CPPI

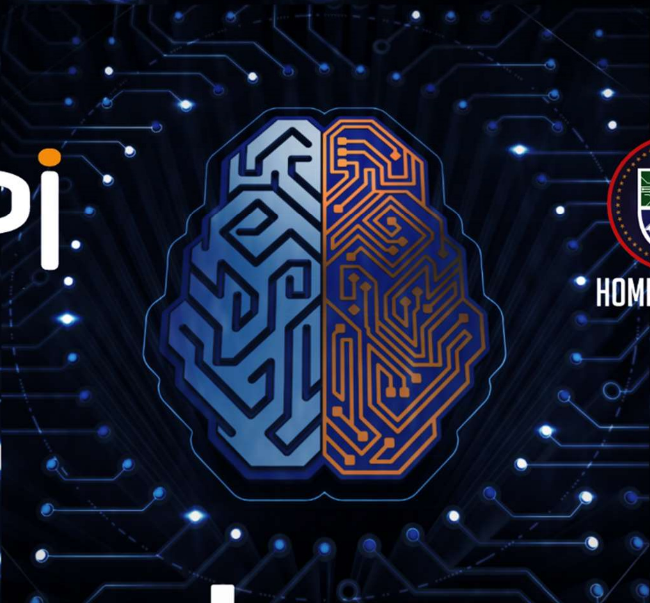

HSDF
HOMELAND SECURITY DIALOGUE FORUM

Artificial Intelligence: Transforming the Government Mission

Virtual Symposium, October 2020

Highlights and Key Findings

Event Partners:

DELL Technologies

intel

Microsoft Federal

ORACLE

servicenow

Table of Contents

Word cloud generated from event transcript

PART I: **INTRODUCTION**

Building out artificial intelligence (AI) capabilities is a national security imperative. The rate of adoption of AI across the U.S. government is increasing and the vast amounts of data collected can deliver tremendous insights to enhance security, improve citizen services, and transform decision-making processes. The confluence of AI and edge computing will enable the federal government to take better advantage of cloud capabilities and achieve better mission outcomes. The U.S. cannot afford to fall behind in developing and adopting AI technology as U.S. adversaries are investing billions to stay competitive in this area and are using AI to their advantage.

To examine these issues in detail, the Center for Public Policy Innovation (CPPI) and Homeland Security Dialogue Forum (HSDF) convened a virtual symposium in October 2020—***Artificial Intelligence: Transforming the Government Mission***. The following report includes major takeaways from a series of keynotes and panel discussions with experts from across government and industry.

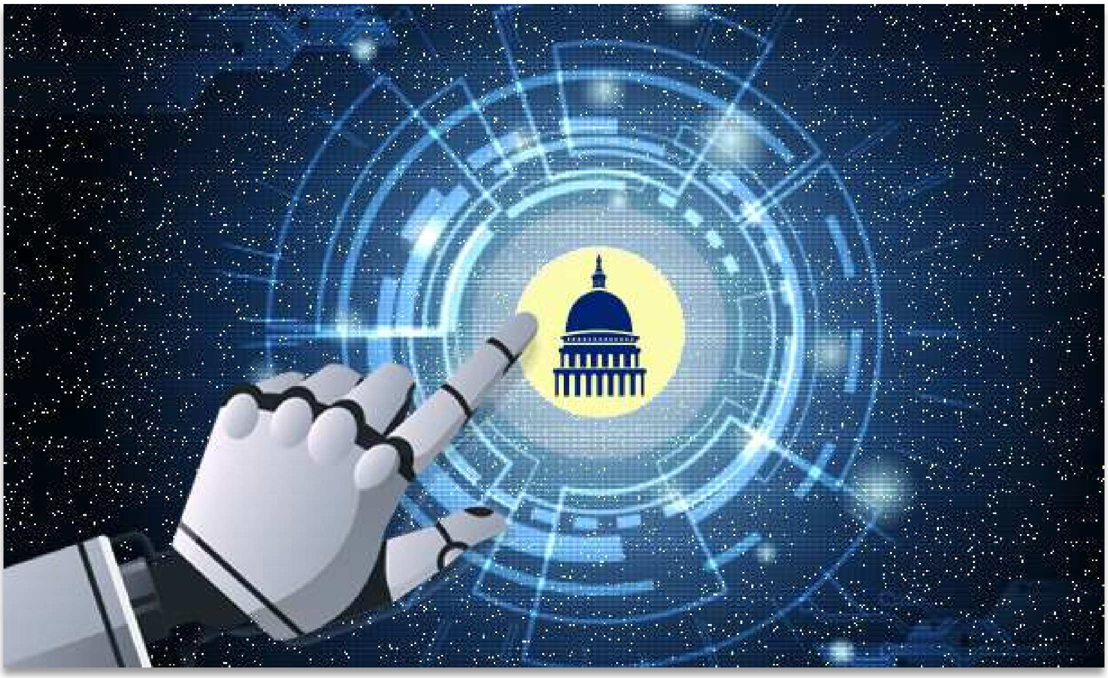

PART II: BACKGROUND ON ADVANCING ARTIFICIAL INTELLIGENCE IN GOVERNMENT

Defining Artificial Intelligence

Within the federal government and private sector there are varying definitions of AI and machine learning (ML). One baseline definition for AI/ML is “the study of computer algorithms that allow computer programs to improve automatically through experience”, which was offered by Amanda Mercier with Microsoft. Instead of hard coding and telling programs how to behave under certain scenarios, programs can learn to improve outcomes on their own. Mercier noted that this is just one branch of artificial intelligence and helps to illustrate why AI and ML terms are used interchangeably. Jonathan Alboum, Principal Digital Strategist for the Federal Government at ServiceNow, added that AI will continue to evolve and grow in sophistication. “What we think of AI today will not qualify as AI in the future—I think that is important to understand,” Alboum told the audience.

Cheryl Ingstad, who leads the Artificial Intelligence and Technology Office at the Department of Energy (DOE), offered clarification on how robotic process automation (RPA) is different from AI. RPA is about optimization of a process; it is not true AI. When an organization layers on predictive analytics, for example, now it is not just an optimized process, insights can be derived from the data to drive efficiencies—that is where the AI comes into the equation.

Executive Leadership on AI and Data Governance Efforts Across Federal Agencies

“The President’s Executive Order on Artificial Intelligence has been helpful in driving American leadership on AI, ensuring there is substantial R&D investment, and creating an AI ready workforce,” commented Ingstad. She also mentioned the National Security Commission on AI, created by the 2019 National Defense Authorization Act, has been releasing quarterly reports including one in October 2020.

For DOE, standing-up their AI & Technology Office was in response to the Executive Order’s focus on improved visibility. Agency leaders wanted to do a better job of understanding all the different AI activities taking place across the department. Ingstad said they have a six-point strategy that very much focuses on DOE’s coordination function to advance the nation’s AI leadership. This includes creating a database of AI projects across DOE, to include all the national labs; creating lists of approved AI tools; and developing a better strategy around data and related cross-cutting projects. These efforts are meant to complement the work being done at the Department of Defense Joint Artificial Intelligence Center (JAIC), the General Services Administration’s Centers of Excellence, and the National Institute for Standards and Technology (NIST).

Ensuring that federal agencies have appointed a Chief Data Officer (CDO), as enacted through legislative efforts, is a critical piece to advance AI adoption. Ingstad commented that CDOs help an agency to develop data maturity models and provide guidance about what data sets can and should be used. Mercier agreed, adding that the CDO has a full spectrum of responsibility when it comes to data from obtaining rights and access, curating the data, and identifying any anomalies or problematic patterns. “Data is the front-end of the machine learning pipeline. One important complement to the CDO is the test and evaluation function. So often our ability to have a ML capability that satisfies the user depends on how we optimize. Making sure we define these testing evaluation metrics means we’re shooting at the same target,” commented Mercier.

Another point raised by Luke McCormack, former CIO at the Department of Homeland Security, was about the importance of where the CDO role lies within an agency to have the greatest impact. Alboum, a former CIO at the Department of Agriculture, pointed out that since each federal agency is different, in some cases elevating the CDO to the Office of the Deputy Secretary might provide greater visibility across a department and potentially access to different data sets to help accelerate AI innovation.

Gretchen Stewart from Intel spoke to the importance for government to develop a comprehensive data strategy that includes security. “The government has the most amount of data, probably across the globe. You need to ensure from the beginning when you think about that data strategy, what is the security component? Is it 2-Bit encryption or not? What are the components you want built-in from the beginning? Then when you're upgrading the infrastructure, you're really ensuring that you have built-in security from the ground-up,” she offered.

The Department of Defense (DOD) recently released their Data Strategy in October 2020. Major Nathaniel Bastian with the JAIC at DOD mentioned how recent efforts have focused on data digitization since so much of the data is not ready for AI/ML applications to be layered on. He added that the JAIC has been focused on developing best practices around data model traceability and data management processes, which are shared with Dave Spirk, the new CDO at DOD.

Congressional Leadership to Advance Artificial Intelligence

Congresswoman Robin Kelly and Congressman Will Hurd introduced a bipartisan resolution calling for a comprehensive, whole-of- government approach to artificial intelligence. The resolution cites recommendations from four white papers they created in conjunction with the Bipartisan Policy Center that focus on workforce, national security, R&D, and ethics. Rep. Kelly commented that AI impacts every sector of our economy and every federal agency. Regarding the resolution she commented that, “American ongoing and future leadership in AI will be based on today’s policies and investments. We need a real and actionable comprehensive plan to engage all levers of national power to create and preserve American AI superiority.”

The House Armed Services Committee recently released a report from their bipartisan Future of Defense Task Force. “I'm proud of our task force’s eagerness to reach across the aisle and develop substantive recommendations so that we have a lasting impact for our war fighters for decades to come,” said Congressman Jim Banks, Co-Chair of the task force. He emphasized that no technology is more critical to maintaining U.S. military dominance than artificial intelligence and information has become our most important weapon. Some of the report’s recommendations include:

- Ensuring all Department of Defense programs and weapons are equipped with artificial intelligence capabilities to provide warfighters with a competitive edge.
- Creating a national STEM workforce that includes a military commissioning source for STEM talent. Congressman Banks co-sponsored H.R. 6526 that would establish a STEM Corps where tuition costs are covered for students in STEM fields in exchange for military service.
- Updating the Pentagon’s acquisition systems to stay current with emerging technologies and make it easier for companies to work with the Department of Defense. This includes accelerating acquisition timelines and ensuring quicker delivery of AI capabilities.

The Congressman acknowledged that the Pentagon is a big bureaucracy and can be reticent to move quickly to adopt new technology, but taking risks and learning to embrace “failing fast” are critical to staying ahead of our adversaries.

He also expressed serious concern about threats posed by China. In particular, the rise of China’s civil military fusion strategy and how the Chinese Communist Party is using its personnel to directly support the military. Even more worrisome is that China is on track to become both the world’s largest economy and leader in global R&D spending by 2030.

China has also sought to expand their technological capabilities through intellectual property theft. The Justice Department has said that nearly two thirds of trade secret theft cases have ties to China. The Chairman of the National Security Commission on AI, Eric Schmidt, commented that China’s view of the relationship between technology and authoritarian governance should be of high concern to Americans.

By accelerating the use of AI across the military, and better leveraging the innovations taking place in the private sector, the U.S. can stay ahead of the threats posed by adversaries and maintain strategic advantage on the battlefield.

Government AI Use Case Scenarios

AI and ML can reduce the cognitive workload on government employees and the warfighter to allow individuals to focus on things that require creativity and problem solving which computers cannot achieve, commented Mercier with Microsoft. There have been numerous examples of how AI can help with burdensome, repetitive administrative or back office tasks, but there are also some more exciting efforts underway that can have a positive impact on mission outcomes.

One example includes the First Five Consortium, which was stood-up with participation from academia, industry, and government. The name reflects the critical first five minutes when responding to a disaster with the goal to apply AI to support first responders and mitigate damage. The Department of Energy and Microsoft are two major players in this initiative, which was formed in response to a call from the White House to improve humanitarian assistance and disaster response. Currently, the Consortium is researching AI applications in the areas of wildfire prediction, damage assessment, search and rescue efforts, and hurricane and tornado response. Mercier mentioned that some disasters cost taxpayers billions in recovery, and deploying AI might help reduce cost-prohibitive response activities, such as detecting wildfire lines.

Ingstad from DOE mentioned they have been deploying AI for a wide range of purposes from detection of seismic activity to better understanding the current COVID-19 pandemic. In particular, she mentioned a partnership between the Department of Veterans Affairs and Health and Human Services to share health data and research to create treatments for COVID-19 and assist in developing a vaccine.

Other areas where AI can have a big impact is in agriculture and smart cities. Microsoft provided some examples of using AI in farming production to maximize production while using less energy. DOE discussed their efforts to apply AI to the smart grid and how it can account for renewable energy sources. "The future for delivering power is going to be dependent on AI because we cannot continue to rely on peak power, we need to manage it at the point of consumption," commented Ingstad.

U.S. Customs and Border Protection (CBP) collects enormous amounts of data, whether it is information about international travelers or cargo manifests, and there are many opportunities to apply AI for enhanced security, commented David Aguilar, former CBP Deputy Commissioner. CBP is one the largest consumers of data within the federal government and layering on data analytics has become critical to their mission.

Facial recognition technology is one area that has made great strides in recent years with improving data reliability, growing public familiarity, cheaper storage, and improvements in camera technology to ensure a near perfect match rate. All these positive developments taken together indicate a speeding-up of progress around biometrics at CBP. Robert Costello with CBP's

Office of Information and Technology commented on ongoing efforts to address privacy concerns with data. “[The public] has started to see that CBP are good guardians of this data and they can trust us with it when we’re using it to catch people that want to harm our country,” he said. There are also a lot of innovation taking place within the airports. In the near future, vetted travelers—such as individuals with Global Entry status—will be able to walk through an airport without stopping at any checkpoints by using biometric applications, commented Costello.

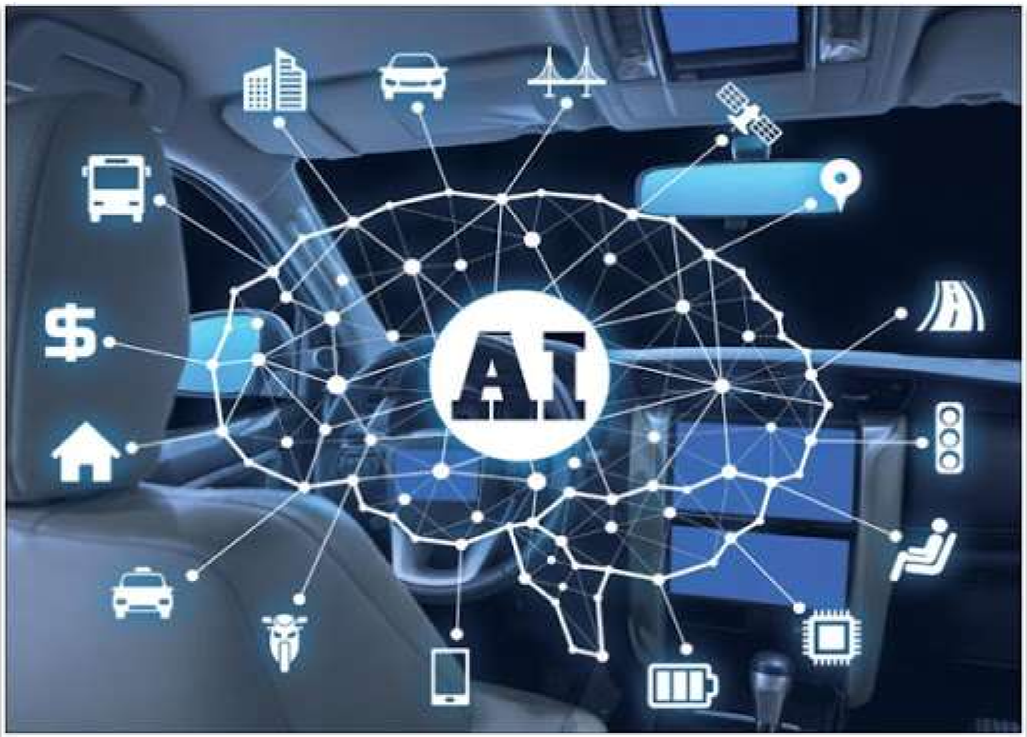

PART III: **TECHNOLOGY ASPECTS OF AI**

Distributed Data Sources and Innovations in Data Storage

Data growth has become so massive at CBP in recent years that data management has become a priority. Data classification, figuring out the best type of architecture, and determining back-end storage needs are all top priorities, commented Sunil Madhugiri, Acting Chief Technology Officer at CBP. Information Lifecycle Management has become a common term within industry, especially given the enormous data growth and the need to classify data and put it in the right place.

Mahtab Emdadi with Dell Technologies echoed Madhugiri's remarks about innovations in cloud storage. Given the growing amounts of data collected by CBP and other federal entities, industry has recognized that traditional storage models are not feasible. Dell Technologies has a platform called Isilon that is designed to handle hundreds of petabytes of data and can help organizations better manage the enormous volume of unstructured data collected.

In speaking about the challenge of distributed data, Emdadi recognized it can be especially challenging for an organization like CBP where much of the mission is happening at edge locations such as the southern border. Dell Technologies has partnered with Intel to create local micro sites that connect to a larger central repository to get data to the edge. Emdadi also mentioned the importance of taking advantage of unique innovations happening within public clouds where integration between public cloud providers have allowed agencies to use the data they need, control it, and bring insights back to their agency. It takes an "ecosystem of partnerships" for industry to deliver the capability the government needs, she commented.

How Edge Computing Enables AI

Costello from CBP said for his organization "the concept of edge core cloud is making sure that you have harmony across all of those elements represented in your ecosystem. That is where CBP's focus resides. For the tactical edge it is about creating these compact platforms that will be able to bring that horsepower to the operators." He also mentioned the important connection between edge sites and central data warehouses so that operators can have access to the "peaks and highs of data utilizations needed for their algorithms."

Working with Clean Data Sets

Alboum with ServiceNow commented that when training systems and developing algorithms, it is important to know as much about that data as possible. He also spoke about the need for strong data governance and understanding how data can get manipulated. For example, certain data when combined with other data sets over time may generate an output that might not be as expected. "You might also get some outputs that are incompatible with the original data set because you have enhanced that data in ways that was not documented," said Alboum. DOE's Ingstad agreed and commented on the importance of data provenance and the need to control who has access to those data sets.

Alboum also commented that each data set is different and it is the responsibility of agency CIOs to understand the risk surrounding the data in a qualitative way—for example, knowing the impact of using the wrong data in an algorithm. Mercier added that concept of "model drift" when it comes to AI algorithms is another potential risk area. This happens when you have an AI model deployed for a very specific task, but as the environment changes, the model may need to evolve as well.

PART IV: **SECURITY CONSIDERATIONS**

Securing the Data

Jeff Alstott with the Intelligence Advanced Research Projects Agency, better known as IARPA, commented that it is not only about protecting data, but also the development process to make sure there are no “backdoors” into data sets or AI systems. “The whole supply chain running into your AI is very long, very distributed and reaches back to China and other places, especially when you consider the chips, algorithms, and the code bases being used,” he commented. Alstott works specifically on the Trojans in Artificial Intelligence (TrojAI) program within IARPA and much of their efforts center on identifying Trojan attacks and existing backdoors in AI systems in order to mitigate these problems.

Oracle has been able to offer security tools in the cloud, which has been a major innovation in terms of securing data. Paul Milsom with Oracle said their tools can detect attacks and repair damage at both the database-level and application-level.

Securing Hardware & Software

Gretchen Stewart with Intel discussed efforts to secure hardware and mentioned their partnership with the National Science Foundation and NIST to develop security standards around hardware and software. Intel continues to encrypt its hardware to the highest level whether that be memory, CPU, or even peripherals that are built by other companies, commented Stewart, who said they are approaching hardware security as if it is a zero trust environment. Once the hardware is secure, the software tools sit on top and need to be optimized for security—this combination of hardware and software security is critical to building trust around AI.

Rise of Adversarial AI & Consequences of AI Misbehaving

The conversation highlighted adversarial AI and the need to build security around data sets. “What we're really talking about here is the potential for an adversary to corrupt our data sets or to corrupt the models that we use. Adversarial AI is something that I think people haven't thought about a lot. They buy their data sets and they don't really know where they came from,” said Ingstad with DOE.

Major Bastian from the JAIC talked about the different paradigms around adversarial AI including:

- **White box:** the adversary has intimate knowledge of the architecture of your system;
- **Black box:** the adversary has no access to the internal details of your system, but can repeatedly probe it with different inputs and outputs to learn more;
- **Hidden box**: the adversary makes some assumptions about the model by observing overall system behavior.

Having this knowledge of types of adversarial AI demonstrates how detrimental a data poisoning attack can be to an AI model. Bastian mentioned an adversary polluting training data, which can skew the overall behavior of an AI system, as one example. He cited spam detection as a scenario where the adversary changes the algorithms to incorrectly classify spam emails as non-spam.

Alstott with IARPA, commented “the consequences of an AI misbehaving are purely defined by what you have the AI hooked to.” In the scenario with the spam detection system being compromised, someone might get an email that they shouldn’t have. An example of something much more nefarious would be AI misbehaving within a lethal autonomous weapon on a battlefield.

The DOD issued AI ethical principles that examines the scale of consequences around AI misbehaving and provides guidance on when a system needs to be shut down. “This would prevent us from getting into these situations where an AI has a heck of a lot riding on it. If we had a security or safety problem occur, that could be very bad. I’m thankful the Department of Defense and Intelligence Community has been thinking about this for a while,” commented Alstott.

PART V: WORKFORCE CONSIDERATIONS AND BEST PRACTICES FOR ADVANCING AI

Creating an AI Ready Workforce

The Department of Energy is bringing on a new leader to help create an AI ready workforce. Ingstad said that while training programs are important, there should also be an emphasis on AI career pathways. At DOE there are a number of nuclear physicists, chemists, and biologists who expressed they felt they are spending too much time managing data; AI can play an integral part of what labs will look like in the future. Ingstad also mentioned that Oak Ridge National Lab has a Manufacturing Demonstration Facility that could help create AI training programs, not just for federal agencies, but for the broader American workforce so the larger economy can benefit from AI applications across a number of industries.

Mercier with Microsoft emphasized that AI career pathways are varied and don't just require people with "hands on a keyboard". Rather it is a "full spectrum of skills" and is a field that continues to evolve rapidly requiring individuals with varied backgrounds and talents to help the U.S. embrace this technology. She also discussed the potential for people to apply "unconscious bias" or assumptions to an AI algorithm, which speaks to the importance of creating a diverse AI ready workforce so everyone can be held accountable.

Alboum with ServiceNow commented that AI and automation *does not necessarily replace jobs, it replaces tasks*. This should be taken into consideration when organizations create an AI strategy and it can be an opportunity to make existing workforces even more efficient and productive.

Stewart with Intel said that even as AI advances, there will always be a need for a human in the loop. "The truth is you really need the domain expertise. For the U.S. Border Patrol, you really need to understand what those folks are having to deal with in the field every day. The combination of great information and great data with a human is the panacea."

Advice from the Experts on Advancing AI

1. **When beginning a conversation around AI tools, start by discussing what an agency or commercial customer will get out of this**. Create something that can be validated and that will help the organization make better decisions. If it works, individuals need to champion that success story within their agency or organization and can open a discussion about other specific needs that can be addressed by AI. (Jonathan Alboum, ServiceNow)

2. **Knowing your audience and having the ability to translate success stories** to be relevant to a customer's particular domain is critical. (Amanda Mercier, Microsoft)

3. **Federal agencies should take advantage of the innovation happening in the public cloud**, that can help accelerate the use of AI. Avoid storing huge volumes of data in the public cloud,

especially because keeping those datasets there for a long period of time can make the cost untenable. “The trick is to take advantage of the innovation, but also maintain your own private data warehouse that resides with your own people.” (Mahtab Emdadi, Dell Technologies).

4. **Strike a balance between data availability and data storage**. Organizations need to ask how much data they really need to train AI models. Vendors are providing pretty sophisticated algorithms that the government should take advantage of rather than re-creating those efforts. (Sunil Madhugiri, U.S. Customs and Border Protection)

5. **Customers need to consider encryption across their entire data journey**. Where do you need the encryption and at what point can you decrypt? Cloud is an important factor in this decision and where the data resides, whether it is in a private cloud, public cloud, or a hybrid environment. (Gretchen Stewart, Intel)

6. **Do not repeat the security missteps made in the creation of the Internet with AI**. In the context of the Internet, IT infrastructure was built without much thought to security and we continue to deal with those consequences today. “It could have been the case that things like memory buffer overflow attacks were never a thing, but we made certain decisions decades ago about design [of the Internet] and now we are stuck with them. We are looking at a lot of great techniques for using AI and we are eager to put these into different contexts, but unfortunately some of them are fundamentally broken. We need to figure out how to make it more secure and safer. It is a race between making the system more secure and trying to use AI for mission.” (Jeff Alstott, IARPA)

PART VI: **CONCLUSION**

AI has the ability to transform the government mission, whether it is helping to respond to wildfires, staying ahead of adversaries on the battlefield, or reducing the burden of administrative tasks, AI adoption will accelerate in the years to come. The success of AI applications in government depends a great deal on data governance, using clean data sets and having an understanding of data sources, and leveraging cloud and storage innovations to deliver AI capabilities to the mission operators.

There are also remaining security concerns when it comes to AI systems. Adversaries are constantly looking for backdoors to manipulate data and influence outcomes. Creating a secure supply chain and ensuring that security is built-in from the beginning into the government's IT infrastructure—including hardware, software, and cloud—is crucial.

Advancing AI in government requires balancing security considerations, while advancing AI applications where appropriate. The U.S. Government has gotten much smarter in approaching AI with efforts such as the creation of the JAIC at DOD, appointing Chief Data Officers within federal agencies, and advancing bipartisan efforts in Congress around AI with input from academia. AI will continue to be a force for good, and bad on the part of U.S. adversaries. It will require continued innovation from the private sector and a whole-of-government approach to understanding the possibilities and limitations AI presents.

APPENDIX

Polling Responses

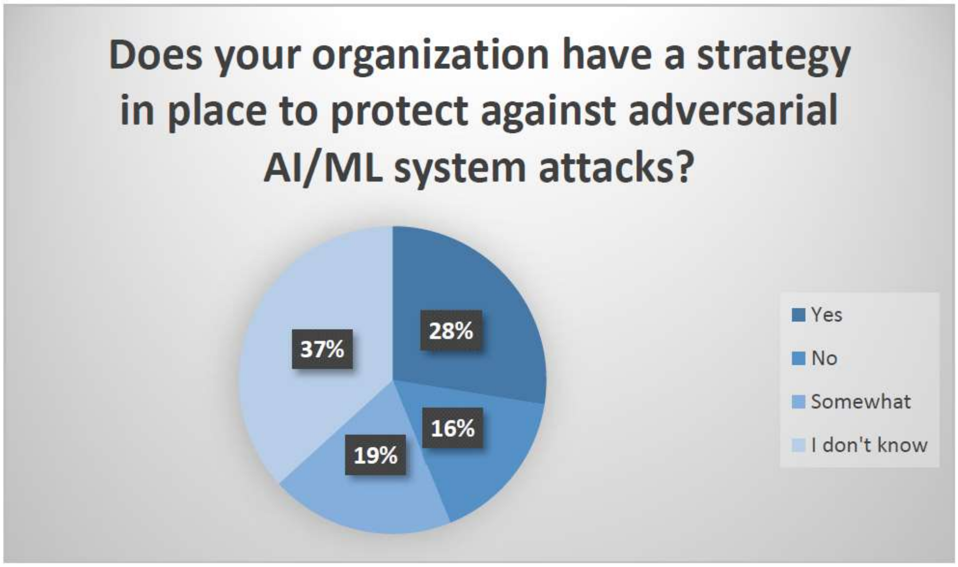

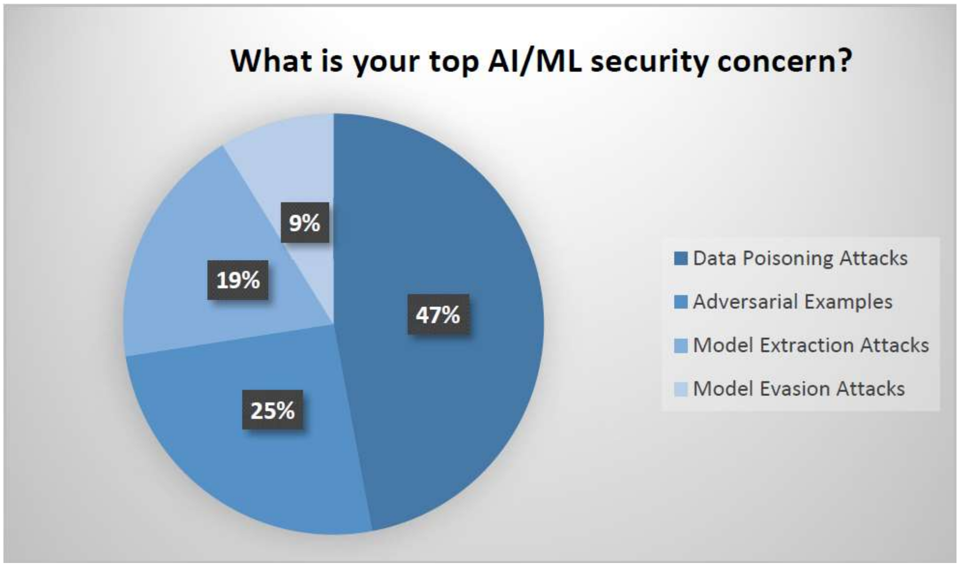

ABOUT THE EVENT ORGANIZERS

Center for Public Policy Innovation (CPPI)

The Center for Public Policy Innovation (CPPI) is a 501(c)(3) not for profit educational think tank whose mission is to assist government officials in addressing the many challenging issues brought on by the rapid advancement of Information Technology. CPPI provides policymakers with groundbreaking thought leadership on transformational technology, informed policy analysis, and innovative strategies to help ensure American competitiveness in the global economy and comprehensive security on the homefront. CPPI convenes educational symposiums, site visits, and other forums that bring together stakeholders from government, industry, academia, and the civic sector to discuss policy issues in a collaborative environment. For more information visit: cppionline.org.

Homeland Security Dialogue Forum (HSDF)

The Homeland Security Dialogue Forum (HSDF) was established in 2003 to boost the level of informed dialogue between the public and private sectors on homeland and national security issues and explore the important role of technology in hardening our nation's vulnerabilities. HSDF has organized more than 450 meetings and other special events with top officials from the Department of Homeland Security; other relevant federal, state, and local agencies; and foreign governments. HSDF is supported by a number of leading technology companies and security solutions providers. For more information visit: hsdf.org.

THANKS TO OUR EVENT PARTNERS:

DELL Technologies intel Microsoft Federal ORACLE servicenow.