner to promote our values. We will consult with leaders from across academia, private industry, and the international community to advance AI ethics and safety in a military context. We will invest in the research and development of AI systems that are resilient, robust, reliable, and secure. We will continue to fund research into techniques that produce more explainable AI, and we will pioneer approaches for AI test evaluation, verification, and validation. We will also seek opportunities to use AI to reduce unintentional harm and collateral damage, that increase situational awareness, and enhanced decision support. As we improve the technology in our use of it, it will continue to share our aims, ethical guidelines, and safety procedures to encourage responsible AI

development and use by other nations." And, again, that strategy was signed by the secretary of defense.

So the last thing I'm going to cover is what else is in that strategy about ethics, so there is a section about a couple pages directly focused on this issue, and I just want to basically say the six things that the department has committed to as part of the AI strategy.

First, developing AI principles for defense, again, which is what we're here to do; second, investing in research and development for resilient, robust, reliable, and secure AI; third, continuing to fund research to understand and explain AI-driven decisions and actions; fourth, promoting transparency in AI research; fifth, advocating for a global set of military AI guidelines; and, sixth, using AI to reduce the risk of civilian casualties and other collateral damage.

And we will post that on our website, so if any of you find that was helpful, which I hope you did, you can use that as a

reference.

And now we get to the real action, which is where you guys get to contribute.

Yeah?

DR. MURRAY: It's 3/14, 1:59, so I'd just like to recognize Pi day today at this particular time, so --

MR. MARCUSE: It's great to know that you all can feel at home with the board members, so they are all members of the same tribe here. Excellent.

So now we will hear from the audience for the rest of our allotted time, and many of you submitted your comments to us on-line or RSVP-ed, and we'll certainly do those who have done so. If we get to those comments, then we'll use the comment cards. I -- I've really explained all this. Let's -- we're just going to get -- let's just get straight to it for now. I think that's great. So -- perfect. All right.

So the first -- first comment we -- we received was from Elli Niewood. Is Elli here? Perfect. Thank you. And after Elli, is William Powers here? William? Okay.

Over to Elli. Thank you. Well, can I continue this later? Elli, over to you.

MR. NIEWOOD: Hi, I'm Elli Niewood, here from the MITRE Corporation.

So, first of all, thank you to the board and to Josh for the opportunity to sort of briefly touch on how we see ethical considerations for military application of artificial intelligence.

At MITRE, we are very committed, both to ethical approaches to modern warfare, but also to enabling our service men and women to have at hands for the best technology that they need, both for the mission and for their own protection, and I think, clearly, as Josh has articulated, artificial intelligence impacts both of those commitments to a great extent.

AI is a key emerging technology that we think can help the joint force fight and win in future wars, yet at the same time I think it's clear that DOD has struggled to field relevant capabilities leveraging this technology, and I think that's for a number of reasons.

One is because I think a lot of the development in this technology comes from areas outside from -- you know, outside the traditional laboratories and companies where DOD has gotten its capability. Some of it, I think, revolves around challenges associated with dirty data sets, with complex system dynamics, and some of which revolves around concerns about how to use this technology in an ethical and open way, and so, clearly, that it needs to be accounted for to move forward.

At the same time, you know, from an ethical perspective, we believe AI is similar to a host of technologies that have preceded it and have been fielded and used in ways that -- that have been ethical, and we believe that integrating AI into military systems and operations -- again, as Josh alluded to, as the strategy alluded to -- in many ways we think can help reduce civilian casualties while at the same time providing a critical military advantage to our forces.

Just as an example, if you take Claymore mines, remotely triggered

anti-personnel devices, that were not banned by the Ottawa convention, were used heavily, for example, in Vietnam, what if a device like that had a sensor that allowed, you know, to determine before detonation if the target was adult sized, was carrying a weapon.

Take the -- you know, the downing and -- tragic downing in 1988 of the -- of the Iranian airliner by the -- by the USS Vincennes, the crew of that ship was forced, you know, to take action, really make a split-second decision, about whether or not to engage an unknown aircraft before they fired that missile. If they had had an AI-based seeker on that missile that could have distinguished after firing whether it was a civilian airliner -- aircraft, you know, we think that would have, you know, helped save lives in that case.

These examples, I think, are kind of, like, two of the three points we would like to bring to your attention, you know, as you move forward in coming up with principles.

The first point is that AI does not

fundamentally change, we believe, the way that we employ advanced weapons. Many of the weapons in our inventory today already select an end point or home in on a target within some given set of constraints after they're fired.

The Tomahawk cruise missile uses seekers and guidance algorithms which correlate the terrain to digital maps.

There are air-to-air missiles that lock on after launch so that the pilot fires them with some expectation of what they will engage, but without knowing for sure what it will do. All of these weapons make autonomous or semi-autonomous decisions about where they go or what to do once a human has decided to go ahead and launch them.

With AI technologies, we clearly have less visibility real-time into what the weapon will decide. We may have more difficulty testing the weapon because of the complexity of the -- what the AI will do, but at fundamental level the human has -- you know, is giving up control of decision making in ways that are consistent with existing

weapons, and that launch decision, with or without AI inside the weapon, obviously needs to be done in an ethical way that balances the risk to others with the risk to the war fighters, and that's been true for a long time.

A second point that I would like to make highlights that, you know, the human is not the ideal decision maker. I think that example with the USS Vincennes makes that clear. According to some reports, the Aegis weapons system on that cruiser knew that there was a civilian transponder code being squawked by the airliner, but the humans did not have time to take that into account as they were making -- as they were making a decision. So used properly, AI technology, we believe, can lead to better decision making, lead to reductions in errors, resulting in less collateral damage, results in less -- fewer unnecessary civilian casualties.

The last point that I think Josh also touched on is that AI technology is not primarily focused on the pointy end of the

spear about launching weapons; far from it. The applications that DOD is considering revolve around better maintenance, around fusing together data from a variety of sources about finding signals and high volumes of data while making strategic decisions.

In closing, we believe it's important to remember there are three ethical commitments we must balance in any set of principles to be developed. We have an ethical responsibility to minimize civil casualties; we have an ethical responsibility to our fellow citizens to find ways to use AI to enhance their security; and we have an ethical commitment to our soldiers, sailors, airmen, and Marines who put their lives at risk to best protect them and give them the capabilities that we need. We think AI can help with all of these. Thank you, very much.

MR. MARCUSE: Thank you, very much, Elli.

Any clarifying questions? Okay. Great.

Okay. So let me just ask, is Kate Crawford or Nick Sinai here? Kate? Nick? Okay. Great.

So next I'm going to ask Bruno to read the statement from William Powers from the MIT media lab.

And I will ask, is Michelle Kinsey here?

Aaron Johnson? Aaron, great, if you wouldn't mind getting ready at the mic, that would be awesome.

Over to you, Bruno.

MR. BRUNETT: Okay. This is from William Powers from MIT media lab. The comment is the text of a January 7th, 2019 op-ed that he co-authored in the Boston Globe called "Beware Corporate Mining Money, Machine Washing of AI."

Back in the late 1960s and early '70s, when the fossil fuel industry and the other corporate polluters came under fire for harming the environment, the polluters launched massive ad campaigns portraying themselves as friends of the Earth. This cynical practice was later dubbed green

washing. Today we may be witnessing a new kind of green washing in the technology sector, addressing widespread concerns about the perniculous (sic) down sides of artificial intelligence, robots taking jobs, fatal autonomous vehicle crashes, racial bias in criminal sentencing, the ugly polarization of the 2018 election. Tech giants are working hard to assure us of their good intentions surrounding AI, but some of their public relations campaigns are creating the surface illusion of positive change without the verifiable reality, call it machine washing.

Last year, Google posted a list of seven AI principles, beginning with be socially beneficial. Microsoft published "The Future Computed", a book calling for human-centered approach to AI that reflects timeless values, and launched a program to support developers working to meet human -- humanitarian needs.

Germany-based SAP, one of the would's largest software companies, now has an AI ethics advisory panel that includes the

theologian -- sorry -- and political scientists and bioethicists. On seeing these initiatives, the natural response is to applaud. If the most powerful tech companies are on the case, surely these problems will soon be solved, or will they?

Facebook's response to the intense public scrutiny it has received since the election has been to treat it as its public relations challenge, and after a sell off of stock in the wake of the Cambridge Analytica scandal in the early 2017, Facebook spent $1.7 million on an ad campaign in subway stations and trains in the Boston area. The slogan was "The Best Part of Facebook Isn't on Facebook," and the accompanying images showed people engaging in healthy, fun, off-line activities such as hiking and dancing. The message, Facebook is all about making our world better, more -- and a more harmonious place; yet, as the New York Times recently reported, the company had also hired lobbyists and opposition research firms to combat Facebook's critics, shift public anger toward rival companies, and ward off damaging

regulation.

As experts on the societal effects of ethics of AI, a term that broadly refers to all technologies that use decision making algorithms, we are keenly aware of how much work remains to be done in understanding how this new form of intelligence work -- works once it's released to the real world.

The tech industry has a long history of humanistic intentions and pronouncements, and, in fact, is responsible for all kinds of progress; yet, somehow, we've gotten into the most serious AI crisis since -- since the dawn of these technologies, and as with climate change and environmental degradation, if we leave oversight of intelligent machines solely to the companies that build and sell technologies, we'll see many more crises in the coming decades.

MR. MARCUSE: Great. Thank you, very much, Bruno.

So I just want to remind everyone that we really want you to use these comment cards and hand them to my friend, Aaron, here, on the corner, and will give them to me, because

I would really love to hear from all of you who came all this way today, and not just have Bruno read all of our friends who have written in, so please don't hesitate.

In the meantime, I will ask, is Bobby Cunningham present? Bobby? Or Matthew Dodd? All right.

Well, you know what to do next, Bruno.

And I'm very glad to hear from you in person, Aaron. Raise your hand.

MR. JOHNSON: Hi. Thank you all for -- for coming, and I think it's great that the Defense Innovation Board is -- is tackling this important issue. I especially want to thank Dr. Murray, whose textbook I used in my class in the fall.

DR. MURRAY: Can you get closer to the mic?

MR. JOHNSON: Yeah.

DR. MURRAY: Thank you.

MR. JOHNSON: Bring that up here. My main concern is with the -- the use of lethal autonomous weapons systems. The issue of AI and ethics is much broader than that, and -- and as you mentioned earlier, that the -- the

many -- the places that the DOD is going to be using AI is much, much broader than this.

The distinction I want to make in particular about lethal autonomous weapons systems is thinking about moral decision making and moral actions, and so, just because if -- if a human had done the same action as an autonomous system, if -- if they had -- if a weapons system does the exact same actions, it doesn't mean that it was a moral decision process, and that thinking about lethality decisions is something that's much deeper than -- than just sort of what the final result was, and so, I wrote probably more coherent stuff in my written comments, but that was sort of the main point that I wanted to bring up, was that these sorts of ethical decisions can't always be measured, and we can sort of just take the weapons system and put it through some kind of ethical tests and make sure it doesn't kill any civilians or -- or whatever sort of test we come up with, because that's not what defines the moral action, and it has more -- actions have more to do with the internal

process and the decision to trade off between, you know, possible actions and -- and make -- potentially make sacrifice from that, so that's all I wanted to say. Thank you.

MR. MARCUSE: Okay. Thank you.

Clarifying questions? No? Great.

MR. BRUNETT: Comment is from Michelle Kinsey, Boston University.

The department of defense, in its push to expand the intelligence, autonomy, and mobility of systems supporting the dismounted soldiers' real-time tactical decision making capabilities will have to address a number of design challenges related to the safe deployment of artificial intelligence learning modules -- models and techniques. Machine learning models are often trained using private data sets that are very expensive to collect, or highly sensitive, using large amounts of computing power.

The models are commonly exposed either through on-line APIs, or used in hardware devices deployed in the field or given to the end users. This gives incentives to

adversaries to attempt to steal these ML models as a proxy for gathering data sets.

While API-based model exfiltration has been studied before, the theft and protection of machine learning models on hardware devices have not been explored as of now.

MR. MARCUSE: Is -- is Alexis Priest Simpson present? Alexis?

Okay. Do you all have these from outside? There were a stack of these outside. Should we bring these inside for you? Do you need them? Please hand them to Aaron, or just waive them in the air. I did not expect such a shy and retiring bunch.

In the meantime, though, I will ask, Bruno, if you would now please read Bobby Cunningham.

MR. BRUNETT: Bobby Cunningham from Omnity.

Omnity is a self-assembling knowledge curation and discovery platform that fuses advanced natural language processing, machine learning, linguistic block chain, and graph math. Omnity detects similarities across diverse intelligence sources, driving rapid

discovery and insight, and has co-founded the Wisdom Tech Society to provide a framework of ethical data curation as data is transformed into wisdom. I've got to slow down a little bit. As artificial intelligence emerges as a means to find patterns, and perform analytic and massive data sets, many organizations, companies, and governments are seeking to leverage this powerful technology for their own applications.

However, it is critically important that those that seek to use this technology also better understand both the strengths and weaknesses associated with the data processing strategies, algorithms, and business practices and enabling machine learning. Such understanding is complicated by the massive hyperbole expressed by many companies, often further amplified by journalists who do not understand the topics about which they are writing. Taken together, these forces create unrealistic expectations and even fear.

Properly applied machine learning can be a useful tool for exploring and discerning

patterns in big data, where human inspection of massive data is not scalable. When seeking to categorize or otherwise sort data into sets that do enable human insight, machine learning processes offer a useful augmentation for human intelligence.

One can think of the algorithm -- algorithms driving these sorting processes as a form of high dimensional curve fitting; that is, applying a structural analysis to find the underlying patterns in a large data set. Used in this manner, AI technology is well suited for useful application.

What computers do poorly is -- what computers do poorly is make judgments. Computers do not understand irony or sarcasm. Computational processes do not well enable abstraction of ideas, generalization, or creative thinking. These areas remain in the realm of the human mind, relying on computer processes, with expectations that the computer can be creative, generalize, or make judgments will lead to disappointment and frustration.

Understanding the limits of machine

intelligence is critical for the effective use of these technologies. It is important to note that the output of a computational process is limited by the quality of its input data. This clearly applies to consistency of data formatting, completeness of data records, and other quality control metrics; yet it also applies to the ethical sourcing and curation of the data sets, themselves.

Where data is sourced from people, organizations, companies, or government agencies, in each case the data sources should be derived in a manner that is both morally appropriate and legally compliant.

Finally, the use of data to form insights is a three-strep process. Data can be defined as numbers, facts, and figures, such as sensor readings or the monitoring of vital signs in a patient; yet, data alone does not afford insight. When data is contextualized, it transforms into information. When that information is contextualized, that information may form wisdom, leading to actionable insight. Each

tier of this transformation process is vulnerable and must be safeguarded ethically so that those actionable insights are consistent with the moral framework of our civilization.

At this time, what is most needed is a framework of ethical data curation as data is transformed into wisdom. This is why we at Omnity have co-founded the Wisdom Tech Society. As the Defense Innovation Board considers ways in which to advise the Pentagon with respect to the ethical use of artificial intelligence, we urge you to meet with Omnity and Wisdom Tech Society which will demonstrate technology that can transform the ways which our nation conducts the intelligence gathering methods.

MR. MARCUSE: Thank you, Bruno.

Aaron, maybe we should hand the cards out again, just a subtle encouragement. Have we received any yet?

UNIDENTIFIED MALE SPEAKER: We have one. Wave a hand over there. If you have a card --

MR. MARCUSE: Great. Yes. If you

have -- yeah, if you have a card, can you hand -- great. Oh, this is very exciting. Wonderful. Thank you. I'm so relieved. Come to this mic. So, Maggie, you -- you've rescued me. Thank you. Please take the microphone. That's phenomenal. Thank you, very much.

MS. OATES: Hello, my name is Maggie Oates. I am Ph.D. student here at Carnegie Mellon in societal computing, and I work in Cylab, which is the cyber security laboratory here.

Let me first say that I disagree with the very enterprise and existence of the Defense Innovation Board as it seems like a tool to lend credibility to the project of advancing military efficiency, and further escalating the baseline that we consider defense.

Second, I object to CMU's heavy involvement with the DOD, and I'm hard pressed to name ethical and responsible use of AI at all.

That said, I would like to focus on something else today, and that is what I view

as a externality of the project of AI in the DOD, and that is the growth of civil surveillance, both domestically and abroad. The development of machine learning algorithms relies on massive amounts of data, of course, and more methods are being developed to, first of all, reduce the amount of data required, and to reduce the amount of labeled data required. These methods often correlate with having the downside of being hard to explain and more difficult to verify, making them an unlikely use in the DOD's context. Beyond that, the project of labeling data often rests on -- rests on exploitative labor practices. So I stand here to assert that any responsible principles must address the effects that DOD AI will have on surveillance, not only from the state, but also from the tech companies that will be the first line in building that AI. This topic is absolutely not out of the scope. Thank you.

MR. MARCUSE: Thank you. I hope we have a few more, and in the meantime, Matthew Dodd from the National Institute of

Health.

MR. BRUNETT: Matthew Dodd from the National Institute of Health.

I would like to thank the board for their leadership and sage counsels to the department, industry, academia, and, indeed, the nation. It is a profoundly overwhelming undertaking of responsibility. I thank you all, and I am humbled by the purity of competence demonstrated by the board. This stands in stark contrast to the anti-expert virus that has infected every corner of society, one participation trophy at a time.

While demonstrating that, primarily one has a duty to serve their country and one has a unique -- has unique experiences and tools with which to innovate the notion of service. One has the moral imperative obligation to find out how to manifest it, innovate, so as to be of service. Thank you, again.

MR. MARCUSE: Great. Bruno, thank you, very much. Next we will hear from Dave Zubrow from DSEI.

Oh, yes, hold on, as long as you

promise to give it back to me when you're done, thank you, very much, and after -- and after Mr. Zubrow, we'll hear from Pat Houston next. You can go to the mic, if you want, sir.

MR. ZUBROW: Hi, I'm Dave Zubrow from DSEI, and, Josh, your remarks at the beginning struck me, because there was a lot of emphasis placed on adhering to the law and the legal framework that's been built up around conducting military operations and war, and I just -- it struck me that our legal framework and our laws may not adequately express our values and morals as a society, and so, simply adhering to the law may not be enough, so I would ask that in the work and deliberations of the board and others, thinking about how to deploy artificial intelligence into -- into our lives, as well as, you know, national security, that we push at that and say is that enough, is that the right framework, is it really expressing our set of values as a society, so that's my comment.

MR. MARCUSE: Thank you, so much.

Any clarifying questions? Thank you, very much, sir, greatly appreciate it.

Bruno, if you would, Julian Klein, and then over to you, sir, right after that.

MR. BRUNETT: Julian Klein from Klein Studios, LLC.

Dear DIB, here many my suggestions: One, while AI is very powerful for information gathering, data analysis, and reactive cyber defenses, any -- any AI dealing with humans should have a margin of error for living creatures' mistakes, confusions, and improvisations. No human nor animal should be held to mathematical expectations.

Two, humans cannot be backed up, re-downloaded, or rebooted. Any physical world AI should do its best to preserve human life in a humane way, despite any further coded tasks.

Three, any powerful AI should come with equally powerful AI minimum of three, which can create a checks and balances scenario. If one AI begins acting strange, the other two can fix it with permission

override keys and regulations. One cannot override the other two. In the event all three are corrupted, the owner or developer should have a kill switch and back door code to -- to delete all three.

Four, we need a clean -- we need to clean the internet with AI. This sort of cyber regulation will rely on a communicative tech society. Segregating tech expert and what they know may seem safer, but is detrimental to our minimum bar of tech education. A collaborative cyber community will have the creative answers to defense issues.

Five, using AI to affect people's ways -- way of life, sway people's opinions, or cause chaos or confusion is different than a media campaign based of AI-collected and configured data. The former is systematic, intentional, and cultural intrusion. The latter is an educated broadcast. Thank you for your time and work.

MR. MARCUSE: Bruno, because it's so brief, why don't you just quickly also read Joshua Darrow.

MR. BRUNETT: Joshua Darrow, Department of the Navy.

Innovation begins and ends with technical capability. Over decades, the government has out sourced its most technical work, eroding the technical skill of -- of the government workforce. To have an innovative government workforce dealing with the distribution of technical work, i.e., designing, building, testing, redesigning, and testing is essential.

MR. MARCUSE: Terrific.

Over to you, sir.

MR. HOUSTON: Good afternoon. My name is Pat Houston, and I'm a soldier, and I would like to provide four observations from my perspective in my personal capacity based on five tours in Iraq and Afghanistan, and approximately three decades in uniform.

First, I want to assure those here today that from what I've seen, the Pentagon is very sensitive to these valid concerns, and is deeply committed to addressing them. Leaders at all levels are totally dedicated to ensuring that AI-enhanced systems are

developed and used in compliance with the law and with ethics, and also done in a responsible manner.

Number two, I acknowledge that no AI system will ever be perfect, but I also would offer that no military system is perfect. In fact, one of the most unpredictable systems we can employ is the individual soldier. As we all know, human behavior can never be reliably predicted, especially when someone is tired, cold, hungry, and scared, as is often the case with -- with soldiers, but I do think AI can help.

One of the measures for determining whether to leverage AI is to ask whether it's as good as, or in some cases better than, what humans can perform by themselves, so we're looking at human machine teaming that leverages the best of both, and that helps preserve appropriate levels of human judgment to address some of the concerns and issues that have already been raised here today.

Third, just to address one concern head on, this concern about killer robots, the law of war is very clear that commanders

can never unleash systems over which they'll lose control, and that commanders always are responsible for any weapons that they employ.

Fourth, and finally, I would suggest that cooperation between the government and academia and industry is absolutely essential to the responsible way ahead in addressing all of these issues. The best and brightest AI researchers out there should insist and have a right to insist on legal and ethical conduct by the governments that they're working with, and once you have that -- that confirmation of ethical and legal performance by those governments out there, then I don't think those AI researchers should boycott the endeavor, as some have suggested that they should do. In my -- in my view, if the best and brightest AI researchers out there who are concerned about ethics boycott the process, that's going to just leave a void that would be filled by other AI researchers who are less ethical or less capable, or both, and I think that would be a recipe for disaster. Thank you.

MR. MARCUSE: Thank you, sir.

Any clarifying questions? No? Great. Okay.

I think I've seen some evidence of cards stirring, so please just feel free to just push -- you know, give them to your neighbor to pass them to me. We can be as casual as this podium stage arrangement permits. That's right. All right.

So I'll ask, Bruno, if you would read the statement from Thomas Creeley, and then Henry Hargrove, if you would, please feel free to approach the microphone.

MR. BRUNETT: Thomas Creeley, United States Naval War College. The U.S. Naval War College has recently established a special graduate certificate program in ethics and emerging military technology. As its director, I work with a dozen competitively selected students to engage in ethics and technology-related course work, as well as conduct research on the ethical implications of emerging technologies. Each student produces a lengthy professional paper, analyzing some aspects of the ethics technology nexus, many of them dealing with

the various forms of AI. We have developed connections with the DARPA, ONR, Boston Global Forum, and the director for defense intelligence, and a number of academic institutions exploring the ethics of AI.

MR. MARCUSE: Please.

MR. HARGROVE: Thank you for allowing the time for public comment. I wanted to reaffirm the other gentleman's comments before that one about the soldier, in the course of understanding the soldier's behavior, which my comments sort of echo.

So in my personal opinion, I -- I just want to echo the need that you've already identified for realistic and continuous testing evaluation, and the reason for this is upon the systems will co-exist with humans, individually and in large social networks, and if we are truly pursuing an understandable AI, the response to stimuli shaped by this humanity, then our DOD model and simulation and live and virtual constructive programs need to keep pace in portrayals of the humans and human technical behavior in those simulations.

Secondly, this testing evaluation improvement should be continuous, and that's because our society and the society of our peers -- correction, allies and adversaries is also changing, and so, that behavior is dynamic and depends -- demands that we are current and honest in our own assessments of that portrayal. Thank you.

MR. MARCUSE: Thank you. Good. Great. Okay.

Now, if you wouldn't mind, that would be great, and bring that back up, and does anyone else need a comment card? Great. Thank you, very much, sir. If you wait, you can just -- that's perfect. You can approach the microphone there on the -- right here would be also excellent. Thank you.

MR. CHEN: Thank you, and hello. My name is Kenny Chen, and I would like to share my appreciation that explainability and transparency have been emphasized in the DOD's initial AI strategy, and because this strategy and its components remain at -- in understandably at -- at nascent phase, it's understandable that those terms are not yet

clearly or rigorously defined, and so, to that end, I encourage the DOD to pay especially close attention to the question of to whom these AI systems are designed to be explainable and transparent, because there are substantial differences between how accurate information might be delivered to a computer scientist, versus a war fighter on the field, versus a journalist or a politician, and the DOD cannot be too careful in assuring that misinformation and the risks of misunderstanding are minimized.

Having talked to people at the U.S. -- U.N. Office of Disarmament Affairs, or considerations -- given the pace of decision making within a highly automated, you know, decision system when it comes to international conflict, setting those kinds of standards and understanding will be existentially important, and just as we are currently emphasizing interoperability across information systems and devices and technology, we should apply the same values and scrutiny to the way that organizations communicate with one another. Thank you.

MR. MARCUSE: Great.

UNIDENTIFIED MALE SPEAKER: Yeah, it's not quite done, but --

MR. MARCUSE: Well, sure. That's great. We just -- we just want to make sure we have your -- your contact information. That's great.

All right. Bruno, Mark Gubrud's statement.

MR. BRUNETT: Mark Gubrud, University of North Carolina, Chapel Hill.

We stand today at the start of a revolution, the rapid advance and wide use of AI. Because this technology replaces human intelligence and judgment, it has the potential to cause catastrophic errors with consequences in proportion to the responsibilities being delegated to machines. The specific causes of error in AI may be foreseen, but in general will not be and may not even be identifiable.

In complex systems, it becomes, in principle, impossible to foresee all exceptional situations that may and will arise. Computer algorithms are particularly

brittle, but also complex networks of analog and living systems exhibit unpredictable collective behavior and sudden crises. As AI advance towards human-level capabilities, its potential for instability is clear.

The most severe danger arises from the unforeseeable, untestable interactions of networked complex competing and adversarial systems. Experience with such networks such as the instability of high-speed trading, which has produced several hugely expensive stock market flash crashes has demonstrated the likelihood that confronting and interacting adversarial networks will erupt into crisis or open combat spontaneously -- spontaneously, or as triggered by unforeseen circumstances, and in any case, once ignited into a condition of ongoing violence, may escalate -- and execute and escalate that violence so rapidly in such a complicated and opaque way as to resist or frustrate any human attempt to intervene.

The ongoing confrontation, competition, and adversarial -- adversarial nature of such systems, and in the

adversarial relationship of their creators, contradict and will frustrate any effort to coordinate between them so as to mitigate the risks of unauthorized or uncontrollable conflicts.

Why would nations undertake to construct and rely on such obviously dangerous systems? For the same reason as in the Cold War, as in the arms race today, and under the competitive pressure of an advancing technology that is already able to aggregate and correlate more data than any human could, and is increasingly able to integrate that information and make high-level decisions, particularly when signals unambiguous, more rapidly than any human.

We must avoid taking that road, but, in fact, it is the road we are already on, so we must avoid going further. The global community must undertake ambitious arms control initiatives, including a mandate of real-time accountable human control over all weapons systems. The DOD and the U.S. cannot do this alone, but America's preference

should be for arms control, we should say so, and everything we do should be consistent with that. Unilaterally disarmament or renunciation of strategically decisive weapon technology would -- would be -- would not be effective or possible, but the opposite extreme of trying to win the arms race should be equally strongly rejected.

Questions about AI and autonomous weapons are too often framed only in terms of ethics. It is -- is it right to use such weapons? We need to consider the ways in which these weapons are eroding our control and creating a threat to ourselves. We must avoid accelerating the arms race toward the loss of human control and the occurrence of war by accident or misconceived design. Is it ethical to lead a global race to oblivion, instead of leading toward a strong regime of binding verified arms control global governance than human security?

MR. MARCUSE: Thank you. That's all. Great. Thank you. Excellent. Thanks a lot.

MS. GALLAGHER: I'm April Gallagher, I'm with the DSEI, and when we're talking

about ethics and responsible use of AI and machine learning, it -- oh, closer? Louder? Okay. All right.

The ethics and responsible use, where I see that, it's really in the details and the implementation, that that's where the ethics are really going to happen, and I've got a list of a couple of examples just to highlight this.

So any time you're talking about machine learning, or AI, there's going to be a base rate of error, and if you took an intro to statistics class, 5 percent is acceptable error. Well, so, Tumbler recently put a adult content filter out, and let's pretend that they actually got their 5 percent error in this classifying images. They didn't get anywhere near that, but if they had gotten the 5 percent error, that's still billions and billions of images that were misclassified, and millions and millions of unhappy users. Now, 5 percent in a DOD application, that's not okay.

So, now, the second point, training data has to match the conditions of use. So

Alexa and Siri and lots of other voice applications are really bad with women's voices because their training data had a bunch of male engineers in it. They didn't collect samples from women, and so, women with slightly higher voices, the equipment doesn't work as well for them, and so, you have women soldiers, that's going to be an issue, but let's take that a step farther. If you train a system in the United States, let's say it's going to detect a threat, whether or not a particular person is a threat, based on the emotions that they're displaying, if you train that on data collected here and then you deploy it in theater, it's not going to work because the training data doesn't match the use case, so, again, the details of how it was put together and how it was designed, that's going to change the responsible use. I'm going to skip one that's written down.

The last point --

MR. MARCUSE: You have time for it, if you want.

MS. GALLAGHER: I have time? Okay.

All right.

Current efforts in explainable AI and people who are the end users of these things, those current efforts are going to be insufficient because -- and there -- there's a lot of proof here. One study in particular took a best-case scenario, all the studies that are working on explainable AI, they work perfectly, they get the results, they presented such results to a largely college-educated audience, and less than 60 percent of those users could understand the output, and you've seen the study, so we've got to think about the human computer -- human AI interaction. Explainability is not enough. We have to think about interpretability, and commuting that efficiently.

Last point, many predictive systems -- so talking more about machine learning here -- they can be harmful or beneficial depending on what you do with the predictions. So if I said I had a system that predicts whether a particular person is going to be a criminal, I think most of us

would be horrified if we took that prediction and arrested them and locked them up, a la Minority Report, that's kind of -- they didn't do anything yet, but such systems are actually already being deployed in high schools.

If we predict whether or not a particular person is a threat of dropping out, if they're a success or at risk, what happens with that information? Is that student supported? Is that student given extra support so that they can actually complete high school? That's a good use. If they are treated as "Oh, you're already a dropout," then that's irresponsible, and so, what we do with the predictions, it's not the ML, it's what we do with the information.

So all of these examples are to highlight that the ethics and responsible use, it really comes down to the details in each particular case. Thank you.

MR. MARCUSE: Thank you.

MR. MEDIN: Sorry, could I ask a question? You --

MS. GALLAGHER: Yes, sir.

MR. MEDIN: -- commented about 5 percent as acceptable.

Where did that come --

MS. GALLAGHER: Oh, that's -- so that's the -- the standard T test. It's been kind of the -- and the -- the 5 percent, it -- it comes from just a tradition of statistics. If you do something -- like, if you flipped a coin, when people start getting suspicious that maybe it's not fair is about -- "Oh, that seems weird," it's at about a 5 percent error, and so, that's a good threshold for a T test, but what's happened is because that's the standard statistical threshold, people have, without thinking about it, transferred that 5 percent over to all kinds of machine learning application as, "This is a great threshold," and it's not for a lot of applications, and so, that's an example of carelessly transferring something from one domain to another domain.

MR. MARCUSE: Thank you.

MR. MEDIN: Thank you.

MR. MARCUSE: Someone here is thinking

about whether they would like to make a comment. I'll try to nudge you to the side of giving it a shot.

While you're thinking about it, Bruno, if you wouldn't mind reading the statement from Michael Duggan -- Duggan?

MR. BRUNETT: Michael Duggan, Booz Allen Hamilton, this is text taken from a Booz Allen white paper called "Analyst 2.0, Redefining the Analysis Trade Craft." It's a little lengthy.

MR. MARCUSE: Well --

MR. BRUNETT: Analyst 2.0. I'm going to read --

MR. MARCUSE: Sure, just a reasonable --

MR. BRUNETT: Yep.

MR. MARCUSE: -- quantity of it, but no more than five minutes.

MR. BRUNETT: Listen for the microphone to be hit.

MR. MARCUSE: Yeah.

MR. BRUNETT: Making sure official intelligence works for the mission. Artificial intelligence and other advanced

analytic approaches are rapidly becoming integral to the intelligence mission. As our nation's security posture grows more complex, and we need to keep our eyes on more people and places, the volume of critical intelligence data is expanding exponentially.

It is becoming difficult for analysts alone to keep pace. There is simply too much data to be brought together and analyzed in the short time frames required by the mission. The military and intelligence communities recognize that advanced analytics hold great potential, and they are the beginning -- and they are beginning to adopt these emerging technologies.

With AI, for example, instead of an analyst spending hours poring over a stream of satellite photos looking for any significant changes, the computer might complete the task in seconds. This frees up the analyst to spend more time on higher-level analysis, reviewing what the computer has found, and then preparing reports for decision makers that are both timely and comprehensive.

In essence, the machines are doing what they do best so that people can do what they do best, but this shift, turning over much of the repetitive work to -- to a computer, is also presenting defense and intelligence organizations with a significant challenge.

How can they be sure that the outputs of the computer are both accurate and relevant to the mission? How can organizations be confident that the analytic tools are working for them? The stakes of this are of the highest order. The expertise of the analyst is vital to national security, and if it is lost or diminished in the human-machine connection, the risk can be significant. What if the computer doesn't have it quite right, and faulty analytic output -- outputs are used by commanders or other decision makers down the line?

Yet another challenge is that analysts may not accept and use AI formed -- informed analytics, either because they don't trust the outputs, or because they fear that the computers will put them out of a job. There

are already examples of this in some organizations. New technology systems are introduced with great fanfare, and then promptly ignored by analysts who are free to pick the tools they want. And yet, with the new technologies, decision makers won't be able to take full advantage of the available data, something that is essential to keep pace with today's threat.

Unfortunately, most current approaches to AI and other advanced analytics don't resolve these dilemmas. In fact, they only make them worse. With all -- with all the hype around AI, data scientists and others are caught up in what the technology can do.

For example, they try to build better and better models for pattern recognition or object identification, but this research is largely academic and theoretical, and not tied to the specific mission at hand. Yes, the tool can look for changes in photos, but is it the kind of change that the analyst is looking for? Too often, such contextualization is missing, and when that happens, the tools simply can't be relied

upon and support decision making. Automation and speed count for nothing if the computer gets it wrong.

MR. MARCUSE: Thank you.

Any other takers of this opportunity? Once, twice? And we are going to adjourn.

So we ask you to stay seated for a moment while the board retires to the green room. If you are a member of the media and you'd like to participate in the media availability, please come forward.

Let me close by saying the comments on-line at that website, it's quite easy to do, it takes a couple of clicks, we would really love to hear from you, and let me also close by saying that we will have another listening session like this one in California on April 26th at Stanford University. You will be able to tune in to that live stream and get information about that by signing up for our invitation. We'll send that information to you.

And let me just close by saying thank you all, very much, for coming to participate in this conversation and listening to the

discussion, and -- and particularly thanks to those who went to the microphone to share their views, greatly appreciate it. Thank you all.

- - -

(Thereupon, at 2:50 o'clock p.m., the proceeding was concluded.)

- - -

C-E-R-T-I-F-I-C-A-T-E

I, Eileen L. Drake-Ober, the undersigned, do hereby certify that the foregoing seventy-three (73) pages are a true and correct transcript of my stenotypy notes taken of the proceeding held at Carnegie Mellon University, Studio Theater, 5000 Forbes Avenue, Cohon Center, First Floor, Pittsburgh, Pennsylvania 15213, on Thursday, March 14, 2019.

Eileen L. Drake-Ober
Notary Public in and for the
Commonwealth of Pennsylvania
My Commission expires: December 6, 2020

- - -