S. Hrg. 118–135

GOVERNING AI THROUGH ACQUISITION AND PROCUREMENT

HEARING

BEFORE THE

COMMITTEE ON HOMELAND SECURITY AND GOVERNMENTAL AFFAIRS UNITED STATES SENATE

ONE HUNDRED EIGHTEENTH CONGRESS

FIRST SESSION

SEPTEMBER 14, 2023

Available via the World Wide Web: http://www.govinfo.gov

Printed for the use of the
Committee on Homeland Security and Governmental Affairs

U.S. GOVERNMENT PUBLISHING OFFICE
53–707 PDF WASHINGTON : 2024

COMMITTEE ON HOMELAND SECURITY AND GOVERNMENTAL AFFAIRS

GARY C. PETERS, Michigan, *Chairman*

THOMAS R. CARPER, Delaware	RAND PAUL, Kentucky
MAGGIE HASSAN, New Hampshire	RON JOHNSON, Wisconsin
KYRSTEN SINEMA, Arizona	JAMES LANKFORD, Oklahoma
JACKY ROSEN, Nevada	MITT ROMNEY, Utah
ALEX PADILLA, California	RICK SCOTT, Florida
JON OSSOFF, Georgia	JOSH HAWLEY, Missouri
RICHARD BLUMENTHAL, Connecticut	ROGER MARSHALL, Kansas

DAVID M. WEINBERG, *Staff Director*
LENA C. CHANG, *Director of Governmental Affairs*
MICHELLE M. BENECKE, *Senior Counsel*
EVAN E. FREEMAN, *Counsel*
LIANA S. KEESING, *Research Assistant*
WILLIAM E. HENDERSON III, *Minority Staff Director*
CHRISTINA N. SALAZAR, *Minority Chief Counsel*
KENDAL B. TIGNER, *Minority Professional Staff Member*
LAURA W. KILBRIDE, *Chief Clerk*
ASHLEY A. GONZALEZ, *Hearing Clerk*

CONTENTS

WITNESSES

THURSDAY, SEPTEMBER 14, 2023

ALPHABETICAL LIST OF WITNESSES

APPENDIX

GOVERNING AI THROUGH ACQUISITION AND PROCUREMENT

THURSDAY, SEPTEMBER 14, 2023

U.S. SENATE,
COMMITTEE ON HOMELAND SECURITY
AND GOVERNMENTAL AFFAIRS,
Washington, DC.

The Committee met, pursuant to notice, at 10 a.m., in room 562, Dirksen Senate Office Building, Hon. Gary Peters, Chair of the Committee, presiding.

Present: Senators Peters [presiding], Carper, Hassan, Sinema, Rosen, Ossoff, Blumenthal, Paul, Lankford, Scott, Hawley, and Marshall.

OPENING STATEMENT OF SENATOR PETERS[1]

Chairman PETERS. The Committee will come to order. Today's hearing is the third in a series that I have convened on artificial intelligence (AI). At our first meeting in March, we discussed the transformative potential of AI, as well as the possible risk that these technologies may pose.

At our second hearing in May, we considered the role of AI in government. How AI tools can improve the delivery of services to the American people, and how to ensure they are being used both responsibly and effectively.

Today, we are going to do a deeper dive into how government will purchase AI technologies, and how the standards and guardrails that government sets for these tools will shape their development and use really all across all industries.

The Federal Government is already using AI and its use across agencies is only expected to grow in the coming years. These systems could help provide more efficient services, assess potential security threats, and automate routine tasks to enhance the Federal workforce.

For example, the Department of Homeland Security (DHS) is using natural language processing to evaluate employee surveys and improve workplace experience. The Federal Aviation Administration (FAA) deploys machine learning to update the weather models that help land planes successfully. Other technologies that are continuing to develop, such as generative AI, offer the potential to improve government services even more.

For example, many agencies, from the Office of Personnel Management (OPM) to the Department of Health and Human Services

[1] The prepared statement of Senator Peters appears in the Appendix on page 41.

(HHS), to the Department of Education, have rolled out chat bots to provide better service to Federal employees and the larger American public. AI is here and is already being put to good use.

Many of these systems are not developed by the government, but rather the private sector. Over half of the AI tools used by Federal agencies have been purchased from commercial vendors. This collaboration between the public and the private sector is crucial. It ensures that government is using the most effective AI systems.

American companies are breaking new ground with these technologies, and we have a chance to share in the benefits of that incredible innovation. But these tools also bring potential risk and policy implications. They require new knowledge from procurement officials, as well as increased coordination across agencies.

In order to successfully and effectively purchase and use AI tools, Federal agencies have to be prepared to address issues like privacy concerns about the use of Federal data to train commercial models and bias in government decisionmaking. We must be nimble, whenever the government collaborates with a private sector, but this is especially true with AI, where new developments emerge almost every single day.

The tools that are purchased are often actively learning and are changing as they are used. Last Congress, I authored and enacted a law that requires officials that procure AI tools, be trained in their capabilities as well as their potential risk. This year, I introduced legislation that would extend this training to all Federal managers and supervisors.

I have also introduced legislation that would designate a chief AI officer at every Federal agency so that they have leadership and expertise to maximize the potential of these technologies and effectively address risk. These guardrails are more important than ever. Federal agencies are inundated with sales pitches and technology demos, promising the next big thing.

While the Federal Government must be forward thinking, we also have to be cautious in procuring these new tools. We must continue to work past the initial purchase, testing and fine tuning our models to ensure that they are effectively serving the American people. As AI development accelerates, private industry has yet to standardize practices for evaluating AI systems for risk, for trustworthiness and responsibility.

Through the Federal procurement policy, the government has a unique opportunity to shape these standards and frameworks for development and deployment of these technologies across the private sector more broadly. I look forward to hearing from our expert witnesses here today.

We look forward to working with you not just today, but in the weeks, months, and years ahead, and continue our bipartisan work to help encourage American development of AI and ensure that it is being used appropriately. I would now like to turn the microphone over to Ranking Member, Senator Paul, for his opening statement.

OPENING STATEMENT OF SENATOR PAUL[1]

Senator PAUL. Thank you. In 2021, the Pentagon, through the Defense Advanced Research Projects Agency (DARPA), asked for proposals for real time comprehensive tools that established ground truth for how countries are conducting domestic information control.

DARPA's goal in developing AI technology for measuring the information control environment was to help the U.S. Government better understand how digitally authoritarian regimes repress their populations at scale over the internet via censorship, blocking, or throttling.

Of course, the solicitation made it clear that the Pentagon did not want the proposals to look at activities of the U.S. Government. The Pentagon and the U.S. Government as a whole enjoy professing moral superiority over authoritarian governments when it comes to upholding basic democratic values.

American politicians have no qualms about criticizing foreign governments like Russia and China for their suppression of civil liberties and efforts to eliminate dissent. Yet there seems to be a complete unwillingness to have an honest conversation about the disturbingly similar actions our own government is actively engaged in and financing.

For decades, the Pentagon and other Federal agencies have been quietly partnering with private organizations to develop powerful surveillance and intervention tools designed to monitor and influence narratives on social media.

For example, a 2021 Pentagon program called Civil Sanctuary sought to use artificial intelligence tools to scale the moderation capability of social media platforms to create what it describes as a more stable information environment. In other words, the goal of this Pentagon program was to exponentially multiply the government's ability to coordinate censorship of online speech.

The Pentagon has invested millions of dollars to develop these tools, not only for use by social media companies, but also the intelligence community (IC) and law enforcement. Meanwhile, the Department of Commerce (DOC) is awarding million dollar grants for cognitive research into how the U.S. Government can foster trust in artificial intelligence with the general public.

While the Federal Government is using taxpayer dollars to develop AI to surveil and monitor Americans' online speech, it is also spending money to figure out how to get you to trust the government with AI.

Over the last year, starting with the Twitter Files, journalists started to expose the deep coordination between the Federal Government and social media. When it comes to content moderation, these decisions in policing the speech of Americans, we have seen this enormous connection between Government and private entities.

As Michael Shellenberger rightly points out, the threat to our civil liberties comes not from AI, but from the people who want to control it and use it to censor information. It is not the tool, it is the corruption of power that is always the problem.

---

[1] The prepared statement of Senator Paul appears in the Appendix on page 44.

Last week, the Fifth Circuit affirmed the government likely violated the First Amendment, a big deal, violated the First Amendment by coercing social media companies to remove speech that the government disagreed with related to the origins of Coronavirus Disease 2019 (COVID–19), the pandemic lockdowns, vaccine efficacy, and the Hunter Biden laptop story.

The court cited numerous examples of U.S. Government officials engaging in domestic information control on social media. Our concern is not that they are doing it, but they are going to do it even more efficiently and even more ruthlessly if they get artificial intelligence and are able to comb through the entire internet.

They are already doing this. Our concern with artificial intelligence is they will take that tool and much more efficiently go through millions and millions opposed to say, that is not allowed. Government officials demanded that the platforms implement stronger COVID misinformation monitoring programs, and then they threatened the platforms. They threatened them with taking away Section 230.

They threatened them with antitrust action. It is amazing. This was no sort of please take down some information. It is, take it down or else. That is why the court enjoined and said to the Biden Administration, you must stop. Currently, the Biden Administration is not meeting with them.

They have had to cancel their meetings with the Federal Bureau of Investigation (FBI), with the Department of Homeland Security. But this is a big deal. But it seems to be only one side of the aisle has been concerned at all with what has happened because some of it involves politics. But it should not. I mean, free speech should be something that both parties really are concerned with trying to protect.

After one meeting with Federal officials, one platform, social media platform, committed to reducing the visibility of information that was skeptical of the government's COVID vaccine policy, even when it did not contain any misinformation.

They were saying, even if it is true, we want you to take it down, because some people might not get vaccinated because you said something that actually did occur, but we do not want people to know that because it would lessen people's enthusiasm.

That is a crazy notion. Facebook promised to label and demote a popular video after officials flagged it, even though they acknowledged it did not qualify for removal under its policies. I fear that we are likely only in the beginning stages of understanding the extent of the Federal Government's involvement in content moderation and the decisions that private social media platforms make.

What we do know is that our government is funding the development of powerful artificial intelligence tools for monitoring and shaping online discourse. I want to be clear, AI is not inherently malicious. It has the potential to revolutionize basic aspects of society, from health care to education.

However, in the hands of unchecked government, AI can be weaponized as a tool to suppress fundamental values like speech, things that our country was founded upon—the open exchange of ideas, the freedom to question, and the right to dissent. This should not be a partisan issue.

Chairman PETERS. It is the practice of the Homeland Security and Governmental Affairs Committee (HSGAC) to swear in witnesses. If each of you would please rise and raise your right hand. Do you swear that the testimony that you will give before this Committee will be the truth, the whole truth, and nothing but the truth, so help you, God?

Mr. GHANI. I do.

Ms. LI. I do.

Ms. RAJ. I do.

Mr. ROBERTS. I do.

Mr. SHELLENBERGER. I do.

Chairman PETERS. Thank you. You may be seated. Our first witness is Professor Rayid Ghani. Professor Ghani is a Distinguished Career Professor in the Machine Learning Department at the Heinz College of Information Systems and Public Policy at Carnegie Mellon University (CMU). At CMU, he leads the Data Science and Public Policy Group, and the Data Science for Social Good program, as well as the Responsible AI Initiative, which he co-leads. Professor Ghani, it is great to have you here at the Committee. You are recognized for your opening statement.

TESTIMONY OF RAYID GHANI,[1] DISTINGUISHED CAREER PROFESSOR, MACHINE LEARNING DEPARTMENT AND THE HEINZ COLLEGE OF INFORMATION SYSTEMS AND PUBLIC POLICY, CARNEGIE MELLON UNIVERSITY

Mr. GHANI. Thank you. Thank you, Chair Peters, Ranking Member Paul, and other Members of the Committee. Thanks for hosting this hearing today and for giving me the opportunity to present this testimony.

As Chair Peters mentioned, my name is Rayid Ghani. I am a Professor of Machine Learning and Public Policy at Carnegie Mellon. I am here today because I believe that AI has enormous potential in helping us tackle critical societal problems that our governments are focused on.

Much of the work I have done over the last decade has been in this space through working extensively with governments at the Federal, State, and local level, including helping and using AI systems to tackle problem across health, criminal justice, education, public safety, human services, workforce development, particularly on supporting fair and equitable outcomes.

Based on my experience, I believe that I can benefit every Federal, State and local agency. However, any AI system, or any other type of system affecting people's lives, has to be explicitly designed to promote our societal values, such as equity, and not just narrowly optimized for efficiency.

I think it is critical for us, Government agencies, policymakers, to ensure that these systems are designed in a way that they do result in promoting our values. Now, while the entire lifecycle of AI systems, from scoping, to procurement, to designing, to testing, to deploying needs to have those guidelines in place that maximizes societal benefits and minimize potential harms, there has been a

[1] The prepared statement of Mr. Ghani appears in the Appendix on page 46.

lack of attention to the earlier phases of this process, particularly on the problem scoping and procurement parts.

As Chairman Peters mentioned, many of the AI systems being used in government are not built in-house. They are procured through vendors, consultants, and researchers. That makes getting the procurement phase correct critical. Many problems and harms discovered downstream can be avoided by a more effective procurement process.

We need to make sure that the government procurement of AI follows a responsible process and in turn makes the AI vendors accountable for the systems they design. They themselves have to promote accountability, transparency, and fairness.

Government agencies often go on the market to buy AI without understanding and defining and scoping the problem they want to tackle, without assessing whether AI is even the right tool, and without including individuals and communities that will be affected. AI systems are not one size fits all.

Procuring AI is first and foremost procuring a solution that is helping solve a problem and should be assessed in the ability to better solve the problem at hand. In that respect, procuring AI is not that different from procuring other technologies. Now, there are a few areas where it is different.

One, AI algorithms are neither inherently biased or unbiased, nor have inherent fixed values. The design of these systems requires making hundreds and sometimes thousands of choices that determine the behavior of the system. If these choices explicitly focus on outcomes we care about and we evaluate the systems against those intended outcomes, the AI system can help us achieve what we want to achieve.

Unfortunately, today, those decisions are too often left typically to the AI system developer who defines those values implicitly or explicitly. The procurement process needs to define these goals and values very explicitly. AI requires that. Society requires that. Ensure that the vendors address those appropriately in the system being procured and provide evidence of that.

Building responsible AI systems requires a structured approach, and the procurement process needs to set expectations, enforce transparency and accountability from vendors in each of these steps.

That includes defining goals, translating them into requirements that the vendors to design the system to achieve, and setting up a continuous monitoring and evaluation process, because the system will both itself change, as well as have to live and function in an ever changing world.

It is critical and urgent for policymakers to act and provide guidelines and regulations for procuring, developing, and using AI, in order to ensure that these systems are built in a transparent and accountable manner, and result in fair and equitable outcomes for our society.

As initial steps, here are some of my recommendations. No. 1, focusing the AI procurement process on specific use cases rather than general purpose, one size fits all AI, both to support intended outcomes around that use case, as well as to prevent harm through misuse.

No. 2, development of common procurement requirements for AI and templates the government agencies can start from. That does not exist today. No. 3, create guidelines that ensure meaningful involvement of the communities that will be impacted by the AI system, right from the inception stage and continuously.

No. 4, and last, creating trainings, processes, and tools to support the procurement teams within government agencies. As the government teams expand their role and start procuring AI augment the systems more regularly, they will need to be supported by increasing their capacity to fulfill this role.

I recommend creating a set of trainings and processes and collaboration mechanisms and tools to help them achieve that. The overall goal behind these recommendations is to set some standards around procurement of AI by government agencies, and to support and enable the agencies to implement those standards effectively and procure AI systems that can help us achieve their policy and societal goals. Thank you.

Chairman PETERS [continuing]. Dr. Li is the Sequoia Professor of Computer Science at Stanford University, and Co-Director of the Stanford Institute for Human Centered Artificial Intelligence (HAI). Before that, Dr. Li spent five years as the Director of Stanford's AI Lab. During that time, she was also Vice President at Google and the Chief Scientist of AI–ML for Google Cloud.

Dr. Li is the inventor of ImageNet and the ImageNet Challenge, a large scale data set effort that contributed to significant advances in deep learning and computer vision. Dr. Li, thank you for being here today. We look forward to your opening comments.

TESTIMONY OF FEI–FEI LI, PH.D.,[1] SEQUOIA PROFESSOR, COMPUTER SCIENCE DEPARTMENT AND CO–DIRECTOR, HUMAN–CENTERED AI INSTITUTE, STANFORD UNIVERSITY

Dr. LI. Thank you. Thank you, Chair Peters, Ranking Member Paul, Members of the Committee. Thank you for the privilege of appearing before this prestigious body. It is truly an honor.

I have spent my life working in the field of artificial intelligence, over 20 plus years, studying, developing, and understanding the technology that has entered the public consciousness due to recent breakthroughs. There is no doubt that we have arrived at an inflection point in AI, largely powered by generative AI, including large language models (LLM) and my own field of computer vision, where we essentially teach computers to see.

For example, a notable application is health care, where AI is augmenting the capabilities of caretakers and medical professionals by detecting anomalies in medical imagery such as X-rays or MRI scans, thereby aiding early diagnosis and treatment. Most importantly, AI presents many opportunities to the U.S. Government.

One area is to streamline the efficiency of government. Many Federal agencies are already experimenting with AI powered tools. For example, HHS has initiated a pilot program that employs AI to enhance the efficiency of fraud detection within the Centers for Medicare and Medicaid Services (CMS).

[1] The prepared statement of Dr. Li appears in the Appendix on page 53.

Second, AI in health care reduces the burden on public health care resources, including Medicare and Medicaid. Medical AI tools can decrease the frequency of unneeded emergency medical interventions and hospital readmissions, which are significant cost drivers in health care expenditures.

However, while AI, like most technologies, promises to solve many problems for the common good, it can also be misused to cause harm and carry unintended consequences. Let me just give you two examples of when the harms of AI can affect how the government approaches it.

First, bias in AI is well-documented. For example, in credit risk scoring, research shows that predictive tools used to approve or reject loans are less accurate for low income minority groups in the United States due to the lack of data in their credit histories. To ensure that AI applications deliver reliable results for all Americans, the availability of high quality representative datasets is crucial.

Second, in an area of heightened public concern over data collection and misuse, we must integrate strong privacy and security protocols into these applications from the beginning. To build Privacy-preserving technology, we must engage diverse stakeholders in health care. This includes AI developers, public sector, health care professionals, patients, and more.

This is why I founded Stanford Institute for Human Centered Artificial Intelligence, where we study AI and its impact not as a field exclusive to computer science, but instead as a multidisciplinary field that includes the social sciences, engineering, law, medicine, humanities, and more.

The Federal Government should adopt a similar approach to properly understood the future of AI. It falls upon the U.S. Government to spearhead the ethical procurement and deployment of these systems, setting norms for AI development and ultimately shaping the field of responsible AI.

I applaud this Committee and the work that has been done thus far on AI, including the AI Training Act and the AI Lead Act, which create powerful tools for Federal Government to set such norms. In the United States, government spending on AI related contracts—as the U.S. Government's funding has surged in AI related contracts, it is more crucial than ever to closely examine these vendors to ensure their goals align with those of the Federal Government. One key component is evaluation, especially in key areas like health care, education, agriculture, finance, and more.

Having created one of the most consequential evaluation datasets for AI models, ImageNet, I firmly believe that evaluation should consider every factor in a holistic way, from accuracy, to fairness, to the reliability of models performing under real world conditions. Second, we must build-in transparency measures.

Vendors should disclose key information about their systems, including how they collect and annotate datasets, what potential risks their systems pose, and how they mitigate those risks. But the procurement is just one piece of the puzzle.

For the United States to maintain its leadership in AI, the Federal Government must make the needed critical public investments in AI. Due to the vast amount of compute and data required to

train these systems, only a select few industry players can shape the future of AI, leaving an imbalance in the innovation ecosystem that lacks the diverse voices of public sector and government labs.

The lack of public sector investment in AI hampers not only in thoughtful regulation, but also proper Federal procurement. Without the ability to train AI talents, the Federal Government will not have the necessary human capital to create meaningful regulation, ensure ethical AI procurement, and be the true AI leader it has the potential to be.

This is why I am unequivocally a strong supporter of the Create AI Act, a strong bipartisan legislation introduced this summer in both chambers. The Create AI Act will establish a national AI research resource that provides the needed resources to allow public sector researchers to innovate and train the next generation of AI leaders.

In June, I personally shared with President Biden how I believe the United States is not prepared for this imminent AI moment. What we need right now is a coordinated moonshot effort for the Nation to ensure America's leadership in AI for the good of humanity.

This task will be no small feat, but with meticulous coordination, significant investment in scientific AI research, and robust collaboration across government, public sector, and industry, we can rise to meet this challenge and ensure America's leadership in AI is both impactful and enduring.

Thank you to the Chairman, Ranking Member, and all the Members of the Committee for allowing me to testify today.

Chairman PETERS. Thank you. Our next witness is Devaki Raj. Ms. Raj is the former Chief Executive Officer (CEO) and co-founder of CrowdAI, a computer vision startup that has been contracting with the U.S. Government since 2019. Ms. Raj was recognized in 2019 on the Forbes 30, under 30, as an AI leader to watch.

Previously, she worked for Google as a Data Scientist in their maps and android sector. Ms. Raj, thank you for being here today, and we look forward to your testimony.

TESTIMONY OF DEVAKI RAJ,[1] FORMER CHIEF EXECUTIVE OFFICER AND CO-FOUNDER, CROWDAI

Ms. RAJ. Chairman Peters, Ranking Member Paul, and distinguished Members of the Committee, thank you for this opportunity to testify on governing AI through acquisition and procurement from the perspective of a small business.

My name is Devaki Raj. I am honored to be here representing CrowdAI, a startup developing no-code artificial intelligence tools since 2016. Until a recent acquisition by Saab, I served as CrowdAI's CEO and Co-Founder.

I admire Chair Peters and Ranking Member Paul respectively for their leadership on AI initiatives and on improving Small Business Innovation Research (SBIR) procurement. Today, there is an emphasis on procurement of commercial technology. However, AI needs to be procured in a manner that reflects this novel technology.

[1] The prepared statement of Ms. Raj appears in the Appendix on page 58.

First commercial off the shelf AI solutions need government curated data to be mission ready. Second, AI procurement needs to include ongoing AI retraining to support it. Third, commercial AI technologies must undergo rigorous testing and evaluation. Finally, it is important to establish paths to programs of record for small businesses through project transition milestones.

First, for government missions be it public health or homeland security, commercial AI does not just transfer out of the box. While the tools to create, modify, and operate AI are available commercially, the algorithms themselves are trained on commercially available datasets.

For example, imagine an AI algorithm built for self-driving cars used to analyze nighttime drone video during a maritime search and rescue mission. Yes, both models were trained to identify vehicles, but their domains, the operational context, and sensors are different.

Robust AI must learn from specific domain mission data it is applied to. If not, models remain brittle. It is important to note that there are constraints that come with the use of government data, statutory limitations, privacy, access, security, etcetera. Appreciating this, offices desiring to implement AI should curate their domain specific data to accelerate development, testing, and transition to operations.

Second, AI is a journey not a destination. AI procurement must include ongoing AI model retraining for continued operational relevance. In 2018, CrowdAI collaborated with the California Air National Guard to automate wildfire mapping using MQ–9 drones, a collaboration we proudly continue.

Our predictive models performed extremely well in Northern California's forested regions where wildfires were common, and we had access to government furnished data. However, as evolving wildfire epicenters shifted to urban areas in the South, our models required retraining to maintain operational relevance.

While AI models are flexible, they still require contracting officers to include model training and contracts, ensuring alignment with evolving mission data. It is a dynamic process akin to software updates, not just a one-time procurement. Third, open publicly available AI code requires rigorous testing and evaluation. Today, anyone can go online and download an open source AI model.

This poses a challenge for government selection panels because they have little expertise of verifying commercial claims. For example, we developed an AI model to identify remote airstrips often used for drug trafficking in South America. From a statistical standpoint, our model performed expertly, finding 100 percent of the airstrips.

But from a mission perspective, it generated excessive false positives by also identifying significant amounts of dirt roads, which was fundamentally detrimental to intelligence analysis. I share these negative results with you to show that evaluating AI is not simple.

It is therefore critical that procurement activities include both qualitative and quantitative evaluation for evaluation metrics for AI throughout both the solicitation process and post-delivery

phases. Finally, small business AI procurement should have project transition milestones.

While SBIR's phased approach offers a structured path for validation to transition, the increasing number of awardees contrasted with fewer transitions suggests misalignment of incentives. I believe that is crucial for any government procurement to have clear transition milestones for a path to a program of record.

The Naval Air Warfare Center's record of including transition milestones in SBIR awards is a shining example. In conclusion, the needs and resources of government missions are unique, requiring tailored AI solutions.

Procurement vehicles must be modernized to reflect the iterative nature of AI for government missions and introduce stringent standards for testing and evaluation. Thank you for your time and I look forward to answering questions.

Chairman PETERS. Thank you. Our next witness is William Roberts. Mr. Roberts is the Director of Emerging technologies for ASI Government. He previously served as the Head of Acquisitions for the Joint Artificial Intelligence Center (JAIC), and Chief Digital and AI Officer (CDAO) for the Department of Defense (DOD).

Before that, Mr. Roberts was the contracting officer for the Office of the Secretary of the Air Force and worked on acquisition policy for the Department of Defense Education Activity (DODEA). Mr. Roberts, welcome to the Committee. You may proceed with your opening remarks.

TESTIMONY OF WILLIAM ROBERTS,[1] DIRECTOR OF EMERGING TECHNOLOGIES, ASI GOVERNMENT

Mr. ROBERTS. Thank you, Chairman Peters, Ranking Member Paul, and distinguished Members of the panel. I am Will Roberts. As Chairman Peters mentioned, I am the Director of Emerging Technologies for ASI Government.

Previously, I was the Director of Acquisition for the Department of Defense Joint AI Center (JAIC), and I have a particular passion for government contracting. During my time at the JAIC, I became very aware of the procurement related challenges to buying and delivering AI to the end user for adoption, for true, real adoption.

I also developed a genuine belief that the acquisition professional, and specifically the contracting officer, holds a very historically important role for us in this space, but they have to be trained so they can step up into this historic role. Currently, they are not, and this is the basis of my testimony.

Considering the historic role of the contracting officer, you could look at U.S. history as a string of transactions. It is a series of deals that led us and brought us to the Nation that we are today. That is because since the Revolutionary War, before we became a nation, we have always relied on industry to deliver.

The government has always connected the services and the innovations to the mission for the benefit of the citizens, but it is the ingenuity of industry that has brought us the innovations, the stuff, the airplanes, the ships, the tanks, the things that brought us to the moon in the 1960s.

[1] The prepared statement of Mr. Roberts appears in the Appendix on page 66.

This is all spelled out in the four corners of a string of business transactions between the government and industry created by the American Government dealmaker. Now we are facing a new chapter in our history with technology we have never seen before, and our American Government dealmaker is called to action once again.

The question that the government must ask itself is not how to develop it. If we are to follow the historic path that led us to be the successful nation we are today, and the question that the government must ask itself is how to buy it.

This question of how to buy it was my life for the last four years. It is complicated. It requires people who take their jobs very seriously. It led myself and my team and others in this space to rethink the way we do procurement within the bounds of the law. It required a special talent and one that had to be learned and developed.

First and foremost, the diligence AI acquisition official must realize that AI is a means to the end, and the end is always the mission. We are never really buying AI. We are buying an enhancements to our mission. Until acquisition professionals realize that they will never really deliver anything truly valuable to the end users.

But within the four corners of the contract vehicle, the parties negotiate very important intellectual property terms, which require knowledge of the various components of AI, the data rights, the cloud, the platforms, the infrastructure, the trained and untrained model, all of which could have a separate intellectual property strategy, which could make or break the project.

Within the four corners of the contract, the parties decide upon the perimeters of the responsible and safe use of the AI and its associated risk mitigation. Within the four corners of the contract, the parties agree upon how this will actually be delivered, really delivered and adopted—not just talk.

Through iterative methodologies, through flexibility, and being able to pull out if the project is not working to prevent wasting taxpayer dollars. Within the four corners of the contract rests the fate of our success in truly delivering AI into the government. The American Government dealmaker and their role becomes very important.

I will close with my two recommendations. One is to provide more contract authorities for contracting officers across the entire government. We had the privilege of using a lot of various authorities, but they were only available to the Department of Defense. These are tools that should be used by a skilled contracting officer in her tool belt. The skilled as is the emphasis there because the tools are useless if you do not know how to use them.

My second and more important point is there needs to be a robust training program that completely re-skills the acquisition workforce. The AI Training Act was a great step forward, but more needs to be done. There needs to be three core competencies. One on AI technical knowledge, which the Act covers, the functionalities and the risks. But the second two actually lead to true delivery.

One is AI business acumen, knowledge of the unique AI market. Many of the members in the market which have never worked with the government before. The third is AI unique contract domain

knowledge, to include internet protocol (IP), ethics, agile contracting, and use of these various contract authorities.

I think it is only then that we would really realize the benefits and the savings of the costs and savings to the lives, and the benefits to the welfare and defense of our Nation. Thank you.

Chairman PETERS. Thank you. Our final witness will be introduced by Ranking Member Paul.

Senator PAUL. We are pleased today to have Michael Shellenberger here with us. He has become one of the most prominent contemporary writers on the scene with regard to censorship. He has been a long time writer on the environment.

But has come to prominence to a lot of people's attention because of his being chosen by Elon Musk to look at the Twitter Files and to see firsthand the interaction between government and a large social media company, and how censorship was brought about by the government pushing and forcing a social media company to adhere to a government interpretation of policy events.

We are excited to have him. Michael graduated with a Bachelor of Arts (B.A.) in Peace and Global Studies from Earlham College, has a Masters of Arts (M.A.) in Cultural Anthropology from the University of California Santa Cruz, is the Founder and President of Environmental Progress from Berkeley, California. And Michael, we are happy to have you today. Thanks for coming.

TESTIMONY OF MICHAEL SHELLENBERGER,[1] FOUNDER, PUBLIC

Mr. SHELLENBERGER. Thank you very much. Chairman Peters, Ranking Member Paul, and Committee Members, thank you for your stated concern with the implications of AI for our civil liberties and our Constitutional rights, and for requesting my testimony. I am honored to provide it.

The ability to create deepfakes and fake news through the use of AI is a major threat to democracy, say many experts. The Washington Post recently reported that AI generated images and videos have triggered a panic among researchers, politicians, and even some tech workers who warn that fabricated photos and videos could mislead voters in what a United Nations AI adviser called in one interview, the deep fake election.

Never before in the United States have we been better prepared to detect deepfakes and fake news than we are today. In truth, the U.S. Department of Defense has been developing such tools both for the creation and the detection of deep fakes for decades. Before elaborating on this point, I want to emphasize that I view AI as a human not machine problem, as well as a dual use technology with the potential for good and bad.

My attitude toward AI is the same fundamentally as it is toward other powerful tools we have developed, from nuclear energy to biomedical research. With such powerful tools, democratic civilian control and transparent use of these technologies would allow for their safe use, while secretive, undemocratic, and military control increases the danger.

---

[1] The prepared statement of Mr. Shellenberger appears in the Appendix on page 75.

The problem in a nutshell is not with the technology of computers attempting to emulate human thinking through algorithms, but rather, who will control it and how they will do so. There is a widespread belief that users already choose their content on social media platforms.

In reality, social media platforms decide a significant portion of what users see. The heavy lifting of censorship, or what we call content moderation, was by 2021 already overwhelmingly determined by AI. Mark Zuckerberg, the CEO of Meta, said that more than 95 percent of the hate speech that Facebook took down is done by AI, not by a person, "98 or 99 percent of terrorist content that we take down is identified by AI."

Similarly, 99 percent of Twitter's content takedowns started with machine learning. The problem with AI technology today, funded by the U.S. Government, whether DARPA or National Science Foundation (NSF), is fundamentally around the control of these technologies by small groups of individuals and institutions unaccountable to the citizens of the United States.

The censorship industrial complex of government agencies and government contractors has its roots in the war on terrorism and the expansion of surveillance after 9/11. In 2003, DARPA told Congress that National Security Agency (NSA) was its experimental partner, using total information awareness and AI to detect false information.

In 2013, the New York Times reported on the NSA's use of AI, which foreshadowed how counter disinformation experts would nearly a decade later describe fighting misinformation online. In 2015, DARPA launched the funding track that directly resulted in AI tools that leading internet and social media companies used today.

Their goal was to develop a science and practice for determining the authenticity and establishing the integrity of visual media. DARPA's warning eight years ago is identical to the Washington Post warning about deepfakes last month. The adoption of AI has grown alongside alarmism about deepfakes, and misinformation and disinformation more broadly.

Also in 2019, a new non-governmental organization (NGO) called the Deep Trust Alliance launched a series of events called The Fix Fake Symposia. The Deep Trust Alliance described itself as, "the ecosystem to tackle disinformation" and its website invited audiences to join the global network, actively driving policy and technology to confront the threat of malicious deepfakes.

Yet, the goal of this Deep Trust Alliance appears to be to advocate for policies to censor and even criminalize digital harms. The head of the organization said that laws needed to be extended to digital harms. There needs to be a set of practices across social media platforms.

It was during this period that the U.S. Government, DHS, created Election Integrity Partnership (EIP) to censor elections skepticism. The year afterwards, it created a project, The Virality Project, to censor COVID skepticism and COVID criticism, or most famously, the Biden White House demanded widespread censorship by Facebook of what Facebook itself called, often true documentation of vaccine side effects.

While social media platforms use AI to identify and censor content, the decisions of what to censor, of course, remain in the hands of humans. The Federal Trade Commission (FTC) in June of last year warned Congress specifically about the dangers of using AI for censorship, urging great caution.

Good intentions were not enough, said the FTC, because it turns out that even well-intended AI uses can have the same problems like bias, discrimination, and censorship, often discussed in connection with the uses of AI.

My recommendations to the Congress, rather, are that we have much stricter oversight of these programs, making sure that we have greater understanding and greater control of how these censorship technologies are used. They should be in the hands of users, not in the hands of big platforms working with big government. Thank you very much.

Chairman PETERS. Thank you. Mr. Roberts, you spoke about the responsible use of AI in your opening comments. My question for you is, from your experience at the Joint AI Center, would you tell the Committee more specifically about what responsible risk management actually looks like at each stage of the procurement process?

Mr. ROBERTS. Yes, Senator. I will mention that it was an evolving process when I started at JAIC. We started with the two levels of trustworthiness, which was the trustworthiness and the functionality itself, whether it will work or created distrust.

But the trustworthiness and the safety and responsible use became a big theme for us. It started in the very planning phase. This is the first instance that we noticed we had to have a more balanced team.

It was not like a relay race where the money people, sent it to the contracting, and the program manager sent it over, and the testers and evaluators were at the very end. We needed to have the input from testing and evaluation from the ethics policy professionals and the end users at the very planning stage, so that the team worked more like a football team where we all ran the ball down the field at the same time.

This was a cultural change for us. This is not common in acquisition. We had to adjust accordingly for that. The two areas where we can provide the most meaningful value in the procurement process is in the evaluation phase and the testing phase. Evaluation, where we can make the responsible use of AI a discriminator for source selection, for award. Then the testing, of course, we can make as a metric for successful performance.

This was challenging for us though, because we did not want to be too restrictive at the time. We wanted to bring in players and companies. We started our efforts by turning it around to industry and asking them.

We had five principles for responsible use of AI the Secretary of Defense released, and we asked them to give us a quality control plan of how they were going to live out those principles. For the most part, it became this ongoing dialog, since we were all in uncharted territory, about how to do things responsibly.

But I will also mention that aside from some specific projects, especially those that dealt with health and warfighting, so much of

our projects had low risk in terms of responsible use. We were able to talk with our ethics person and they were able to weigh the risks.

But for those that were high risk, we treaded carefully, and we worked hand in hand with industry, specially making sure that when the contract was created, the post award performance, which is where it really matters, was all set up to make sure that as this AI was being delivered, it was being done in a responsible way.

Chairman PETERS. Thank you. Ms. Raj, right now, we are hearing from our Federal procurement officers that they are basically being bombarded by companies wanting to demonstrate the promise of their products, and they have no shortage of promises that they are presenting.

My question for you is, how can we ensure that the Federal agencies avoid being caught up in what is clearly AI hype right now around the country and spend taxpayer money responsibly on services that will actually deliver tangible results for the American people?

Ms. RAJ. Thank you, Senator, for the question. As a former small business owner, we have to continuously rise against the same noise to provide the best quality computer vision for our customers.

We first started working with the U.S. Government in 2017, and in 2017 we were one of the handful of Silicon Valley companies working with the U.S. Government. Now we see large tech companies with massive internal resources, or a handful of companies that are being massively funded by venture capitalists, in the hopes that they can land meaningful contracts from the U.S. Government in AI.

I think I have two points. The first is point one and three of my oral testimony. For Federal agencies looking to bring in AI technology, it is important to have datasets that are curated and readily available during the contracting solicitation time, as well as during when that contracting is evaluated.

Models can be benchmarked even before the technology can be acquired. Often when we worked on contracts, waiting for government furnished data was the longest lead time on contracts. For faster evaluation, the less taxpayer money is spent. I want to commend National Geospatial Intelligence Agency (NGA). They have been doing a great job getting their data ready for AI.

Second, to your point, right, anyone can download AI models off the internet and claim AI expertise. Again, having testing and evaluation datasets ready, both at the contracting process and throughout the entire execution of the contract.

It will be easier for procurement officials to be able to both find and evaluate tech through quantitative and qualitative means.

Chairman PETERS. All right. Thank you. Dr. Li, as legislators work to establish key values for the American use of AI, the idea of explainability has received substantial attention. However, building responsibility and trust into the use of AI seems like it is going to require more than just an understanding of the math behind a model.

Could you offer some specific suggestions as to how we ought to prioritize when evaluating AI systems for potential use by Federal agencies?

Dr. LI. Yes, thank you for the question. Actually, explainability is a very much used word in the talks of AI. If we think about it, it is actually a very nuanced word. I will give you an example that is not AI. For example, you talk about math.

When a bottle drops, we can actually use a Newtonian mathematical equation to explain why the bottom drops. But when it comes to the usage of Tylenol, other than some doctors and biomedical researchers, even as a consumer, I do not know how to explain how Tylenol works, yet there is an explainable process of how the Federal Government has regulated it so that I can trust it.

Or another example of using Google Map to go from point A to point B. There, the explainability is not in mathematical equations nor in regulatory framework. It is more about the options it has provided to me, the fastest route, or the avoid the tolls, and so on. I am using these three examples to show you explainability is a very nuanced term and depend on the use cases. It depends on the system approach.

We have to think about it carefully. Again, I advocate for a systematic approach in thinking about explainability, but do put the human values, as well as our society's value at the foreground of this, and depending on how we use it, will require a different kind of explainability.

Chairman PETERS. All right. Thank you. Ranking Member Paul, you are recognized for your questions.

Senator PAUL. If it is OK, could I defer my questions and allow Senator Hawley to go?

Chairman PETERS. Yes, that is fine. Senator Hawley, you are recognized for your questions.

OPENING STATEMENT OF SENATOR HAWLEY

Senator HAWLEY. Thank you very much. Thank you, Senator Paul. I thank you, Mr. Chair. Thanks to all the witnesses for being here. Mr. Shellenberger, I want to start with you, if I could. I am so glad that you are here with us today, and you are here at a significant time. I am looking at a piece from yesterday, I think it is, yes, that you published, U.S. Intelligence Dangerously Compromised, Warned CIA and FBI Whistleblowers.

You are not the only one to report this, of course, but I was reading your report on it this morning. This is something that you have been warning about for quite some time. The allegations stem from a whistleblower who has come forward to the House, a whistleblower from the Central Intelligence Agency.

I have the letter, the relevant letter here from the House Oversight Committee. The whistleblower alleges that a CIA team was paid to change its assessment of the origins of COVID–19. Do I have that broadly correct? Is that your understanding of the report?

Mr. SHELLENBERGER. Yes, sir.

Senator HAWLEY. This is obviously a bombshell report. Deeply troubling. I am glad that the House is going to look into it. We should look into it. What caught my attention, as you point out in your article on this, that the government has deliberately violated the COVID Origins Act, which this body passed unanimously, which the House passed, the President signed into law.

Maybe wasn't so happy about signing it into law, but he did. It is the law of the land, and which required that all of the government's intelligence on the origins of COVID be made public. Instead, what the Administration did was offer up a summary, which they then in turn heavily redacted.

You point out that in addition, the Administration refused to report the names of scientists who fell ill at the Wuhan Institute of Virology in 2019, despite the fact they know the names.

The intelligence community knows the names. Now, you are absolutely right to say this is a violation of the COVID Origins Act, and I would know because I wrote it. I am not very happy about the fact that this Administration continues to flaunt, flout, completely ignore public law passed, again, unanimously by the U.S. Senate.

For what end? I cannot tell. I cannot figure out why in the world. I do not know what partisan gain there is to it. Why in the world they want to lie to the American people. You conclude your article by saying, the government has become extremely comfortable with lying to us. Just explain what you mean by that and tell us why you think this is so significant.

Mr. SHELLENBERGER. Sure. Just on the very specific point, if we were the first to identify the three people that contracted the coronavirus in China. They were the people working on gain of function research in the Wuhan Institute of Virology.

The Wall Street Journal confirmed our reporting two weeks later, and then I think it was one week after that or a few days after that, the Director of National Intelligence (DNI) report came out and it did not reveal this information. We had multiple sources. We have no idea if the Wall Street Journal's sources were the same.

But I think we are clearly seen a lot of abuses of power occurring in multiple executive agencies. We have seen it with the FBI. One of the things that we noted yesterday was that we saw perverse incentives in the FBI to go after so-called domestic violent extremism (DVE), pulling an agent off of things like child exploitation, onto really hyping a set of cases that particularly appeared to be aimed at spreading disinformation around the idea that there is a significant increase of domestic extremism when we do not think that the evidence shows that.

Now we see this report that came out that suggests that there is an FBI whistleblower who says that six of the seven analysts had said it was a laboratory origin and that they had reversed their position in some exchange for some sort of a salary bonus or some sort of a financial incentive.

We keep documenting it. We just keep finding agencies and agencies, DHS involved in trying to create a disinformation governance board. The censorship industrial complex, we just keep finding new parts of it. In the research for this testimony, we discovered this Deep Trust Alliance that had what appears to be ties to the security and intelligence agencies of the U.S. Government, appears to be trying to set itself up, although it is now kind of ghosted after 2021.

But it appeared to be trying to set itself up to decide what is reality and what are fakes for people, and I think it should have a chilling effect that is not how we do free speech in America.

We do not have government agencies, we do not have cut outs or front groups that appear to have support from those agencies, telling the American people what is true, what is false, or telling social media companies behind the scenes what they should be censoring.

Senator HAWLEY. To that last point, we now know, thanks to the case of Missouri v. Biden, that that is exactly what this Administration, from the White House, to the FBI, to the State Department, to the CDC, to Cybersecurity and Infrastructure Security Agency (CISA), have all been meeting with the social media companies for years now, giving them direct commands about what to censor and takedown, naming specific accounts and specific speech they want suppressed, threatening the social media platforms if they do not do it.

Remarkably, and I am quoting the court here, the Fifth Circuit Court of Appeals, and there is a huge evidentiary record. Everybody don't take my word for it. Go read the record. It is all on the record from the District Court. What the Fifth Circuit said, it is remarkably that the social media platforms all complied. All of them.

They all agreed to be tools of the U.S. Government and to censor what they were ordered to censor, to suppress the speech they were ordered to suppress. You are a journalist. Tell us about the threat to the First Amendment—and by the way, just for the record, I think it is important to establish, the Federal Court of Appeals said directly in no uncertain terms, this was a clear violation of the United States Constitution.

The First Amendment does not allow the Federal Government to use private companies to censor what they would not be able to do it themselves, and that is exactly what this Administration has done.

Tell us, as a journalist, the threat to free speech, to freedom of the press from this kind of collusion between a very powerful government trying to hijack every media company it can get its hands on.

Mr. SHELLENBERGER. Sure. If you start on the issue of the COVID vaccine, for example, public interest advocates spent a very long time requiring the pharmaceutical companies to list the side effects of their drugs in their advertisements.

Here we saw a situation where people were sharing information about the side effects of the vaccine on Facebook and other social media platforms, and the White House demanding that it be taken down. Facebook complying, acknowledging that it was often true information.

We also saw that Facebook's own internal research showed that actually it increases vaccine hesitancy when you censor those stories. That people, if they want to be comfortable with a new drug, they need to be able to talk it out of it. Facebook told the White House that actually it would backfire.

The White House insisted. Facebook caved in because according to the Facebook executive Nick Clegg, he said, well, we have this other business that we need to do with the White House, which is

the data flows. Meaning we need the White House to help us negotiate with the Europeans to bring our data back to the United States.

I think the Fifth Circuit Court did a great job in identifying the clearly coercive measures, but I do not think it went far enough because the First Amendment, it prevents the government from abridging or infringing on free speech.

Offering an incentive to social media platforms such as helping them with their dispute with Europe in exchange for censoring often true content, though, of course, the First Amendment also protects false content, I think is a very chilling effect.

I think it is very disturbing. Anybody that cares about holding powerful entities to account should be disturbed by what we saw take place on Facebook, on Twitter. I think we have to remind ourselves—and what disturbs me, when I hear sort of the conversation around AI coming into it sort of with a beginner mind, I hear a lot of talk about how to protect the public from harm.

We have to protect the public from harm. What people are saying is that we need to censor speech, censor certain voices, censor disfavored voices because of this idea that it will cause real world harm. This is a well-documented phenomenon that psychologists have measured where over decades people have grossly expanded their definition of things that cause harm.

I think this should be a moment for a reset. That free speech is almost absolute in the United States, with a few exceptions, around immediate incitement to violence, around fraud, around child exploitation. But we allow very open conversation in the United States. It is what makes us so special. It has been a chilling effect. As a journalist, I personally have been censored by Facebook. I think the platforms are out of control.

Senator HAWLEY. Thank you, Mr. Shellenberger. Thank you, Mr. Chair.

Chairman PETERS. Thank you, Senator. Senator Blumenthal, you are recognized for your questions.

OPENING STATEMENT OF SENATOR BLUMENTHAL

Senator BLUMENTHAL. Thanks very much, Mr. Chair. As you may know, Senator Hawley and I have authored a framework for protecting the public against some of the perils, I would like to think all the perils, of AI through an oversight entity that would license new models, would require testing and red teaming, require transparency, a notice when AI is used, but also accountability on the part of AI companies, and a means of enforcement.

That is a very rough summary. But the point is that we have been working through the Judiciary committee, the subcommittee of the Judiciary committee on Privacy, Technology and Law, and I want to thank Senator Peters for his focus on AI in this Committee as well. But we had a very useful and productive forum yesterday.

My hope is that many of the people who came before us in that forum will agree to come before our Subcommittee and testify in public, under oath, and give us the benefit of their views, which they expressed to us privately as Senators, because I think the

public has a right to know, and we should be putting these views on the record.

Transparency in our process is as important as transparency in the disclosure of how algorithms work, how AI works, so the public has a better understanding of it. I think we are going to be pursuing our framework, putting it into legislative form.

We have gotten tremendously positive response. I would say that almost all the provisions of our framework, in fact, all of them were endorsed by one member of that group yesterday or another, and the vast majority of the group endorsed the core provisions of our framework.

We are making progress. I think what is important about the chair's actions here is he has sponsored a bill called the Transparent Automated Governance Act. I do not know whether he has mentioned it yet, but I am going to be joining as a co-sponsor. What it requires is more disclosure to the public about when they encounter AI.

In fact, a Subcommittee of this Committee, which I chair, the Permanent Subcommittee on Investigations (PSI), held a hearing recently on Medicare Advantage, which is a government program, health insurance, that provides key coverage for people who are eligible for Medicare, and they can choose to go into this program.

I am vastly oversimplifying, but not in any way diminishing the key point here, which is, Medicare Advantage insurance companies are using AI to make decisions about what they will cover or not. Some of these decisions cause denial of coverage to people who then have to try to access the system to get overruled a decision which essentially is made through AI.

I know that I am somewhat simplifying, but the key here is that AI is making decisions that hugely impact people's lives, often without their knowing it. That is why the chairman's bill I think is so important. I see a number of you nodding, I hope it is in agreement, with my basic point, which is that disclosure is very important here.

I will just ask you, perhaps beginning with Dr. Ghani, will you support a legal obligation for AI companies to disclose when a person is interacting with an AI or decisions being made about them using AI?

Mr. GHANI. Absolutely. I think it has to be critical. I think more than that, it is not just disclosing when you are interacting with the system, because you might be interacting with a human that is informed by an AI.

It is nuanced, but you want to make sure that if a decision is being supported by AI, it is not just the person knows but I think even more importantly, they need to have recourse. That exists in other areas. In financial services, we have these things called adverse action notices.

When a decision is made against you for denying you a loan, the bank is supposed to tell you why, and then allow you to change those characteristics, and then give you the loan if you changed that. I think the same thing needs to be present as one example of extending the current AI systems. I think procurement, again, is the perfect vehicle to force that.

Senator BLUMENTHAL. Thank you. Dr. Li.

Dr. LI. Thank you, Senator. I have to say, when you talk about Medicare Advantage, this is the life I live in, home caring two elderly parents who are chronically ill, and I have personally experienced claim denials and getting on the phone forever to talk about all these cases.

Yes, I think disclosure is part of a systematic approach to how to use these powerful tools intelligently and also responsibly. It is really important to recognize AI is a tool, and every tool can be used to our advantage but can also have unintended consequences. For example, we all take airplane rides, and we know there is autopilot in the airplane.

Yet we know there is enough regulation and disclosure and responsible use to feel safe to a large degree about this. Right now, this technology is so new, we, multistakeholders, need to get on the table and have a nuanced approach to these critical issues in high risk areas, like disclosure, trustworthiness, privacy preserving, all these issues.

Yes, thank you for your effort.

Senator BLUMENTHAL. Thank you.

Ms. RAJ. Thank you, Senator. These are issues that are important to the American people. They are common sense measures. We realize that we do not want to stifle innovation, but there are guardrails in place that make it important, so thank you.

Mr. ROBERTS. Yes, Senator, I also agree with your statements. But I would say also that this is one of the challenges of government contracting, is putting the language in the contract versus once the contract is awarded, making sure that these perimeters are followed.

What I have seen is usually when we award a contract, the focus of trustworthiness on the functionality of the tool itself usually takes a front seat, and then all the planning of how we would do this responsibly sometimes takes a backseat. We focus so much on will this work, will this benefit the users.

We are the largest, the Federal Government is the largest buyer in the world. We have so much ability to put forceful language in our contracts. But it is also what happens after the award and the quality assurance measures we are taking to make sure this responsible use is not taking a backseat.

Senator BLUMENTHAL. Thank you. Mr. Shellenberger.

Mr. SHELLENBERGER. Yes, absolutely.

Senator BLUMENTHAL. Thank you. Mr. Chair, again, thank you for your leadership. The chairman's bill was reported out of the Commerce committee in July, and disclosure is also an essential part of the framework that Senator Hawley and I have advanced.

I am looking forward to supporting his bill, and perhaps in our efforts to combine our ideas and our forces, taking them up together. But this panel has offered some really important insights and I really want to thank you all for being here today. Thanks, Mr. Chair.

Chairman PETERS. Thank you, Senator. Thank you for co-sponsoring the bill. We will hopefully move it through the Senate as quickly as possible. I need to step aside briefly for another committee hearing that is going on to ask questions, so I am going to

pass the gavel to Senator Hassan. But before I do that, Ranking Member Paul is recognized for his questions.

Senator PAUL. Thank you, and thanks to the panel for being here today. I think we have had a good discussion, pros and cons of AI, how we can kind of control it to make sure it does not lead to abuses.

But I think as we have, and if you think about potential uses, we had a hearing a month or two ago talking about classified data. We have like 25 million bits of data or something sitting out there. It is supposed to be declassified after 25 years. No human is ever getting through it. We need help.

Something like AI could be of great benefit. But as we have talked about different ways to try to control AI, either it is contracting, or transparency, or this rule or that rule, I think what we are missing is how far apart we actually are. We think we are all together on let us have some controls, but I do not think we are very much together, if you think about the most basic of rights would be our Bill of Rights.

Among the Bill of Rights, the first one is supposed to be one of the most important. The Supreme Court says it has special scrutiny for the First Amendment, and yet we have absolute complete disagreement on the First Amendment.

The idea of whether or not it is in breach of the First Amendment for the FBI and the Department of Homeland Security to meet on a regular basis with Twitter and Facebook and talk about the content, what people are saying, people's speech, and limit that.

If you look at the court case, it is worse than that. It is not just talking about or notifying, it is actually threats to say that we may take away your Section 230 if you do not comply and take down this. The Section 230 is the liability protection, we may get rid of that.

We may institute antitrust rules to try to break your company up, which they are apparently going after Google now anyway. We also may notify the top guy, the President will be informed that you are not taking down this information.

But even more than that, I think they went even further. It is my understanding that when Twitter said, oh, yes, we will take this down, but it is a heck of a lot of work, can you pay us? I think the FBI actually paid them $3 million. Mr. Shellenberger, can you comment on did Twitter take money from the FBI to take down content?

Mr. SHELLENBERGER. It did. You are absolutely right that that happened. There was some controversy around it because Twitter was being reimbursed for the time that it was spent in helping FBI. But nonetheless, it appeared as though Twitter had refused that money previously because they recognized there would be a conflict of interest.

After the former Chief Counsel for FBI, Jim Baker, came to Twitter, they changed that policy and they did take that money from the FBI after working with them on this.

It was one of many, what the Fifth Circuit Court recently called it, kind of close nexus, which is the kind of thing that you are worried about, a close nexus in terms of censoring content online that we saw. To see the financial incentives there was very troubling.

Senator PAUL. This is sort of what concerns me, because we think, oh, through contracting or through rules or transparency, somehow we are going to control this.

When I am concerned that I do not think there is one Senate Democrat that has criticized the FBI and Department of Homeland Security meeting on a regular basis to discuss taking down speech.

Even true speech, as you mentioned, the Virality Project, stuff they said, well, so-and-so had the vaccine and they are in the hospital today, and it was verifiably true. Now, cause and effect, everybody has an opinion on what caused what, and yet they were taking that down, even admitting that it was true, taking that down, and that we do not seem to have any concern on one side of the aisle.

How are we going to get to fixing this problem with contracting and transparency if we cannot even agree on what the First Amendment is. This is a big deal, and I hope this case, it has now been decided at the appellate court level. I hope it makes it to the Supreme Court level so it can be basically adhered to across the land.

But, as much as I want to say, oh, let us all get together, we are going to hold hands and have transparency and fix the contracting in AI, I am more than worried than ever. I am not an opponent of AI. My son does AI. I am a fan of technology. I think AI can do great things.

Yet, I am worried that one entire party, about half of our country, representative wise, does not seem to have any concern about the First Amendment, about the FBI meeting. Now, if this were the 1960s, it would be interesting where it would be.

In the 1960s, if and when people heard that J. Edgar Hoover was meeting and looking through MLK's mail, looking through Vietnam War protesters, at that point in time, the left was much better than the right, and they were absolutely exercised. They wanted to defend the First Amendment.

Somehow we have lost that. It really should not be partisan. I do not care who is the President and I do not want any President sending the FBI. The way I try to put it so people can comprehend this. I will do interviews on television.

Let us say I am on there and the woman on television, the man on the television interviewing me and I say, after you are done, what would you think if the FBI called you and wanted to sit down and discuss my interview, and because I said I do not think masks work in the general public and do not prevent the trajectory of the virus, that is my opinion, I can give you 25 studies to support it, but how would you feel if the FBI sat you down and said, we are worried about that and we think that is misinformation.

No broadcast TV would ever stomach that. No newspaper would ever. The Washington Post would not stomach that. Yet nobody in the left seems to care at all about the FBI meeting with Twitter, which is arguably maybe more powerful than all of the traditional legacy media anymore.

You mentioned that already in their algorithms, Mr. Shellenberger, that 99 percent of it is being taken down through artificial intelligence. You mentioned that a lot of this originated in DARPA. I assume you mean the funding for originating and dis-

covering how to do artificial intelligence. Is that what you are saying came out of DARPA?

Mr. SHELLENBERGER. Yes. Also, and I did not get in as much detail about it, but also obviously DARPA has been—or maybe not obviously, but in the process of creating deepfakes and then creating technologies to detect those deepfakes, we have also seen, we know there is cases where they are actually identifying persons of interest to be developing deep fakes around.

I think there is a set of things going on that people are not aware of that they should be aware of in the development of these technologies. Then there is an ostensibly civilian process been going on to try to govern these deepfakes, to try to kind of establish some separate authority to decide what is deepfake, what is not.

Senator PAUL. To point out the problem, though, if DARPA is in charge of a lot of this, they are a small, secretive group that won't respond to me. They are the group that was going to fund research in Wuhan to stick a furin cleavage site into a virus, which turns out what COVID looks like.

Yet when I asked them, even today, I have been asking DARPA for their information, they won't give it to me. Now I am supposed to trust DARPA with artificial intelligence, an agency that won't give me unclassified documents with proper request.

That is a real danger that DARPA, Defense Threat Reduction Agency (DTRA), all these defense and intelligence agencies that are developing this to spy on others are so secretive that we don't have oversight. I think they are in government without oversight, and I am worried about that as we move forward, that we are sort of saying, oh, yes, we will just have oversight, we will have contracting.

I am not against that. I think that is good. But how can I contract something when I don't even know what their budget is. Any comment on the secrecy of DARPA?

Mr. SHELLENBERGER. The secrecy in some ways is the main event and that is how the censorship has been taking place, behind closed doors. It is not jawboning. They have sort of said there has been this argument that politicians should be able to get up there and criticize somebody publicly, but that is not what was going on.

It was behind the scenes work. I am actually glad to hear that they have been as unresponsive to you, Senator, as they have been to us. We went to over 50 or 60 of these organizations, most of which had some sort of government funding, to just have a conversation with them about their censorship activities, and not a single one of them agreed to be interviewed. This is not the kind of transparency you would expect from government contractors.

OPENING STATEMENT OF SENATOR HASSAN

Senator HASSAN. Thank you, Senator Paul. I am now going to recognize myself for the next round of questions. I am really glad we are having this hearing. I am really grateful to all of you for being here and being a part of it.

I want to start with a question to you, Dr. Li. I am Chair of the Subcommittee on Emerging Threats and Spending Oversight (ETSO), and I am concerned about potential public safety and national security ramifications of AI. Dr. Li, you have raised concerns

about AI behaving in unintended or unpredictable ways that could be detrimental to Americans.

What are the key considerations for Federal acquisition and procurement policy to ensure the safety of AI?

Dr. LI. Thank you, Senator. I have been known to say there is no independent machine values. Machine values are human values. Especially when it comes to U.S. Government, I think we really need to care about, as Senator Paul said, our Constitutional values, our Bill of Rights, and all that.

In the procurement process, again, I think we must take a systematic approach to ensure that the kind of investment and support of the AI systems we want to develop and deploy reflect our values. This includes starting from procurement of data, the privacy issues, the bias issues, the trustworthiness issues, all the way to the development of the system itself.

All the way to what in machine learning we call inference, which is when you have developed the algorithm, now you are ready to deploy it and do things. Here you, again, have privacy issues, bias issues, and trustworthy issues, and the whole ethical framework.

Again and again, I want to say that this is a powerful tool, and it has good and bad sides. We need to take a systematic approach, and every step of the way we should apply responsible and ethical values to this——

Senator HASSAN. Acquisition—to the standards that we use in acquiring this technology.

Dr. LI. Yes.

Senator HASSAN. Thank you. Mr. Roberts, I want to turn to you. Last year, I worked with Chair Peters to lead the bipartisan effort to codify Federal Risk and Authorization Management Program (FedRAMP), the Federal program run by General Services Administration (GSA) that evaluates cloud service providers (CSP) and their products for use by Federal agencies. FedRAMP promotes efficiency and increases security by having one agency responsible for vetting and approving these companies and their products.

Now, I understand AI is a little bit different, but could a FedRAMP type program, or an entity with aspects that FedRAMP has, that is designed to specifically evaluate AI products be a feasible option for evaluating the safety of those AI products?

Mr. ROBERTS. Yes, Senator. I think that would be beneficial. So much of AI is software centric and we have had our projects challenged by security requirements to include authorities to operate in FedRAMP.

We found ways of working with that, especially with prototypes. But we have been looking for new guidance and new ways in which we can make this easier, especially for the developers and contractors.

Senator HASSAN. OK. Because what we are really looking toward is conserving resources as we try to buildup these standards and apply them in the acquisition process, but also really having a centralized place and one set of standards that are applied across. So, you think something like FedRAMP might be applicable?

Mr. ROBERTS. I think that would be helpful, Senator.

Senator HASSAN. Another question for you, Mr. Roberts. During my time in the Senate, I have focused on reducing the Federal Gov-

ernment's reliance on aging technology in order to save taxpayer dollars, improve security, and obviously provide better service to the American people. I am concerned that Federal efforts to adopt and use AI may not be successful if we continue to rely on legacy IT.

Could the Federal Government's aging infrastructure prevent us from effectively adopting AI technology? Alternatively, are there ways in which artificial intelligence could help agencies convert from costly legacy IT to modern systems that provide better services and are more efficient?

Mr. ROBERTS. Yes, Senator. I do not think legacy systems will necessarily prevent us from implementing AI, but I do think that the Federal acquisition workforce needs to be more trained on making that analysis of whether to sunset a legacy system, which has been done successfully in various areas of the Federal Government, and bring about a more modernized, a flexible system with little disruption, or to try to modify that system through Application Programming Interface (APIs) with AI functionalities.

But I would say that the best approach for an agency that wants to introduce AI into their organization and their systems that leads to the most success is to start small, and start narrow, and start feasible, with minimal risks in terms of responsible use.

We have seen many cases of that where it is just automation through AI of business systems that create huge impact, that bypass any risks toward ethics or responsible use. But also, because they are done in isolation, they can yield immediate impact to the end users.

Senator HASSAN. OK. Thank you. Last question. Again, to you, Mr. Roberts. I helped lead a bipartisan effort to codify the General Service Administration's IT Modernization Centers of Excellence, including the Center of Excellence for Artificial Intelligence.

The AI Center of Excellence assists Federal agencies seeking to use AI tools and acts as a centralized resource center for agencies looking to develop policies around AI. In your view, is the AI Center of Excellence equipped to provide support to all agencies, especially smaller agencies seeking to procure and adopt AI? If not, how could it be improved?

Mr. ROBERTS. Yes, Senator. We had a good relationship in the JAIC with the GSA Center of Excellence for AI. I would say that this is still a very challenging topic to try to at scale create policies and guidance to deliver and procure and deliver artificial intelligence.

I think right now it is still a pocket. They are a good organization, but it is a pocket. Whereas it should not be an anomaly. This should be more mainstream and widespread, and hopefully they will contribute to that effort.

Senator HASSAN. OK. Thank you. Senator Lankford.

OPEMNING STATEMENT OF SENATOR LANKFORD

Senator LANKFORD. Thank you. Thanks to all of you all. This is an exceptionally complicated issue, obviously, as you are trying to wrap your head around that, but some things are pretty straightforward with this.

The most difficult thing that I continue to be able to hear is the continual term everyone throws around, we just want to make sure there is responsible use of AI, to which I always smile and say, define that for me. I have yet to have anyone be able to really wrap a good definition of what responsible use of AI means and does not mean.

Does anyone want to try to jump in on that and give me a concise definition of what responsible use of AI means? I am seeing a long pause. This is not a trick question. I am just trying to because this is an issue that we have to be able to both regulate and be able to wrap around.

Everyone is saying, hey, there needs to be some kind of controls here for responsible use. What does that mean?

Senator CARPER. If I could interrupt for just a moment. I am going to Chair this Committee for just a little while. He is asking a great question. Do not be afraid to answer—and even if you are going to change your answer later on, but give it a shot, if you would, please. Go back to Senator Lankford.

Ms. RAJ. Thank you for your question. AI is a powerful tool, we believe, in large part that can keep people and society safe. But fundamentally, as Dr. Li said, it is really about the human knowledge guiding it.

I think that when you think about responsible AI and implementing responsible AI, it is really about equal protections of the American people. Making sure that everyone can benefit from the AI system in an equal manner.

Dr. LI. Can I answer that?

Senator LANKFORD. Yes, ma'am.

Dr. LI. First of all, great question. AI is a tool, just like many tools. When we ask about responsible AI, I personally start with a responsible tool use. It is not just AI. I do believe there is no one size fits all answer.

A responsible drug is different from a piece of financial service product because every use case when rubber meets the road is different, and when we think about these, we start with the value.

There are very high level, Constitutional values America represents, and within that framework, it is important to look at the different use cases and define responsibility, responsible, and trustworthiness within those framework.

Senator LANKFORD. Who would you suggest actually sets that framework for what is responsible AI in each one of those categories?

Dr. LI. I personally believe in multi-stakeholder approach. I think, for example, going back to health care, right, if it comes to drugs and food and all that, FDA should be a huge part of it. So does the consumers and the industry, as well as public sector.

Senator LANKFORD. It is not a simple issue, to say the least on it. But it is one of those things that hangs out there because as we discuss a regulatory framework, you want to hang a regulatory framework on a set of values and to be able to say these match our values.

We have not been able to get good clarity among all of us on, how do you define what is that value of what is responsible use of AI in each sector, and I encourage anyone who is participating in this,

because I agree, multi-stakeholder is helpful. There is 320 million people smarter than any one of us.

To have that engagement, as people have ideas on that, they need to be able to contribute those because we have to set that value set on this. Obviously, the Constitution is the first set of values, but then we have to be able to set for each one of those, what is that responsible use?

We have another big challenge in the Federal acquisition procurement process that deals with software is who owns the intellectual property. Because if we do weapon systems, there is just, we are going to keep updating it, and we have to know this.

With AI, it becomes particularly difficult on who owns the intellectual property for this and when changes need to be made. How do you know that you have a safe process to be able to do updates on this?

Or if someone were to get into the system and to be able to change some of the system, how do you know it and how do you track that? The first thing on AI for me is on the procurement process, how do we handle who owns it?

It is one thing to be able to buy it off the shelf and to say we are going to use this product. It is another thing to say this is a unique product that we are going to use as a Federal product that has AI that is built into it. Somebody want to try to step into that?

Ms. RAJ. Yes. That is something that we had to deal with as a vendor for the JAIC, licensing of the architecture, licensing of the model, with licensing of the government data. That was something that we had to really navigate alongside JAIC.

I think fundamentally, AI needs to be purchased in a manner that reflects the updating nature of this rapidly changing technology. To my second point, I mentioned that a contracting should have the ability to include ongoing retraining.

We do not want taxpayers buying model 2.0 when weeks later, model 4.0 is released. It is about really taxpayer saving money while ensuring the latest technology is in the hands of the U.S. Government.

Senator LANKFORD. Yes, we have a weapon system that was delivered about five years ago that was delivered with Microsoft 7 on board five years ago, and the entire system was built on Microsoft's 7, and all the software that connects to it is all built on Microsoft 7, which is not even updated anymore.

We cannot have a situation where we are that far out of date, especially with something like this. But the other challenge is, when we do security checks for most Federal systems, we are trying to inspect to make sure that there is nothing within the system that is a problem.

That is uniquely challenging with AI, because we have to get into how has it been trained, what is the decisionmaking and the process, how do we verify that decisionmaking process when we are verifying it for security, for biases, for all those things?

If there are ideas that are out there that you have seen for that type of verification, that would be exceptionally important to us because we are not going to buy an off the shelf item for AI and just assume it has good ethics that is built into it, or it has good security built into it.

One of the thing that I want to be able to mention, then I want to pass this on to other colleagues who want to be able to ask questions. You had mentioned before about AI is a tool, and I totally agree. There has been some dialog even here in this Committee, about putting an AI representative in every agency.

I have real concerns on that because if you have an AI representative in each agency, their job is to increase AI usage. That would be like having a screwdriver representative in each agency to find a way to increase the screwdriver uses. It is a tool.

We need to treat it as a tool. I have concerns when there is a focus on, it is the latest thing we talk about, so let us proliferate that in the Federal Government. If it is a tool that works, fine, but we need to have some ways to be able to verify security, verify on the updating, verifying the IP address, and then verifying what is not the IP, but who owns the IP on it, and then verifying the whole process of updating it, using it, and then also what is responsible use of it.

But I just have to be able to tell this body, I do not think we should have an AI office in each one of our Federal entities. I think we have AI specialists within our technology folks, but not focused on trying to get more of it. Did you want to make a quick comment on that, Ms. Li?

Dr. LI. I want to respond by saying AI is very powerful and it is a horizontal technology. It is really important that many of us leaders of our society have that knowledge.

This is why under the AI Training Act, Stanford HHI is actually committed to creating educational opportunities to our policy leaders and policy members of the policy world, because it is not just one person or one person per agency's responsibility.

It is our collective responsibility, and having that basic level knowledge, and also in some cases, specialized knowledge. It is really important to recognize that.

Senator LANKFORD. To the former chairman and the current fill-in interim chairman of the day, thanks for your gift of an extra minute there in time. I appreciate that.

Senator CARPER. Any time. Thank you for those very thoughtful questions and good responsive answers. Senator Rosen, you are next.

OPENING STATEMENT OF SENATOR ROSEN

Senator ROSEN. thank you, Senator Carper. Thank you to everyone here. I appreciate all the work that you have done and being here today, and I am just going to get right into it because I know we think AI, it is all iterative and it is going to learn on its own and do all of that.

But then we know it does not all build itself, and so we need a Federal AI workforce, one that is trained. The White House Select Committee on AI has worked to enhance the overall effectiveness and productivity of the Federal Government's efforts, of course, related to AI. Earlier this year, the Select Committee on AI released the Strategic Plan for AI, which identified significant challenges with the AI workforce.

I have a two part question for you, Mr. Ghani, and then a follow up for you, Ms. Raj. Mr. Ghani, to that point, how do we ensure

the existing Federal AI workforce has the necessary skills to buy, procure as we are talking about, build, and field these AI enabled technologies?

Are there programs that currently exist to upskill our Federal workforce? You can think about it, Ms. Raj, your follow up will be, if you can speak to it, how can academia and industry talent be leveraged to meet these dynamic needs? Because it is not a static industry for sure.

Mr. GHANI. Thank you for the question. It is a little bit related to the previous conversation on, do you train everyone on everything, or do you have specialized people? I think this is a dilemma that even the private sector has been focusing on for the last 20 years in this space.

Where do you build one central team that is a tech team, an IT team, an AI team, and then you have them help everybody else? Or do you enable each agency, each department, upskill them? For me to answer that—what has worked is the latter. It is because, again, AI is not a generic tool, it is not Microsoft Office, it is not Windows, it is not a word processor. It is different. Then what you need to have is you need to have it de-configured and used for specific applications, specific programs, specific policy areas.

If you are bringing in tech people in a centralized way, you need to now have them be trained in every single thing that every single agency might do.

That is not possible. I think in my mind, the way to do this is to augment the existing trainings of people within different departments, within agencies, and enable them to do their work in an AI augmented world rather than AI first training.

To your question, Senator Rosen, on do these types of programs exist? Not exactly. There are programs that exist in pockets. The majority of programs that exist are in upskilling and in several universities and several nonprofits, including Carnegie Mellon, and Coleridge Initiative that I have been part of.

We have created programs to train government agencies and workers in the use of these types of technologies. But not enough of them exist and they are not at the scale that we need very quickly.

We need to enhance the training, but I think there are two pieces there, and I will echo some of the things my colleagues have said.

Senator ROSEN. I am a former coder. I wrote software for a living. I get it.

Mr. GHANI. They need to be kind of experience centered, grounded in real problems, and have people solve these things, which requires access to data, which requires access to real problems and access to experts who can help.

These programs are expensive, and these programs do not scale because you cannot just put them up and have Massive Open Online Courses (MOOCs).

I think we need an investment in training government agencies and not just in using them but being at the forefront of helping design them. Because, again, the needs are very different for governments than they are for the private sector.

Senator ROSEN. Ms. Raj.

Ms. RAJ. Thank you, Senator, for your question. It is evolving at a rapid pace that even folks like myself in the field have a hard time keeping up. There have been pockets of incredible initiatives across the U.S. Government, namely the MIT Air Force AI Accelerator, which is a multi-disciplinary team of embedded officers and enlisted airmen who joined MIT faculty, researchers, and students.

But what is really interesting about it is that leadership at all levels participate in these workshops that are reflective of the mission need. I believe it is interdisciplinary teaming through AI workshops and training alongside both industry and academia that make this possible.

But again, it is not a one-off workshop. It is a continuously evolving relationship between academia and industry.

Senator ROSEN. Thank you. I want to move on a little bit to AI procurement, because we have heard testimony in this Committee about, of course, we all know how quickly AI is evolving. This is no secret there.

Giving the breakneck pace of the AI evolution and, of course, much iterations, how quickly is going to learn on its own, I am concerned that our Federal Government's acquisition process is just going to not be able to keep up with the rapid pace of the development of the AI tools.

This question is for you, Mr. Roberts. How should the Federal procurement process, how could we improve it, streamline it, to be sure that AI products are purchased at the speed of relevance, meaning that they are not rendered obsolete before we even benefit from their use? How should our Federal contracts account for the need to retrain a procured AI system? They often require ongoing updates, security patches, and monitoring.

I am a former computer programmer myself. I have been advocating for what we call a software bill of materials (SBOM) so you know—I guess like if you look at the back of a cereal box, you know all the ingredients that are in there.

We should be able to know that for our government software, so we have to make the appropriate changes so we know the list of ingredients, if you will. If you want to speak to that.

Mr. ROBERTS. Yes, Senator. I would start by saying I think the best way to change the processes is to invest aggressively in our talent. I would say that there should be nothing short of a mandatory AI training for all interns coming into the acquisition workforce.

I think this is the only way that professionals coming in will really understand. The reason why is twofold. I think one is if we really truly believe this is transformative technology, which I think it is. It will affect every mission. Every contracting officer in the future will have to know what this is.

Right now, they are set up for failure. The second reason is, to acquire AI, it requires a major reskilling of the way processes are done. There is a mantra in contracting that says, poor planning equals poor contracting. What that is translated into is that if you do not have all your ducks in a row about exactly how this project is going to pan out, then you should not release your solicitation.

Obviously, that does not work with artificial intelligence. What we have, what we need is a team that works in an iterative fash-

ion, which is hard right now in the government culture. What we need is a team that focuses on, we have something called modular contracting in Federal Acquisition Regulation (FAR) Part 39.

It is underused, but it is there in the Federal Acquisition Regulation. But we also need performance that is based on value add to the actual end user. There is a surprising lack of focus in the end user just because there is so many rules and regulation and compliance.

Senator ROSEN. I am the last one here, so I am going to ask my follow up question, which you have led me into perfectly. We have some authorities at the Federal Government to acquire these tools, maybe do this training. What authorities are missing to help move us forward?

Mr. ROBERTS. Yes, Senator. In the Department of Defense, we had a lot of the authorities we needed. The Federal Acquisition Regulation did a lot for us, but we also needed other transactions when it required.

We needed public, private partnerships, which is a great tool for AI because it incentivizes industry for dual use application. But that was the Department of Defense. These authorities are not given to other civilian agencies who will also be entering this field.

My recommendation would be to also provide that authority to those civilian agencies, especially other transnational (OT) authority, and things like partnership intermediary agreements, public, private partnerships. AI is so diverse that these are truly tools that match perfectly depending on the situation. We have used them all.

Senator ROSEN. Thank you. I see our chairman is back, so I will turn it over to you. Thank you.

Chairman PETERS. Thank you, Senator. Senator Carper is recognized for his questions, so thank you for holding the gavel.

OPENING STATEMENT OF SENATOR CARPER

Senator CARPER. God, it was fun doing that again. I used to be him. Now I am me. [Laughter.]

Right. I got to Chair a Committee on Environment Public Works (EPW), which is pretty good gig as well.

All right. My name is Senator Carper. I have been called worse. Welcome, everybody. We are delighted that you are here. We had this big seminar yesterday, as some of you probably participated in, or at least followed it.

I was asked by the press, the press stakeout after it was over and I was leaving, and they said, well, what do you think? Do you know a lot about AI? I told him earlier this year, when we started getting really focused on this, I said I could barely spell AI.

Another reporter said to me, how do you feel? Do you feel like you are coming along on AI now? I said, I quoted the wife of Albert Einstein. She allegedly was asked a long time ago if she understood her husband's theory of relativity. She was very smart woman.

She said, I understand the words, but not the sentences. I feel that pretty well sums up my approach released earlier this year on AI. Hope I am coming along a little bit, taking baby steps, but we will eventually get there. I just remind us all that our members tend to have broad and different backgrounds. I got an MBA a long

time ago, became a naval flight officer (NFO) in Vietnam War, and many years retired Navy captain.

I have been Delaware's Treasurer, Congressman, Governor, and Senator. So, I should be able to spell AI. I am not going to give up and will eventually hopefully be able to make some real contributions through this Committee and others as well. Among the questions I would like to ask you, I have a couple of them, and we will just see how much time we have. Dr. Li.

Where was your family from originally?

Dr. LI. New Jersey.

Senator CARPER. New Jersey. West Virginia, here, so.

Dr. LI. Oh, great. I was a small town. Parsippany, New Jersey.

Senator CARPER. All right, good. Very good. Dr. Li, let me just begin by thanking, again, all of you for being here and for your efforts to help make us, I like to say guided missiles as opposed to unguided missiles. But thank you for your testimony and your thoughtful testimony on the impact of AI.

As we heard this morning, AI has a potential to transform many aspects, not all aspects, but many aspects of our lives, including the speed and effectiveness of government services. We are servants. Our job is to serve the people of this country.

The question here is for me is—how can AI help us to be better servants and to serve the people in this country in a wide variety of ways? Along with helping us with the spread of information, helping us with respect to economic competitiveness as a Nation, and help us with respect to other changing worker responsibilities.

I am curious about how AI will be incorporated throughout the Federal Government, specifically with regards to efforts to streamline service delivery and waste reduction. I do not like to waste my money. I do not like the waste of the taxpayer money. I think that is probably true for all of us who serve here.

Dr. Li, my question. What are some of the ways that you can see artificial intelligence improving the delivery of services to our constituents? That is pretty broad range because we want to provide services in a broad range of ways, but just a couple of examples, Dr. Li, of how AI can help us improve delivery of services to our constituents.

Dr. LI. Thank you, Senator, for the question. First of all, this is actually one of the upsides of this technology, is that it can help productivity greatly. I want to continue on the example earlier that, for example, Medicare Advantage Services, right.

As a user of that for my elderly parents, it is actually highly inefficient right now to have a conversation talking about claims. You can just imagine, especially with these language model technology, that it actually can greatly help our Federal Government to become much more efficient in handling claims and all that.

When I say efficient, I do not mean replacing humans. I mean augmenting humans. This technology can serve as a companion and working assistant to many aspects of our Federal Government's work.

Senator CARPER. If I could use an aviation term, it is more as a copilot.

Dr. LI. Copilot is actually the fantastic term. Even in software engineering now, we call it co-pilot. This is, absolutely there is op-

portunity for a co-piloting. Also, there are what we can call machine in the loop of human work.

For example, there are a vast amount of documents and knowledge we have to sort through. Sometimes AI can become that kind of copilot to help to preliminarily sorting that. These are just simple examples, as long as we adhere to our responsible framework.

Senator CARPER. Doctor, thanks for that. Can you also discuss for us how the adoption of artificial intelligence tools will impact the Federal workforce on a daily basis, and how these tools will impact how Federal agencies plan for the future.

Dr. LI. A lot of Federal Government work is knowledge based. Again, whether we are processing documents, making decisions, and AI right now, especially in the recent breakthroughs, a lot of these language based models are extremely helpful in knowledge work, as a copilot.

In Stanford, we have colleagues working with the Internal Revenue Service (IRS), looking at, for example, tax and looking at fraud detections. I can imagine EPA looking at environmental issues, understanding different aspect of environmental problems. We can use this for example, firefighting and climate help.

There are many ways that AI can help in Federal work. Like my colleague here just said, if we empower our Federal workers, if we continue to educate them, train them, their productivity at work can be really elevated, and that is what I personally hope to see. That is part of American leadership.

Senator CARPER. Mr. Chairman, just one last thought before I yield back to the real Chair. You mentioned the IRS. We have a fairly new commissioner, right, in charge of the IRS. Daniel Werfel is doing a great job, and we are actually beginning to do a much better job, as the chairman knows, collecting taxes that are owed by a lot of folks, including people of great wealth and companies that are highly profitable.

The idea of having better tools for our folks at the IRS is a big plus. I Chaired the Committee on Environment Public Works. We were all about the climate change and the kind of weather that we are seeing, weather conditions we are seeing around the country, and fires, and so forth.

But we have our hands full, and our firefighters have their hands full, and they can use some help. Maybe AI at the end of the day can be of help there too. Thank you. Thanks, Mr. Chairman.

Chairman PETERS. Thank you, Senator Carper. Professor Ghani, a couple of questions. We are going to wrap this up. All of you have been great, but we are starting to wrap this up with just a few more questions. Mr. Ghani, what do you see as essential provisions of contracts involving AI systems that governments are not currently incorporating and should be required?

Mr. GHANI. I think simply what the majority of the procurement process I have seen, it over focuses on the mechanics of the system being procured.

I will give you an example. Many years ago, I was working with police departments on these systems called early intervention systems that were designed to identify police officers who were going to use unjustified force in shootings and detect them early.

When you look at procurement documents, request for proposals (RFPs) for systems, what they talked about was measuring what is the uptime of the system? Can people log into the system? Does it show up? As opposed to the truly functional requirements, which are, does it reduce police shootings? Does it prevent those shootings? Does it help save people's lives?

What is happening is that the focus has been on the mechanics, because that is easy to do. It is easy to measure. It does not require that much thinking and effort, and so we do the easy thing, and we forget the hard stuff.

There are many such other examples where we have contracts that we get stuck into that are unnecessarily long term. We do not allow the systems to get data out and put more data in. We do not have them interoperate, because as we talked about, AI systems do not work in isolation, they are connected to different pieces.

There is a lack of customization and configuration. I think there is a whole set of things. What we need to do is create a much more holistic procurement process that has both requirements around what did you designed the system to do, how did you validate that it did what you wanted it to do, and how are you developing a continuous monitoring process that is continuing to do so? Most of those things, there are no standards that exist today for that procurement, and we need to create those.

Chairman PETERS. A follow up on that. Many others were nodding their head. This is important. What those requirements are. The other question is, can we standardize those requirements across a variety of government agencies, or does it have to be more niche? You want to take the first stab at that.

Then I see Ms. Raj, you are nodding your head, too, so your thoughts on that. Then any others want to jump in. But, Mr. Ghani.

Mr. GHANI. I think it is going to have to be an 80–20 thing here, where there is a series of things that we can standardize, and what we can standardize on is what do we ask for. We can standardize if we need to figure out what values this system is built on.

We need to figure out how you built it. What design choices did you build? What artifacts were produced? Does it work for people? Who does it work for? Who does it not work for? How is it validated? Those are kind of high level questions that we can standardize that need to be there.

Data in, data out, interoperability, configuration. What we cannot standardize on is what specific values it should have. That needs to come from the use case. That is where that collaboration is going to happen, is people with expertise in understanding the policy issue.

We are talking about service delivery. What is the goal of that service? Is it to improve people's lives or is to save money? When there is a tradeoff, who decides that? The questions that we were talking about earlier, those are the things that are not going to be standard. They are going to depend on the specific use case and specific policy and the specific service.

But everything, the process that was used to design that system, to come up with those values, and to validate it, and all the other things, those can be standardized. Again, I think it is going to be

an 80–20 thing, where 80 percent can be standardized, and 20 percent will need to be customized.

Chairman PETERS. Right. If anyone else has thoughts. But Ms. Raj, you have some thoughts.

Ms. RAJ. Yes. Thank you, Senator Peters, for your question. At a high level, I think we can also think about it from the other way, which is what data is available for AI to tackle from a low hanging fruit perspective, right. I think there is a way to organize, hey, this is the data that is available that potentially can be used for automation.

This is the level of the responsible AI that can be applied to these particular questions. Perhaps getting AI slowly integrated via this is the data available, and these are the questions that there are more guardrails around, that could be a good way to start the standardization process. Because starting standards without actually tying it to specific use cases and mission need, then you will have that misalignment.

Chairman PETERS. Mr. Roberts.

Mr. ROBERTS. Yes, Senator. I will piggyback on something Professor Ghani said, too, about obsessing over the mechanics in terms of valuing the performance. This is why it is so important for acquisition professionals to be mission focused and look at AI as an enhancement of the mission.

Because it is so easy to, as Professor Ghani mentioned, to measure performance based on mechanics, based on how is it working rather than is it working, is it actually valuable to the end user? When you have a focus on that, I think you will find that the acquisition team changes the way they approach even risks and responsible use of AI.

They change the way they look at the intellectual property. It all focuses on mission value and value to end user, which trickles down into the way we look at everything.

Chairman PETERS. Very good. Mr. Ghani, back to you. We have heard previously about the need to audit AI systems to account for drift and unintended consequences. The question for you is, what procedures should be put in place to audit the government and these AI systems within the government, and should this need be accounted for upfront in the procurement process? And if so, how?

Mr. GHANI. If so, absolutely, yes. I think the audit has to be there. I think there are, in my mind there are three stages of this audit, right. The first audit, and I think I keep going back to and we are all sort of saying the same word, values, right.

When we are designing a system, the procurement has to ask for a system to help achieve certain values. We need to audit those and validate whether those are the values we should care about. That is the first audit.

That is not a technical audit, that is a values audit. Two, when a system is being procured, we need to audit how did the vendor, consultant, researcher build this system to help it achieve those values? That is a technical audit. Three, once it is deployed, it is not going to work in isolation.

In most cases it is informing, especially high stakes decision, it is informing humans. You can audit the system for what it outputs, but we need to audit, how does it interact with this person, and

how does this human decision change? Because that is eventually what we care about, is this impact on people.

We need to audit the interaction between the AI system and the human system, and then audit the outcomes that it produces. So those are the four pieces. It is not a one-off. It is a continuous thing because it is going to change.

Chairman PETERS. Good. Mr. Roberts, how can the government ensure that the data that is used for testing and training the Government's AI platforms is actually secure and protected? Anyone else can jump in on these too.

I will pick out one individual, but feel free if you want to say anything. Raise your hand as we wrap up. But only a couple of more questions and then you will be free. Mr. Roberts.

Mr. ROBERTS. Yes, Senator. I will start by saying that the more we restrict the free flow of data and access from contractors to the data, there is also a sense in which that becomes more problematic as well, especially for the functioning of the model.

We have seen instances just on the other side of over-classification, over-regulation, over-protection of data that has killed projects. However, having said that, data security, protection, and privacy, is essential, and especially in areas that I have worked with, with the classified data that affects national security and personally identifiable information (PII), especially with health records.

We have seen some things that have helped, especially with health records. The use of synthetic data was beneficial to us that we were able to use. There are other sources of data. But not to oversimplify, I think the most important thing for the acquisition field is to put the ethics professional, the person who is dealing with privacy, the security professional, into the planning phase.

Again, having a balanced team that has all these professionals involved at the very beginning. Supply chain risk is another big problem with security. It is something that is not looked at much. We are finding with AI, because a lot of these rules such as supply chain risk were always rules, but they are reemerging in much more important ways when we are looking at the risks, the adversarial threats.

It is looking at all these risks in new ways, and it is making sure you have a full, balanced team to be with you at the planning stage, on how to deal with it.

Chairman PETERS. Right. Ms. Raj.

Ms. RAJ. Yes. I want to talk about it from the perspective of small business. CrowdAI has worked with the U.S. Government on AI initiatives on a wide range of sensors, all unclassified. As we move forward and mature, we started working with more with the Defense Counterintelligence and Security Agency (DCSA) to make sure that we could be ready for other types of data, more sensitive data.

Many systems that we worked with, with U.S. Government data, were either in bare metal servers or data that remained in government clouds. I believe that it is important that if a company that is in dual use wants to work with the U.S. Government, they also need to ensure that the data is treated with responsibility and privacy to the maximum extent possible.

I think that as companies start putting their technology in a more dual use manner, they also need to comply with privacy and regulation that is so standard across a lot of large companies.

AI is the ever evolving technology, and so the way you make sure that companies of all sizes continue this type of evaluation around privacy and ethics is making sure that there is qualitative and quantitative testing, because often aggregated statistics may not paint the full picture.

Chairman PETERS. Very good. I would like to thank our witnesses for being here today. I am certainly grateful to your contributions. This is a very important discussion, and it does not end here. We are going to have many more discussions going forward. We hope all of you are available to help this Committee work on this important issue.

Certainly, as we heard today, the use of automated systems to help the government provide public services more efficiently is nothing new. We have been dealing with this for a long time.

Mr. Roberts, you have been dealing with it for a long time, as well as everybody on the panel. However, as we enter this age of rapid development of advanced machine learning models and other forms of artificial intelligence, now is the time to ensure that the algorithmic systems that the government buys do not have unintended or harmful consequences.

I think, as each of our witnesses have emphasized, enacting appropriate guardrails and oversight policies for the procurement of AI in Government will shape its development and use across all industry, and industries in the years to come.

Americans deserve a government that is modern, that is efficient, and innovative, as well as one that is transparent, fair, trustworthy, and protects their privacy. As Chair of this Committee, I will continue to work to ensure that government lives up to these principles and that promise.

Your testimony will help inform the Committee's future and legislative activities going forward. Again, we hope this is an ongoing dialog in a very fast moving and challenging area, but essential for us to understand and act appropriately.

The record for this hearing will remain open for 15 days until 5.00 p.m. on September 29, 2023 for the submission of statements and questions for the record. This hearing is now adjourned.

[Whereupon, at 12:01 p.m., the hearing was adjourned.]

APPENDIX

Chairman Peters Opening Statement As Prepared for Delivery
Full Committee Hearing: Governing AI Through Acquisition and Procurement.
September 14, 2023

Today's hearing is the third in a series I have convened on artificial intelligence. At our first hearing in March, we discussed the transformative potential of AI, as well as the possible risks these technologies can pose. At our second hearing in May, we considered the role of AI in government – how AI tools can improve the delivery of services to the American people, and how to ensure they are being used responsibly and effectively.

Today, we are doing a deeper dive into how government will purchase AI technologies, and how the standards and guardrails government sets for these tools will shape their development and use across all industries.

The federal government is already using AI, and its use across agencies is only expected to grow in the coming years. These systems can help provide more efficient services, assess potential security threats, and automate routine tasks to enhance the federal workforce. For example, the Department of Homeland Security is using natural language processing to evaluate employee surveys and improve workplace experience. And the Federal Aviation Administration deploys machine learning to update the weather models that help land planes successfully.

Other technologies that are continuing to develop, such as generative AI, offer the potential to improve government services even more. For example, many agencies – from the Office of Personnel Management, to the Department of Health and Human Services, to the Department of

Education – have rolled out chatbots to provide better service to federal employees and the larger American public. AI is here, and it's already being put to good use.

Many of these systems are not developed by the government, but rather in the private sector. Over half of the AI tools used by federal agencies have been purchased from commercial vendors. This collaboration between the public and private sector is crucial – it ensures that the government is using the most effective AI systems. American companies are breaking new ground with these technologies, and we have the chance to share in the benefits of that innovation.

But these tools also bring potential risks and policy implications. They require new knowledge from procurement officials, as well as increased coordination across agencies.

In order to successfully and effectively purchase and use AI tools, federal agencies have to be prepared to address issues like privacy concerns about the use of federal data to train commercial models and bias in government decision-making. We must be nimble *whenever* the government collaborates with the private sector – but this is especially true with AI, where new developments emerge almost every single day, and tools that are purchased are often actively learning and changing as they are used.

Last Congress, I authored and enacted a law that requires officials that procure AI tools to be trained in their capabilities and potential risks. This year, I introduced legislation that would extend this training to all federal managers and supervisors. I have also introduced legislation

that would designate a Chief AI Officer at every federal agency, so that we have the leadership and expertise in place to maximize the potential of these technologies and effectively address the risks.

These guardrails are more important than ever. Federal agencies are inundated with sales pitches and technology demos promising the next big thing. While the federal government must be forward thinking, we also have to be cautious in procuring these new tools. And we must continue that work past the initial purchase – testing and fine-tuning our models to ensure that they are effectively serving the American public.

And as AI development accelerates, private industry has yet to standardize practices for evaluating AI systems for risk, trustworthiness, and responsibility. Through federal procurement policy, the government has a unique opportunity to shape standards and frameworks for development and deployment of these technologies across the private sector more broadly.

I look forward to hearing from our expert witnesses today, and to continuing our bipartisan work to help encourage American development of AI and ensure it is being used appropriately.

RANKING MEMBER PAUL OPENING REMARKS

Committee on Homeland Security and Governmental Affairs

September 14, 2023

In 2021, the Pentagon, through the Defense Advanced Research Projects Agency (DARPA), asked for proposals for "real-time, comprehensive tools that establish ground truth for how countries are conducting domestic information control."

DARPA's goal in developing AI technology for "Measuring the Information Control Environment" was to help the U.S. government better understand "how digitally authoritarian regimes repress their populations at scale over the internet via censorship, blocking, or throttling."

Of course, the solicitation made it clear that the Pentagon did not want the proposals to look at the activities of the United States government. The Pentagon, and the U.S. Government as a whole, enjoy professing moral superiority over authoritarian governments when it comes to upholding basic democratic values.

American politicians have no qualms about criticizing foreign governments like Russia and China for their suppression of civil liberties and efforts to eliminate dissent. Yet, there seems to be a complete unwillingness to have an honest conversation about the disturbingly similar actions our own government is actively engaged in and financing.

For decades, the Pentagon and other federal agencies have been quietly partnering with private organizations to develop powerful surveillance and intervention tools designed to monitor and influence narratives on social media.

For example, a 2021 Pentagon program called Civil Sanctuary sought to develop AI tools to scale the moderation capability of social media platforms to create what it describes as a "more stable information environment." In other words, the goal of this Pentagon program was to exponentially multiply the government's ability to coordinate censorship of online speech.

The Pentagon has invested millions of taxpayer dollars to develop these tools not only for use by the social media companies, but also the Intelligence Community and law enforcement.

Meanwhile, the Department of Commerce is awarding million-dollar grants for cognitive research into how the U.S. Government can foster trust in AI with the general public.

So, while the federal government is using taxpayer dollars to develop AI to surveil and monitor Americans' online speech, it is also spending money to figure out how to get you to trust the AI.

Over the last year, starting with the Twitter Files, journalists started to expose the deep coordination between the federal government and social media platforms when it comes to content moderation decisions and policing the speech of Americans.

As Michael Shellenberger rightly points out, the threat to our civil liberties comes not from AI but from the people who want to control it and use it to censor information.

Just last week, the Fifth Circuit affirmed the government likely violated the First Amendment by coercing social media companies to remove speech the government disagreed with related to the origins of COVID-19, pandemic lockdowns, vaccine efficacy, and the Hunter Biden laptop story.

The Court cited numerous examples of U.S. government officials engaging in domestic information control on social media. Government officials demanded that platforms implement stronger COVID misinformation monitoring programs, modify their algorithms to avoid amplifying misinformation, target repeat offenders, and magnify communications from certain trusted sources.

After one meeting with federal officials, one platform committed to reducing the visibility of information skeptical of the government's COVID vaccine policy even when it does not contain actionable misinformation. Facebook promised to label and demote a popular video after officials flagged it even though they acknowledged it did not qualify for removal under its policies.

I fear that we are likely in only the beginning stages of understanding the extent of the federal government's involvement in the content moderation decisions of private social media platforms.

What we do know is that our government is funding the development of powerful artificial intelligence tools for monitoring and shaping the online discourse.
Now, I want to be clear. AI is not inherently malicious. It has the potential to revolutionize basic aspects of society, from healthcare to education.

However, in the hands of an unchecked government, AI can be weaponized as a tool to suppress the fundamental values our country was founded upon – the open exchange of ideas, the freedom to question, and right to dissent. This should not be a partisan issue.

GOVERNING AI THROUGH ACQUISITION AND PROCUREMENT

Testimony by
Rayid Ghani,
Distinguished Career Professor
Machine Learning Department and the Heinz College of Information Systems and Public Policy
Carnegie Mellon University

Before the
United States Senate Committee on Homeland Security and Governmental Affairs Hearing on "Governing AI Through Acquisition and Procurement"

Thursday, September 14, 2023

Chairman Peters, Ranking Member Paul, Members of the Committee, thank you for hosting this important hearing today, and for giving me the opportunity to submit this testimony.

My name is Rayid Ghani and I am a Distinguished Career Professor in the Machine Learning Department and the Heinz College of Information Systems and Public Policy at Carnegie Mellon University. I've worked in the private sector, in academia, and extensively with government agencies and non-profits in the US and globally on developing and using Machine Learning and AI systems to tackle social and public policy problems across health, criminal justice, education, public safety, human services, and workforce development in a fair and equitable manner.

The promise of AI in helping build a better society

Artificial Intelligence has enormous potential in helping tackle critical problems we face in society today, ranging from improving the health of our children by reducing their risk of lead

poisoning[1], to reducing recidivism rates for people in need of mental health services[2], to improving educational outcomes for students at risk of not graduating from school on time[3,4], to improving police-community relations by identifying officers at risk of adverse incidents[5], to supporting proactive inspections to improve health and safety conditions in workplaces[6] and in rental housing[7], to improving healthcare practices[8], to designing more effective organ exchange systems[9]. There is tremendous potential for every federal agency to use AI - in helping them design, implement, and evaluate their programs to help improve outcomes for everyone and result in a better and more equitable society.

However, any AI system affecting people's lives has to be explicitly designed to focus on increasing equity and promoting our societal values, and not just narrowly optimizing for efficiency. AI can have a massive, positive social impact but we need to make sure that we put guidelines in place to maximize the chances of that positive impact. If not designed, deployed, and used appropriately, it can risk harm to people who have been traditionally marginalized in society. An AI system, designed to explicitly optimize for efficiency, has the potential to result in leaving "more difficult or costly to help" people behind, resulting in an increase in inequities. It is critical for government agencies and policymakers to ensure that AI systems are designed, developed, and used in a responsible manner to ensure that they result in supporting equitable outcomes for everyone. In a policy brief published by the Responsible AI Initiative at the BlockCenter at Carnegie Mellon University, we highlighted some of the unique challenges of AI Accountability and lay out a set of policy recommendations[10].

Scoping and Procurement of AI systems needs to be a focus area for policymakers

While the entire lifecycle of AI systems - scoping, procurement, designing, testing, deploying, and using needs to have guidelines and best practices in place that maximize the societal benefit and minimize potential harms (such as the efforts around the AI Risk Management Framework[11]

---

[1] Predictive Modeling for Public Health: Preventing Childhood Lead Poisoning. Potash et al. Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2015)
[2] Reducing Incarceration through Prioritized Interventions. Bauman et al.. ACM SIGCAS Conference on Computing and Sustainable Societies, 2018.
[3] A Machine Learning Framework to Identify Students at Risk of Adverse Academic Outcomes. Lakkaraju et al. Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining
[4]http://www.dssgfellowship.org/project/identifying-factors-driving-school-dropout-and-improving-the-impact-of-social-programs-in-el-salvador/
[5] Early Intervention Systems – Predicting Adverse Interactions Between Police and the Public. Helsby et al. Criminal Justice Policy Review, 2017.
[6]http://www.dssgfellowship.org/project/improving-workplace-safety-through-proactive-inspections
[7] http://www.datasciencepublicpolicy.org/projects/public-safety/san-jose-housing/
[8] Kilic A, Dochtermann D, Padman R, Miller JK, Dubrawski A (2021). Using machine learning to improve risk prediction in durable left ventricular assist devices. PLOS ONE 16(3).
[9]https://aaai.org/tuomas-sandholm-wins-2023-aaai-award-for-artificial-intelligence-for-the-benefit-of-humanity/#:~:text=This%20year%2C%20the%20AAAI%20Awards,on%20both%20practice%20and%20policy.
[10] https://www.cmu.edu/block-center/responsible-ai/index.html
[11] https://www.nist.gov/itl/ai-risk-management-framework

being developed by NIST), there has been a lack of attention to the earlier phases of this process, specifically scoping and procurement. Many of the AI systems being used in federal, state, and local agencies are not built in-house but procured through vendors, consultants, and researchers. This makes getting the procurement phase correct critical - many costly problems and harms discovered downstream can be avoided by a more effective and robust procurement process.

We need to ensure that government procurement of AI follows a "responsible" process, and in turn requires AI vendors to follow a "responsible" process in designing such systems, and results in the selection, deployment, and use of a system that promotes accountability and transparency, and leads towards equitable outcomes for those impacted.

Procuring solutions to specific problems rather than procuring "AI"

Too often, organizations go on the market to buy "AI" without completely understanding, defining, and scoping the concrete problem they want to tackle, without assessing whether AI should even be part of the solution, and without including individuals and communities that will be affected. AI systems are neither applicable for all problems facing government agencies, nor are they one-size-fits-all. By starting with the concrete problem at hand, and understanding how it's being tackled today, an effective, collaborative, and inclusive scoping process can help determine the requirements that the AI system needs to fulfill. For example, consider procuring a system to support matching unemployed individuals with training or skilling programs that they can be enrolled in to get them back into employment most effectively. The requirements for such a system are not primarily AI requirements, but rather come from the intended goals of the program being administered - whether this system will result in increased employment rates for the individuals impacted, whether it will propagate existing disparities in training and employment outcomes, whether it will enable and empower employment agencies, unemployment counselors, and unemployed individuals, among other requirements that will be surfaced through the scoping process, such as the one we've previously developed[12].

AI systems optimize for what their developers tell them to optimize for (and the procurement process needs to tell the vendors what to optimize for)

AI algorithms are neither inherently biased nor unbiased (in the societal sense) or have inherent, fixed, "values". The AI developers designing and building the AI system (implicitly or explicitly) make hundreds of design choices that result in the eventual system and its behavior. If the developers make design choices that explicitly focus on societal outcomes we care about and evaluate their systems against those intended outcomes, the AI system can help achieve what we want it to achieve. The procurement process needs to define these goals and values, and ensure that the vendors address those appropriately in the system being procured.

[12] http://www.datasciencepublicpolicy.org/our-work/tools-guides/data-science-project-scoping-guide/

AI is forcing us to make societal (and public policy) values explicit (and the procurement process needs to define what those values should be)

Because an AI system requires us to define exactly 1) what we want to optimize it for, 2) which mistakes are costlier (financially or socially) than others, and 3) by how much, it forces us to make these ethical and societal values explicit. These values are implied in any decision-making process, including all the human decision-making processes that exist today, but are not necessarily made explicit. These implicit values (coded through human decisions) when biased and unfair, result in inequitable outcomes.

For an AI system to be built, these values need to be provided as a critical input[13]. For example, for a system that is recommending lending decisions, we may have to 1) specify the differential costs of highlighting someone as unlikely to pay back a loan and being wrong about it versus predicting that someone will pay back a loan and being wrong about it, and 2) specify those costs explicitly in the case of people who may be from different gender, race, income, or education level groups. While that may have happened implicitly in the past and with high levels of variation across different human decision makers (loan officers in this case), with AI-assisted decision-making processes, we are forced to define them explicitly.

One key question the procurement process has to answer is who and how we should come up with these sets of values for a given problem setting, what information to ask for around these values, and how to evaluate the correctness of the values, and the fidelity of the procured and designed system to these values. Unfortunately, today, these decisions are too often left essentially by default to the AI system developer or an arbitrary set of individuals who define those values in an AI algorithm (explicitly or implicitly). The recommendations at the end of this testimony go into more detail on what I recommend should be done but it certainly should not be left to the AI system developer making those choices alone; the team and process should include all stakeholders including policymakers and the community being impacted by this system.

What does it take to create responsible AI systems for society?

The following steps need to be taken to create Responsible AI systems for society and the procurement process needs to set expectations and accountability for vendors in each of these steps:

1. **Defining** the goals and policy and societal outcomes the system needs to help achieve(which includes the societal values and a collaborative, multi-stakeholder process).

---

[13] From Preference Elicitation to Participatory ML: A Critical Survey & Guidelines for Future Research M. Feffer, M. Skirpan, Z. Lipton*, and H. Heidari. The AAAI /ACM Conference on Artificial Intelligence, Ethics, and Society (AIES), 2023.

2. **Translating/Mapping** those desired outcomes and values into analytical and technical requirements that the vendors should design the AI system to achieve.
3. **Building** an AI system that fulfills those analytical requirements and releasing documentation and additional artifacts enumerating all the design choices (including around the choice and use of data, the AI algorithms, and the downstream use and impact), demonstrating and providing evidence of how it was built to achieve those goals.
4. **Validating** through a trial (and providing evidence) that the AI system did, in fact, fulfill those requirements and achieve the initial outcomes defined in step 1 before deploying the system.
5. **Continuous Monitoring & Evaluation** of the entire system (the AI system followed by human decisions) during its lifetime to ensure that it continues to achieve equitable outcomes from step 1.

Moving Forward to Governments Procuring and Using AI Systems that Result in a More Equitable Society: Our Recommendations

It is critical and urgent for policymakers to act and provide guidelines and regulations for both the public and private sector organizations procuring, developing, and using AI in order to ensure that these systems are built in a transparent and accountable manner and result in fair and equitable outcomes for society. As initial steps, we recommend:

1. Focused Procurement for Specific Use-Cases

We need to ensure that AI systems are procured for specific use-cases, and to support intended outcomes around that use case, rather than as generic AI systems. This is intended to both promote better outcomes as well as to prevent harm through misuse. An AI system that yields beneficial and equitable outcomes in one context might yield just the opposite in another. While AI algorithms across different areas have a lot in common, developing a generic and complete framework for AI that works well across all possible uses is likely to be an unrealistic proposal. Rather, the need for application-grounded procurement processes is important to achieving policy and societal goals across different government agencies.

2. Development of common procurement requirements and templates

While this may seem contradictory to the previous recommendation, the documentation and artifacts needed to assess the appropriateness and effectiveness of an AI system are common across many use cases. The specifics of the use case define the concrete values, the goals, and the evaluation criteria, and the common procurement requirements and templates are used to assess how well the system is able to achieve them. These common procurement and RFP templates should include the set of artifacts that should be provided during the evaluation of the AI software. This includes information on:

1. How the system was built and what it was designed to optimize for
2. What tests were run to check if it did do what it was intended to do?
3. What types of people was it effective for? Who does it fail for?
4. How long was it in trials for, when, and how did the effectiveness change over time?
5. What risks need to be considered and what are the mitigation plans for each of these risks?
6. Any extended data collection process and infrastructure that may need to be set up to collect additional data attributes (such as race, gender, or income) that may not already be collected but are necessary to measure equity outcomes
7. How to set up evaluation standards to compare the performance of these systems to the human decision-making processes (if any) currently being used.
8. How the vendor supports explainability and interpretability of the AI systems in order to provide recourse to individuals who may be adversely impacted by the decisions made using the system.
9. A continuous improvement plan to ensure that the system continues to not only be evaluated but also improved upon to achieve the desired outcomes.

RFPs for AI systems should include an explicit initial project phase to gather requirements for the values and goals of the system. This process should include a diverse team and work with stakeholders including: developers who build and deploy AI systems, decision-makers who implement the systems in their workflows, and the community being impacted by these systems.

Ideally this should be put in place for any process involving decision making of any kind, whether human decisions or AI-assisted decisions but becomes critical in cases where the scale of deployed AI systems increases the risk. This is not an exhaustive set of questions and will need to vary based on the problem being addressed and the impact this system can have on people's lives.

3. Community Participation

Create guidelines that ensure **meaningful involvement of the communities that will be impacted** by the AI systems right from the inception stage. Engage in continuous dialogue and feedback to understand their concerns, values, and suggestions which should guide the design of RFPs, and of the design, deployment, and use of the AI systems.

4. Create Trainings, Processes, and Tools to Support Procurement Teams

As the procurement teams expand their role and start procuring AI-augmented systems, they will need to be supported by increasing their capacity to fulfill this role. We recommend creating trainings, processes, collaboration mechanisms, and tools to help them:

1. Understand where existing processes may and may not be well-adapted to systems using AI.
2. Understand and define what process and outcomes standards to set.
3. Understand how to evaluate whether the requirements created for an AI system were in fact aligned with the identified societal equitable outcomes.
4. Understand how to evaluate whether the AI system did in fact do what it was designed to do.
5. Develop a continuous monitoring and audit process and tools.
6. Create standards for when a system should "expire" and a corresponding renewal process.
7. Create technical software tools to support the end-to-end procurement process - help with scoping the need, to write RFPs, to identify issues with RFPs, to analyze responses, to conducting technical evaluations of AI software from vendors.
8. Avoid common pitfalls that result in costly downstream impact such as:
    a. Locking into unnecessary long-term contracts
    b. Inability for the underlying data systems to ingest new or external data,
    c. Inability to export data into other systems for further linkage and policy analysis
    d. Lack of interoperability with commonly used systems within and across federal, state, and local government agencies
    e. Lack of customization and configuration based on changing needs of the use case
    f. Hidden financial costs that may be involved in various processes such as for scaling the system, or in customizing it.
    g. Over-focus on trivial, software metrics such as up-time, at the expense of use-case focused metrics around effectiveness or equity.

These steps need to take a phased approach and be iterative:

- Near term: Getting started by partnering with external organizations such as universities to help create these guidelines and to provide training to government agencies
- Medium term: Collaboratively developing more rigorous procurement processes and tools to support the agencies in the medium term
- Longer Term: Providing the agencies with the resources to expand their internal capacity.

The overall goal behind these recommendations is to set some standards around procurement of AI by government agencies and to support and enable the agencies to implement those standards effectively and procure AI systems that help us achieve their policy and societal goals.

Stanford University
Human-Centered
Artificial Intelligence

Governing AI Through Acquisition and Procurement

Fei-Fei Li[1]
Stanford Institute for Human-Centered Artificial Intelligence (HAI), Stanford University

Testimony presented to the U.S. Senate Committee on Homeland Security and Governmental Affairs on September 14, 2023.

I. Introduction

Chairman Peters, Ranking Member Paul, Members of the Committee, thank you for the privilege of appearing before this prestigious body. As a proud American, few things could make me more honored than offering service to our government leaders.

I have spent my life working in the field of artificial intelligence (AI) with over 25 years studying, developing, and understanding the technology that has just entered the public's consciousness due to recent breakthroughs. I have approached the study and development of AI not only as a pioneer, academic scientist, and teacher, but also as the child of two parents with chronic health issues.

Right now you are hearing wild claims about AI from two ends of the spectrum. Some, propelled by industry hype, have said that AI will fix all of the challenges humanity faces. At the other end of the spectrum, some have claimed AI will lead to the end of the world through a biological or nuclear catastrophe. As a proud, self-proclaimed, "nerd," I, too, love a good science fiction story. However, one goal of my testimony today is to demystify some of these wild claims for you.

We are indeed at a moment where the right investments in AI could fundamentally transform the human condition for the better, and if we are not careful those same investments could exacerbate some of the worst parts of human nature such as authoritarianism, racism, and crime. So let me tell you what AI can do at this time, where we will likely see progress in the near future, and how the federal government can steward the development of AI in a way that allows us to realize its benefits and mitigate its harms.

II. Benefits of AI

We have arrived at an inflection point in the world of AI, largely propelled by the breakthroughs in generative AI, including increasingly sophisticated language models like GPT-4. These

[1] Sequoia Capital Professor in Computer Science, Stanford University; Denning Co-Director and Senior Fellow, Stanford Institute for Human-Centered Artificial Intelligence (HAI), Stanford University; Co-Founder and Board Chair, AI4ALL; Member of the National Academy of Engineering, the National Academy of Medicine and American Academy of Arts and Sciences. All views expressed in this testimony are provided in an individual capacity and do not represent the views of any affiliated organization.

models have revolutionized various sectors from customer service to adaptive learning. However, the scope of intelligence is far broader than linguistic capability alone. In my specialized field of computer vision, we have also witnessed remarkable advancements that empower machines to analyze and act upon visual information—essentially teaching computers to 'see.'

I'd like to particularly highlight today how AI advancement is augmenting the capabilities of caretakers and medical professionals. Healthcare is a domain that I have dedicated my entire career to and I am honored to contribute to the field as a member of the National Academies of Medicine. Earlier, I mentioned that I am the primary caregiver for my aging parents. This personal experience has provided me with invaluable insights into the challenges that caregivers face daily and how AI can transform the healthcare landscape. For example, algorithms can detect anomalies in medical imagery such as X-rays and MRI scans, thereby aiding early diagnosis and treatment.[2] AI-enabled documentation assistance can reduce the administrative burden on healthcare professionals but also minimizes the risk of errors in patient records, thereby improving patient safety and care quality.[3]

Most importantly, such an innovation presents many benefits to the U.S. government.

First, government entities have long been consumers and stewards of public-use technology to streamline the efficiency of governance. Nearly half of all federal agencies have experimented with AI and related machine learning (ML) tools.[4] As technology continues to advance at an unprecedented pace, it is imperative that our government remains at the forefront of innovation to better service delivery to its citizens and address complex societal challenges more effectively. When it comes to healthcare, for example, the Department of Health and Human Services has initiated a pilot program aimed at combating Medicare fraud.[5] This program employs AI-based models to enhance the efficiency of fraud detection within the Centers for Medicare & Medicaid Services (CMS), which handles over one million transactions daily.

Second, the use of AI in healthcare not only improves the quality of life for the elderly population, but also reduces healthcare costs.[6] Medical AI tools can reduce the frequency of emergency medical interventions and hospital readmissions, which are significant cost drivers in healthcare expenditures. By facilitating proactive, preventive care, the technology can mitigate the financial burden on public healthcare resources, including Medicare and Medicaid, and

---

[2] Nikki Goth Itoi. 2022. Could Stable Diffusion Solve a Gap in Medical Imaging Data? Stanford Institute for Human-Centered AI. https://hai.stanford.edu/news/could-stable-diffusion-solve-gap-medical-imaging-data.
[3] Grace Hong et al. 2020. Clinicians' Experiences with EHR Documentation and Attitudes Toward AI-Assisted Documentation. Stanford University School of Medicine and Google Health. https://med.stanford.edu/content/dam/sm/healthcare-ai/images/Stanford-Google_AI-Scribe_WhitePaper.pdf.
[4] David Freeman Engstrom, Daniel E. Ho, Catherine Sharkey, and Mariano-Florentino Cuéllar. 2020. "Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies." Administrative Conference of the United States.
[5] Nihal Krishan. 2023. HHS CIO Mathias Says Tree-based AI Models Helping to Combat Medicare fraud. Fedscoop. https://fedscoop.com/hhs-cio-mathias-says-tree-based-ai-models-helping-to-combat-medicare-fraud/.
[6] Milstein A. Haque and L Fei-Fei. 2020. Illuminating the Dark Spaces of Healthcare with Ambient Intelligence. *Nature* 585. https://www.nature.com/articles/s41586-020-2669-y. 193–202.

facilitate more efficient allocation of resources in healthcare settings, enabling medical professionals to focus on other critical areas of public health and welfare that require attention.

III. Harms and Unintended Consequences

However, while AI, like most technologies, promises to solve many problems for the common good, it can also be misused to cause harm and carry unintended consequences. The very same technology that could potentially save countless lives from hospital infections could one day be repurposed as an active form of surveillance against people. I know Senator Paul and others on this committee's longstanding concerns related to surveillance and I, too, appreciate these very same concerns as an AI developer.

Let me give you two examples of when the harms of AI could affect how the government approaches AI. First, bias in AI is well-documented.[7] But that's not the only problem that can lead to harm to marginalized communities. Take credit scoring as an example. As credit risk scoring tools increasingly leverage AI, research shows that predictive tools used to approve or reject loans are less accurate for low-income, minority groups in the United States due to the lack of data in their credit histories.[8] To ensure that AI applications deliver reliable results for all Americans, we must ensure the availability of high-quality, representative data sets.

Second, healthcare AI also presents considerable challenges to privacy and data security. In an era of heightened public concern over data collection and misuse, it is vital that we build strong privacy and security protocols into these applications from the beginning. Achieving this necessitates a multidisciplinary effort that engages experts across various fields. Developers, policymakers, healthcare providers, and patients should all be proactively involved throughout the entire development and implementation phases to ensure the AI tools are both effective and secure.[9]

This is why we must ensure a diverse ecosystem in the development of AI and why I helped found Stanford's Institute for Human-Centered Artificial Intelligence (HAI), where we study AI and its impact not as a field exclusive to computer science, but instead as a multidisciplinary field that includes the social sciences, engineering, law, medicine, and the humanities. The federal government should adopt a similar approach to properly understand the future of AI.

---

[7] See David Danks and Alex John London. 2017. "Algorithmic Bias in Autonomous Systems," *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. https://doi.org/10.24963/ijcai.2017/654; Joy Buolamwini and Timnit Gebru. 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," *Proceedings of Machine Learning Research* 81. https://doi.org/https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf; Yolande Strengers, Lizhen Qu, Qiongkai Xu, and Jarrod Knibbe. 2020. Adhering, Steering, and Queering: Treatment of Gender in Natural Language Generation. *In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM. https://doi.org/10.1145/3313831.3376315.

[8] Laura Blattner and Scott Nelson. 2021. How Costly is Noise? Data and Disparities in Consumer Credit. arXiv. https://arxiv.org/abs/2105.07554.

[9] Jenna Wiens et al. 2019. "Do No Harm: A Roadmap for Responsible machine learning for health care." *Nature Medicine* 25:9: 1337-1340. https://pubmed.ncbi.nlm.nih.gov/31427808/.

IV. What Can the U.S. Government Do

The adoption of AI introduces unique complexities and issues, not only for the public bodies acquiring these technologies but also for the communities they impact. It falls upon the U.S. government to spearhead the ethical procurement and deployment of these systems, with the aims of both safeguarding the rights of individuals and communities, as well as encouraging industry innovation through responsible AI guidelines. Indeed, responsible federal acquisitions and procurement have the true potential to set the norms for AI development and ultimately shape the field of responsible AI in a more immediate and direct way than any future regulation that may or may not come from this Congress.

I applaud this committee and the work that it has done thus far on AI, including the AI Training Act and the AI Lead Act, which create powerful tools for the federal government to set such norms.[10] The AI Training Act can up-skill procurement officials and equip them with a nuanced understanding of AI capabilities and limitations. I am proud that Stanford HAI has tailored a government education program to fulfill the Training Act requirement specifically for the U.S. General Services Administration, in partnership with the Office of Management and Budget.

As the U.S. government's spending on AI-related contracts has surged, it's more crucial than ever to closely examine these vendors to ensure their goals align with those of the federal government.[11] One key component is establishing accountability and transparency requirements. Vendors should disclose key information about their systems, including reporting how their AI systems work; how their dataset was designed, collected, and annotated; what potential risks their systems pose; and what their strategy is to ensure their models can be evaluated against federal government standards.[12]

V. U.S. Public Sector Investment

For the reasons I have just outlined, it is imperative that the federal government makes the needed critical investments in AI. While the numerous benefits of AI such as efficiency and productivity gains in government that allow it to be more responsive to its citizens make AI attractive, this future is not foretold and it may not be dominated by the United States.

---

[10] See S.2551 - 117th Congress (2021-2022): AI Training Act, S.2551, 117th Cong. 2022. https://www.congress.gov/bill/117th-congress/senate-bill/2551/summary/55; S.2293 - 118th Congress (2023-2024): AI LEAD Act, S.2293, 118th Cong. 2023. https://www.congress.gov/bill/118th-congress/senate-bill/2293.

[11] See Nestor Maslej et al. 2023. "The AI Index 2023 Annual Report." AI Index Steering Committee, Stanford Institute for Human-Centered AI.

[12] See Emanuel Moss et al. 2021. Assembling Accountability: Algorithmic Impact Assessment for the Public Interest. Data & Society Report. https://datasociety.net/library/assembling-accountability-algorithmic-impact-assessment-for-the-public-interest/; Margaret Mitchell et al. 2018. Model Cards for Model Reporting. arXiv. https://arxiv.org/abs/1810.03993.; Rishi Bommasani, et al. 2023. Ecosystem Graphs: The Social Footprint of Foundation Models. arXiv. https://arxiv.org/abs/2303.15772; Rishi Bommasani, Daniel Zhang, Tony Lee, Percy Liang. 2023. Improving Transparency in AI Language Models: A holistic evaluation. Foundation Model Issue Brief Series. Stanford Institute for Human-Centered AI. https://hai.stanford.edu/foundation-model-issue-brief-series.

That is why with its democratic values, commitment to the rule of law, and spirit of innovation and entrepreneurship, America must lead in AI. And while U.S. industry currently enjoys unique advantages in AI, there is a deep imbalance in the U.S. AI innovation ecosystem that hinders leadership in the field.

Because of the vast amounts of compute and data required to train these systems, only a select few industry players can currently work at the frontiers to shape the future of the technology, leaving an imbalance in the AI innovation ecosystem that lacks the diverse voices of academia and government labs. In fact, of the 32 significant industry breakthroughs in AI last year, only three originated from academia, and none from government labs.[13]

The lack of public sector investment in AI means that not only thoughtful regulation but also proper federal procurement and acquisition processes are at risk. Without the ability to train AI talent, the federal government will not have the necessary human capital to create meaningful regulation, ensure ethical AI procurement, and be the true AI leader it has the potential to be.

In June I personally shared with President Biden how I believe the United States is not prepared for this imminent AI moment and how the federal government needs to change its thinking about AI and adopt a moonshot mentality. If the United States is to truly lead in AI, we must not only adopt a robust regulatory and procurement framework, but must also invest in scientific AI research.

This is why I am unequivocally a strong supporter of the CREATE AI Act, a strong bipartisan legislation introduced this summer in both chambers.[14] The CREATE AI Act will establish a National AI Research Resource which will ultimately provide the needed computational infrastructure and data resources to allow academic researchers to innovate and train the next generation of AI leaders.

What we need right now is a coordinated moonshot effort for the nation to ensure America's leadership in AI for the good of humanity. And that moonshot will have to include infrastructure investments such as the National AI Research Resource, as well as national labs focusing on solving the hardest problems. This task will be no small feat, but with meticulous coordination, significant investment, and robust collaboration across government, academia, and industry, we can rise to meet this challenge and ensure America's leadership in AI is both impactful and enduring.

Thank you to the Chairman, Ranking Member, and all the Members of the Committee for allowing me to testify today.

---

[13] Nestor Maslej et al. 2023.

[14] See S.2714 - 118th Congress (2023-2024): CREATE AI Act of 2023, S.2714, 118th Cong. 2023. https://www.congress.gov/bill/118th-congress/senate-bill/2714; H.R.5077 - 118th Congress (2023-2024): Creating Resources for Every American To Experiment with Artificial Intelligence Act of 2023, H.R.5077, 118th Cong. 2023. https://www.congress.gov/bill/118th-congress/house-bill/5077.

Written Testimony

of

Devaki Raj

Former Chief Executive Officer & Co-Founder, CrowdAI

Before the

U.S. Senate Committee on Homeland Security & Government Affairs

Introduction

Chairman Peters, Ranking Member Paul, and distinguished members of the committee, my name is Devaki Raj. I am here representing CrowdAI, a Silicon Valley based start-up leading the development of no-code artificial intelligence tools since 2016. Until a recent acquisition by Saab Inc., I was CrowdAI's CEO and co-founder.

Thank you for this opportunity to testify on "Governing AI Through Acquisition and Procurement" before the Committee on Homeland Security and Government Affairs. I would like to thank Chairman Peters for his leadership on AI initiatives and Ranking Member Paul for his leadership on improving the Small Business Innovation Research (SBIR) / Small Business Technology Transfer (STTR) program for small businesses and startups like mine.

I am a proud American—born in Ohio and raised in Massachusetts, Connecticut and California. So, it is an honor to be here and to present my testimony from the perspective of a small business and startup, the lifeblood of American ingenuity. Before getting started, I want to thank my team for their tireless efforts, as well as my family for their relentless support.

CrowdAI proudly serves the U.S. federal government across multiple mission areas, notably for disaster response and fighting wildfires in California, as well as countering narcotics trafficking in South America. The government recognizes the clear need and value of AI to these critical homeland security missions.

There have been multiple notable efforts towards using AI, especially during the early days of the AI pathfinders, such as Project Maven and the Joint AI Center (JAIC), where "failfast and adapt" was a shared rallying cry.

Indeed, we live in remarkable times with novel challenges: we are witnessing the dawn of generative AI.[1] The pace of technological advancement is such that everyday Americans cannot keep up[2], let alone consider regulations[3] and procurement practices[4]. But, change is needed. To remain stationary at this moment will only result in the U.S. being left behind by our allies and our competitors.

Should members of the committee take away anything from my testimony today, it is that AI must be thought of as a journey, not a destination. In this sense, I prefer the term "machine learning" to AI, as it better aligns with how we understand human development and how we should be thinking about this technology. It is in this difference between a discrete concept, AI, and a continuous process, learning, that today's procurement generally breaks down.

Today, I will share four main observations about AI procurement:

1. First, commercial off-the-shelf AI solutions need government curated data to be mission ready.
2. Second, AI procurement needs to include ongoing AI model training and the infrastructure to support that training.
3. Third, the rapid growth in open-source AI technologies necessitates rigorous testing and evaluation before and after procurement.
4. Finally, it is important to establish paths to programs of record for small businesses through project transition milestones.

To be clear, my testimony is from the perspective of an AI practitioner, based on the work CrowdAI has done with the U.S. government. But AI research is evolving, and rapidly so. The pace of this change is such that even I feel it moving past me at times. Just like AI, I too must continuously learn; which stands to highlight the importance of this and other hearings on artificial intelligence being held this week. I appreciate the invitation to speak and thank you for your interest.

---

[1] Gmyrek, P., Berg, J., Bescond, D. 2023. Generative AI and Jobs: A global analysis of potential effects on job quantity and quality, ILO Working Paper 96 (Geneva, ILO). https://doi.org/10.54394/FHEM8239

[2] https://www.pewresearch.org/science/2023/02/15/public-awareness-of-artificial-intelligence-in-everyday-activities/

[3] https://www.bloomberg.com/news/articles/2023-03-17/chatgpt-leaves-governments-scrambling-for-ai-regulations

[4] https://www.linkedin.com/posts/michaeljkanaan_pentagons-bridge-to-techs-private-sector-activity-6966421664184025088-P2LS

Commercial Off-the-Shelf AI Solutions need Government Curated Data to be Mission Ready

Often, today, we see solicitations ask explicitly or implicitly for ready-to-deploy automated solutions; however, government missions are inherently unique–their sensors, needs, and environments. Commercial off-the-shelf capabilities require government-furnished information to train to those unique missions, regardless if it is for homeland security, intelligence, or public health. Most commercially available AI systems must learn *a priori* about their operational use or they break down. We call this brittleness.

Now, this is a somewhat nuanced point, and some people may bristle at the notion, as certain government programs and companies have built themselves upon the fallacy of "ready-to-use." To be clear, the tools and architectures to create, modify, and operate AI do exist commercially, but the models are built on publicly and commercially available data and marketed aggressively with a "one-size-fits-all" mindset.

The reason for these assertions is that the knowledge and training data needed for many government missions is not in the public domain.

AI needs exposure not only to the data and domain it is to learn from, but that data must match the intended operations environment. In some cases of algorithm development, we can estimate or otherwise get closer to mission data; but, domain shifts from the training data to operations data, will still make models unreliable or brittle. In short, mission specific data are required for each algorithm because each problem posed by homeland security, intelligence, or public health is inherently unique. Accordingly, AI must faithfully represent the situations not only specified by the government, but validated by it too.

Earlier this year, for example, CrowdAI released a toolset that automated remote monitoring of military airfields using "few shot learning" techniques. The tool scales globally, meaning every air force worldwide, while using only a tiny fraction of the training data, in some cases as few as a single image or even a line drawing. This revolutionary advancement in computer vision still had to be vetted by our government partner because we sourced information on aircraft specifications from the internet. We are experts in AI, not in the complexities of Russian, North Korean, and Chinese militaries.

The problem of training robust models to government missions gains complexity as we move from sensor to sensor. The domain shift within a satellite constellation can cause brittleness. Even more so is the shift from commercial sensors to sensitive ones, such as unmanned systems, like aerial drones, or overhead collection platforms operated by

the National Reconnaissance Office. There is no way for industry to know, in advance, what those systems and their data look like unless they have been or are on contract.

When we start delving into other areas of AI, such as large language models and generative AI, which draw their education from a massive corpus of publicly available information, but scant government information, similar domain shift issues arise as well. The greater the domain shift, the more training that is needed, transforming the initial model to its government purpose and no longer "commercial off the shelf".

To accelerate this process, departments and agencies pursuing AI modernization should compile and then furnish datasets so that during contract execution it enables faster transition from commercial off-the-shelf to government off-the-shelf. The National Geospatial-Intelligence Agency, to its credit, for years has been working to curate and redefine how its structures GEOINT data to be machine readable. A herculean task for sure, but a necessary one. Similar efforts, across the government, could be started, if not already done, to identify which missions across the enterprise can be machine-augmented. Force ranked by priority, start curating mission-specific datasets to allow faster and more accurate validation of commercial models and accelerated model tuning. This will help to promote capabilities from Research & Development and transition them to Operations & Maintenance. It is important that I acknowledge that there are constraints that come with use of government data (statutory limitations, privacy, security, etc). But these, for us, have been largely overcome through strong relationships and the "mission-first" attitude we share with our government and industry partners.

However, once this government information is provided, AI isn't a "set it and forget it" solution.

Ongoing AI Retraining: A Continuous Learning Process

As mentioned earlier, AI is a journey. AI procurement needs to include ongoing AI model retraining and a retraining infrastructure.

Currently, procurement processes often buy AI as a one-off software solution, like a copy of Microsoft Word for your desktop. However, due to the nature of machine learning, it is critical to procure AI technologies with the ability to continuously incorporate new data as sensors and missions change, as they invariably do.

Current AI software that is commercially-produced, needs to integrate with current government systems, and continually update on new data to account for new locations, sensors, ecosystems, and missions.

For example, in 2018, CrowdAI collaborated with the California Air National Guard to automate wildfire mapping using MQ-9 drones, a collaboration we proudly continue.

Our initial predictive models performed extremely well in northern California's forested regions where wildfires were common and we had access to government furnished data.

Unbeknownst to us at the time, when the Air Force drone arrived on station, a different sensor was flown for overwatch. Furthermore, the wildfire's epicenter was in a suburban area. As evolving wildfire epicenters shifted and sensors were updated, our models required retraining to maintain operational relevance. These changes represented a twofold domain shift from the previous year's model, which needed to be retrained. Anticipating these operational inevitabilities, we built model retraining into our contract to avoid unnecessary contract options or extensions that, during disaster missions, could cause fatal delays.

While AI models are flexible, they still require contracting officers to include model retraining in contracts, ensuring alignment with evolving mission data. It's a dynamic process, akin to software updates, not just a one-time procurement.

For that AI development project, we increased wildfire map production from just one per day to updates pushed electronically to firefighters every 30 minutes. This was a momentous improvement for our first responders. And our goal in the next six months is to achieve real-time situational awareness, fully automated, from the overhead sensor to users' mobile devices. With our partners at the 163rd Operations Support Squadron, the Hap Arnold Innovation Center, and California Department of Forestry and Fire Protection, this will be reality before next fire season, further saving lives and providing an exemplar for others to follow.

Model training and retraining is not only inevitable, it is the rule. As the world constantly evolves, it is important for a model to continue learning and not be bound by its original training. Contracts for AI must include continuous learning. This ensures that the government is not using a stagnant model that declines in time, but one that continuously improves and adapts as it is exposed to more data.

For the purposes of operationally-ready AI, there is no functional difference between development and operations, as they are continuously interlinked. Therefore, we must move past AI being funded as one off software purchase, but build procurement vehicles that bake in ongoing updates or service level agreements. This is not only because AI needs re-training but also to provide the procuring officer with technology that is cutting edge. It is important that commercial tools, available today, can be implemented at the enterprise level to operate AI at the pace of mission as both the AI and sensor technology continues to evolve.

Rigorous Testing and Evaluation of AI Solutions

Today, anyone can go online and download a trained computer vision model. For example, it is easy to find detection models using state-of-the-art architecture trained on millions of images collected all over the Internet. For many use cases, it is a great tool for experimentation and use.

Cyber threats aside, the risk of using publically available AI architectures creates room for moral hazard. Companies today are incentivized (if not required!) to claim AI expertise, even when it is limited. There often is little-to-no ability by a procurement officer to verify vendor claims of the robustness of their models: Are these models purpose-built or open sourced? How will it perform on my data? These questions are near impossible to answer *a priori*.

Source selection panels for contracting have the difficult task of interpreting changing rules about AI and inferring proposing companies' credentials and ability to deliver. Because AI moves so fast, industry has been particularly susceptible to research that only exists in white papers and technical proposals. Companies recognize that to be competitive they must sell the art of the possible versus the science of today. The result can cut either of two ways: the company delivers on its promises or it does not.

On account of how procurement works today, companies often project future potential capabilities, regardless if they are possible, to win federal awards. The risk of development failure comes later and falls on the government. The only risk to these companies is failure to secure contracts. This phenomenon can be mitigated through more rigorous market analysis and, more importantly, sound model testing and evaluation—during solicitation of capabilities, such as with AI competitions that grant favorable terms to participants for commercialization, and after delivery on an ongoing basis.

Codifying government-wide standards for AI testing and evaluation would help mitigate unverifiable corporate claims. As we move from mission to mission, department to agency, everyone's thinking today about AI is different. Those differences not only leave room for risk, but drive confusion both in industry and government.

Both quantitative and qualitative evaluations of AI are important: quantitative evaluations give information on how well the model is doing across large amounts of data, but can sometimes obscure important information by only reporting aggregate metrics. We've found that it is equally critical to evaluate data samples to get a sense of how and why a model is performing, and more importantly, qualitatively if the model is accomplishing the task assigned to it.

An example CrowdAI encountered was finding airstrips used by narco-traffickers. We built a tool that found areas meeting our description: long, straight dirt paths cleared of vegetation. However, in addition to finding airstrips in the jungle, it often also found dirt roads, which were prolific throughout the region. The point being that evaluating AI isn't as simple as setting an arbitrary performance score. Yes, those metrics matter, but they represent only part of the story. Performance must also consider if the capability solved the problem, in this case finding runways.

This is why I recommend both quantitative and qualitative metrics. Knowing these upfront can alleviate confusion later about why or how a system performed in an operational environment. For this to happen, we must educate our federal workforce on AI fundamentals and, then, provide specific training on how it relates to their role. I commend Chairman Peters for leading AI legislation to address some of these challenges through the AI Training for the Acquisition Workforce Act.

Establishing Paths to Programs of Record for Small Businesses through Project Transition Milestones

With no-code tooling, such as ours, and the advent of generative AI, it will not be long before we see AI being developed in every government agency, business, university, and home across the country, whether people know it or not. Like when the appstore for the iPhone opened itself to 3rd party creators, I anticipate an explosion of AI tools and models beyond what we can imagine today. So, rather than force innovative solutions through one or a few large prime contractors or government innovation units, we should normalize standards and federate acquisition to benefit all businesses.

Startups, at first glance, may appear agile and resilient; but, on government timelines, they are fragile and exposed. The "procurement Valley of Death" continues to take its toll not only on small businesses, but the government's efforts to source and deploy AI, as well. Even after selection, with funding obligated and sole-source authorization under SBIR, we have seen contracts take the better part of a year to award—a lifetime for a small business. So, we look for other means.

One of the fastest ways to get on contract is through a subcontract with a systems integrator. But, working through a prime contractor is a double-edged sword. Startups, so eager for work, can reveal their innovations and thinking, giving primes new ideas. It is a risk we take in order to get the work—to live to fight another day.

The preferred alternative is direct award from the government. The Small Business Innovation Research program is one such vehicle, and I want to acknowledge the members of this committee for SBIR's reauthorization last year. Thank you for continuing this valuable program.

It is difficult to overstate the transformative capabilities the SBIR program has for small businesses like mine. SBIR's sole-source award policy allows small businesses to compete with large prime contractors. This rule helps level the playing field, increasing the types of new technologies that the government can procure, as well as the speed with which they can test and evaluate them.

SBIR's phased approach spans requirement validation, initial development, and transition, which provides a more navigable path than perhaps any other contracting mechanism in government. The SBIR program has the added benefit of protecting intellectual property and promoting small businesses.

However, SBIR is struggling.

Over the last few years, we've watched the quantity of SBIR awardees balloon, the size of awards shrink, and the quantity of transitioned projects slip. This suggests to me that the current system is incentivized by the wrong metric: recognizing departments and agencies for the quantity of awardees versus the number of transitions to programs of record.

I have found it crucial as a small business owner that any government procurement have clear transition milestones for a path to a program of record. A great example of a transition partner is Naval Air Warfare Center, which includes project transition milestones in its contract milestones.

Conclusion

In conclusion, the needs and resources of government missions are unique, requiring tailored AI solutions. Therefore, procurement vehicles must reflect the iterative nature of AI, and the introduction of standards for testing and evaluation will promote more effective AI adoption. Most importantly, AI education and relevant procurement training are both imperatives for increasing AI adoption across the federal enterprise.

Finally, in all phases of an AI project lifecycle, remember that machine learning, just like human learning, is a journey and not a destination.

Thank you for your time.

Written Testimony
of
Will Roberts
Director of Emerging Technologies
ASI Government, LLC
Before the U.S. Senate Committee on Homeland Security and Governmental Affairs (HSAG)

I. INTRODUCTION:

My name is Will Roberts, and I have worked in the Federal procurement space for the last 15 years. I am currently the Director of Acquisition Solutions and Emerging Technology for ASI Government, LLC. Previously, I was the Acquisition director for the Joint AI Center (JAIC) at the U.S. Department of Defense.

I am particularly passionate about Government procurement, and during my time at the JAIC I became very aware of the procurement-related challenges in delivering AI to government end users. As I will be repeating throughout my testimony, my time in this space instilled a strong belief that the Acquisition Professional - particularly the Contracting Officer - serves one of the most important roles in navigating how the Government will harness AI for the welfare and defense of this nation.

But I didn't always feel this way – my attitude towards AI and emerging technology used to mirror the current sentiments of many federal acquisition professionals. Before joining the JAIC, I was a contract professional at the Air Force, my goal was to vector over to Wright-Patterson AFB and work on major systems alongside some of the best minds in federal contracting. I wanted to help build new airplanes, weapon systems and the like. I wanted to be part of the "next big thing." I wanted to take the defense contract skills I learned and apply them to real tangible challenges. Technology was not tangible to me, and I looked at emerging technology as providing more costs and headaches than any real gains to a program or mission. As a consumer, I am among the worst adopters of technology. I still read a newspaper. My wife and I just recently bought a Roomba this year – we didn't really trust it before then. So, I sympathetically relate to the current aversion to AI adoption that exists in many agencies and offices.

Months before my family and I were heading to Dayton, OH an innovation advisor to the Secretary of the Air Force, whom I had worked some early AI contracts for, convinced me to apply for the lead acquisition position at newly created JAIC under the leadership of Lt. General Shanahan. I took the chance and now look back at the decision as a pivotal one in my own career. During my time at the JAIC, I read more, I had access to more information, and my attitude toward AI completely changed. One of my primary lessons learned was simply the scale and importance of artificial intelligence as a revolutionary technology. AI truly is the "next big thing" – and our federal workforce must be more prepared to realize its incredible benefits but also its risks. But technical knowledge of AI is only one ingredient in the recipe for success. I became aware of some *new* skills that the modern acquisition professional needed to develop to successfully buy AI functions and deliver them into government missions at the speed of relevance. These weren't weapon systems skills or large airplane-buying skills. There was not a clearly written framework for these skills. The AI procurement professional had to think different, fast, and agile. The federal acquisition professional had to truly understand the AI marketplace and recalibrate his existing procurement domain knowledge. Applying government procurement domain knowledge to an

AI functionality is not intuitive, nor can one easily lift and shift their traditional acquisition expertise. It is a new skill that must be *learned*. And, as I will try to demonstrate in my testimony, the Contracting Officer and the Acquisition Team as a whole, are among the most important roles in the effort to modernize the government missions through AI. Success in this new chapter of U.S History rests in the hands of a very diverse acquisition team. And, I will add, it requires a special level of *talent*.

II. THE AI PROCUREMENT TALENT WAR

In the commercial marketplace, the AI Talent war has intensified, brought about in large part by the introduction of large language models and generative AI. A recent Wall Street Journal article described many companies paying as high as seven figure salaries to AI programmers and data scientists. These are all technical skills and it makes sense that various industries are vying for such aggressive investments. Such investments make sense because the private marketplace supplies innovations, so they need the technical talent. Our American industry fosters our current and future inventors and creators.

In the same way, the Government should be involved in a major AI talent effort. However, the Government's aggressive talent investments should not focus primarily on technical expertise – simply because, when it comes to AI, the Government does not *make* it. It *buys* it. And buying it is hard enough. Buying and delivering AI – creating that bridge from the technology to the end user – is not something companies can really do, certainly not as affectively as the Government can. This is an inherently government function. The Government knows its missions, its end users, and its internal bureaucracy. The Government creates the bridge between the product and the mission. Under the bridge is the chasm in which products die and never see adoption. But buying AI is a function that the government is not doing very well, and *this* should be the focus on our hiring, and training efforts. We must be seriously concentrated in cultivating top-notch modern acquisition teams. These are the bridge builders that enable technology adoption.

I want to take a moment to talk about the diversity of a good acquisition team. A typical AI project, from ideation to adoption, requires a special team that operates in a way that is rare in federal acquisition. Technical experts alone, and contractors alone, will not achieve success. Many factors exist – money issues, legal problems, ethical concerns, contractual matters – that threaten a project's momentum and potentially stop a project dead in its tracks. It takes diverse skills, managed by a very capable product manager. This is a government function. And these skills must be applied in very new ways. These various experts must also be tightly knit because the delivery of AI is multifaceted and not subject to traditional phases. In AI government acquisition, the development, procurement, and sustainment all happen at the same time and in cycles. This means the budgetary people talk to the testing and evaluation people, the contracting people talk to the end users, etc. It is not played like a "relay race," typical in traditional acquisitions, where the baton is handed in linear fashion from budget to procurement to testing, etc. It is a team sport, and the entire team needs to run the ball together down the field, pivoting and reacting to the dynamic environment. During my time in this space, I witnessed budget and money experts looking at their fiscal laws and procedures in new ways and forced to make innovations to bridge the technology over to the end users. Experts in the field of law and social sciences became crucial, but they also were entering uncharted territory in their expertise and were confronted with very new concepts relating to human-machine interaction, AI's impacts on the workforce, and a myriad of ethical risks. And, of course, the contracting professional becomes an

essential key to the team. And as with the other team members – contracting officers must apply their contractual domain knowledge in new ways.

At the end of this testimony, I will provide two recommendations for how the Government can get on the right track in cultivating a modern acquisition team – particularly with special focus on the procurement professional.

The first would be granting various contract authorities to all components in the federal government. As I will explain later, the modern Contracting professional needs a full and diverse toolbelt. Every tool, whether Other Transactions, of FAR-based vehicles, Public Private Partnerships, or Partnership Intermediary Agreements – fits a unique need and within the diverse range of AI projects, a skilled contracting professional will use each tool.

But tools are no good if you don't know how to use them, or when the contracting talent is not empowered through trust by their leaders. And so, the second and much more important recommendation is for a much more robust, substantive, and universally mandatory AI Acquisition training program for all current and incoming acquisition professionals.

III. THE IMPORTANCE OF THE CONTRACT FOR AI SUCCESS

I want to take a minute to provide a little bit more of my background, because I think it would provide some benefit to this testimony, particularly as it concerns the importance of our nation's contracting officers. When I was in law school, I became fascinated with contract law. And mainly because I saw in contract law something special. I saw two people agreeing upon something, and in the process creating their own law -which more or less the courts would protect. This "law" could even be written on a napkin – it was the essence of the *deal*. It was the "tit-for tat" – the way that both parties can benefit from an arrangement. It is the foundation of business. And I loved it – because it transcended rules that are created by official lawmakers. Because these laws, these agreements, could be written by two ordinary people on a napkin and retain the full force of law.

As I delved more into these principles, I started to study Government contracting – and became even more fascinated. In this context, the U.S. Government becomes a business partner – they become one of these parties that create a binding agreement with another – again formalized in a written agreement -- which becomes, in a way, a law written by the parties – *not lawmakers*. Starting in the year 1800, Congress began creating rules to curb and control the flexibility of Government contracting officials acting in this capacity – to prevent abuse and maintain stability in the contracting process. But even so, when the Government contracts with a company – the Government enters the market and *engages in business*. The Government creates terms which become *binding*.

But the Government is not a company. Instead of shareholders, this business is financed by American taxpayers. This means that *everyone* invests in the business. And the taxpayers are paid back through a different means of return than rising stock value. The taxpayers seek a return in the form of welfare, defense, peace, and security. And so, this is the ultimate responsibility of the American contracting officer – to bring these returns. It is for this reason that I believe the American Contracting Officer serves one of the most critical roles in our Government.

It is with this interest that I left law school after passing the bar and took an oath of office as a civil servant, to do my part in ensuring taxpayers received such returns from the business of Government.

And in my 15 years I realized that the course of U.S. History can be summarized in a series of transactions. In fact, this is all history is to me… it's a series of business transactions. Starting in the Revolutionary War before the U.S. became a nation, the government has relied on industry. Some transactions, we should not be proud of, and we still carry the scars of these business decisions, many of which were outlined in the four corners of contractual agreements. But, for the most part, historical business transactions formed the great nation we are today. American industry and ingenuity, not the government, was the source. The Government was the means to connect that ingenuity to the mission to strengthen the nation and benefit the taxpayer. But it was the American inventor that created the airplane. It was industry that facilitated the industrial mobilization effort that supplied the planes, ships and tanks that helped us win World War II. It was the ingenuity of industry that took us to the moon in the 1960s. All accomplished through a series of contracts between the government and industry.

But we are turning a new page in our history as a nation. There is indeed a new technology that is powerful. This technology will change the nation, it will change the way families live their lives, businesses operate, and nations interact with one another. It presents numerous advantages and many dangers. We have the opportunity now to become ready for this growing revolution, but currently we are not. And so, the question for the Government is not – "*how do we develop it*?" If we are to follow the historical path that has made our nation successful in the past, the question for the Government must be – "*how do we buy it*?"

This question of "*how do we buy it*" – this was my life for the past three years as I headed acquisition activities for the DoD Joint AI Center. It can be complicated. To every acquisition professional working today, these are very exciting and historical times, and in many ways they will feel like they are navigating unpaved paths in a frontier. Many aspects of the job were *very* unpaved, such as negotiating terms for the responsible use of certain AI functionalities, such as those involving warfighting and medical procedures. Navigating the wild frontiers of this technology is exciting but can also be very dangerous. It is a job that must be taken seriously.

I'm going to mention three main examples of how AI Procurement is different and unique: (1) intellectual property; (2) responsible use of AI; and (3) the incorporation of agile performance language.

When it comes to intellectual property, for example, there are many unique considerations for the prudent Contracting Officer. There must be a balance between rewarding the American inventor and nontraditional company, while at the same time preventing arrangements that are not advantageous to future government operations. It takes unique knowledge of the technical components of AI in order to determine the proper IP strategy. The rights to the data, for example, will probably be different than the rights to the AI model. Even with data, we have input and output data. We have trained and untrained models. We have infrastructure that runs the pipelines and hosts the AI application – each of these components require careful thought into the appropriate IP ownership. Insisting on rights to the wrong things will discourage the right players from providing technology to our end users. On the other hand, giving rights to the wrong things will lock the government into one company, which will balloon costs on a program and prevent any new competition and innovation. It not only takes a knowledge of the underlying technology (technical knowledge) to navigate these waters, it takes knowledge of the market (business knowledge) and adequate knowledge of tailorable IP language (domain expertise). The AI Training Act tackled the technical knowledge for the civilian agencies. The reality is that all three areas of knowledge (technology, business, and contract domain) are missing and are not treated as a priority.

Responsible AI becomes another unique dilemma. When considering the topic of "AI Trustworthiness" , there are two forms of *trust* that must be attained: (1) trust in the *functionality* of the AI model (i.e. will it work?); (2) trust in the *responsible use* of the product (i.e. is it safe/ethical?). Losing trust in functionality will prevent early adoption and create skepticism – something we have historically seen in our slower adoption of airplanes, submarines, and the radio (to name only a few examples). Losing trust in *responsible use* is more serious, as it pertains to safety, privacy, and equal treatment. For now, the four corners of the contract define the mutual agreement on how to handle the parameters of what is "responsible use." This has been an interesting challenge. In procurement, the contractor's quality control of responsible use can be evaluated in very powerful ways – 2 especially: (1) as a discriminator for contract award selection; and (2) as a metric for testing and evaluation. Emphasizing the responsible use of AI (RAI) in either of these two phases of the acquisition sends a clear message to the contractor, but also requires contract professionals to set very clear and objective definitions of what is and what is not responsible. On one end, the parties can mutually agree to keep the terms ambiguous – thereby making any responsibility for RAI meaningless. On the other end, the government may push for terms so restrictive that most companies will grow wary to contributing their talent to the mission. The latter is perhaps even more problematic. As with IP, I will resort to the same three areas of knowledge for the prudent and competent AI procurement professional: (1) Technical; (2) Business; and (3) Contract Domain knowledge. Technical knowledge to know the various types of data biases, risks, and mitigation tactics involved in responsible use of AI. Business knowledge to gauge the attitudes and awareness in this critical topic, including how to speak about the Government's position to cautious companies, and understanding nontraditional companies' resourcing capabilities to comply with any potentially restrictive RAI requirements. Contract domain knowledge to understand when and where such agreements should be articulated in contractually binding language, and which agreements should be worked instead through the ongoing business relationship. In other words, RAI is loaded and intricate. AI is so diverse that RAI risks vary according to the circumstances. Some AI functionalities are extremely low risk, while others impact human life or privacy. It's new and important. But mastering the three areas I mentioned would resolve most of the issues. But *most importantly*: fear of risks should not prevent us from utilizing this technology – as the technology will often prevent many more risks associated with human error (in everything from business processing, medical diagnosis, and even defense activities). If AI Acquisition professionals are unable to skillfully navigate these risks and alleviate anxieties and fears, the true benefits of AI on government missions may never be fully realized.

Finally, the AI procurement professional must understand the concept of "agility." This is a very strange concept in our current procurement environment – but every competent procurement professional that acquires AI must rise above their culture and engage in agile and flexible contracting. In some environments, the word "agile" has become an annoying buzzword. However, agility is essential for successful AI delivery. Agile contracting can be summarized in three sentences: Contract fast. Iterate Often. Fail Early. Contracting fast to keep pace with the speed of relevance in emerging technology. Iterate often implies that all contracts would be results-based instead of requirements-based. In other words, the entire acquisition team was focused on results based in phases, or sprints. As results are recorded and value is measured, the team builds the new iteration to improve what is working and stop what is not. Finally – the "fail early" philosophy was a direct response to the fallacy of sunk costs. Setting up contract agility means you can pull the plug before things get bad. In other words, you prevent wasting taxpayer dollars on bad AI projects. The current procurement and acquisition process is anything but agile. It is more akin to creating a huge barge that is approaching a port. If, within a few

hours of reaching port, it is discovered that the requirement must change – or the underlying technology has changed – it is too late. No time to turn. No time to adjust. There's no stopping that big barge from coming in. Agile contracting creates swift boats that can swerve and pivot among the volatile waves of technological change. This is a paradigm shift in thinking. It is currently practiced in Government, but only by a small percentage. Agile must become mainstream. This requires an aptitude that, again, is not emphasized or required across the board in federal procurement. It is currently not a core competency, and it should be. The longer we keep this tucked away as a niche, the longer we remain completely unprepared for the modern challenges that await us as a nation.

Agile AI contracting also involves a keen understanding and prudent utilization of all the contract authorities available to the procurement professionals. The prudent, trusted business advisor must have a variety of tools in her toolbelt. FAR-based tools, Other Transactions, Technology Transfer agreements, and the like. There is no 'one-size-fits-all solution – and the wrong kind of contract will often negatively impact successful adoption. In many situations, the intellectual property strategy will determine the most optimal contract mechanism. In other situations, the unique aspects of the marketplace would determine the most appropriate course – such as opting for a public private partnership if the Government is merely a contributor to a larger commercial effort. This may lead a prudent CO to use a specialized technology transfer vehicle such as a Public Private Partnership or a Partnership Intermediary Agreement. These types of determinations require AI business acumen which, again, must be *learned*. Through a combination of heighten business acumen and domain knowledge of all the contractual authorities available, the competent CO will then wield her tools in a way that benefits the market, contracts at the speed of relevance, and focuses on results over process.

Within the four corners of the contract, the parties agree upon IP terms that can completely destroy the Government mission or lead a major spark of new AI functionalities in our nation's commercial marketplace. Within the four corners of the contract, the parties agree on the parameters of responsible use for AI. Such agreements could lead to potentially disastrous results that extinguish the momentum or lead to effective uses of powerful technology which paves the way for future progress done safely. Within the four corners of the contract, the parties set up agile performance that can lead to actual measurable value. In essence, within the four corners of the contract rests the fate of successful artificial intelligence adoption. The role of the dealmakers, then, becomes paramount.

IV. AI PROCUREMENT TRAINING CHALLENGE

So what, then, is the level of priority and importance within the executive agencies? As of today, AI acquisition training is sparse. Even the training that exists among certain agencies is not universally mandated to all professionals, but rather seen as a niche – or elective training. New interns entering the workforce are currently not required to learn about the technology, the marketplace, or contractual challenges of AI – or for any new technology for that matter.

This is not good, especially if we are to acknowledge and agree that AI will continue grow into a transformative and revolutionary technology that will impact every mission, agency, and field office in the Government. The end users are incredibly diverse – further testament to the universal impact of AI. They include soldiers protecting us from foreign threats as well as immigrants wanting a more efficient citizenship process. From IRS processing to forestry service activities. From health and human services

to border protection. AI can enhance it all – increasing mission impact, severely cutting costs, and even saving lives.

But with the current lack of priority to build these needed AI procurement skills across the board, we are essentially setting our future workforce up for failure. We are looking over at a commercial gold mine, ready to be mined. And we have no miners.

I was encouraged to see Congress release the AI Literacy requirements for the Department of Defense, as well as the AI Acquisition Training ACT for civilian agencies. But I do not believe they go far enough.

I worry that the execution of these statues will not extend to all acquisition professionals. I also worry about the content of the AI Acquisition curriculum. The various elements of the AI Training Act, for example, are within the competency of technical knowledge. As I stated earlier, this is critical information for any AI Acquisition program, but it is only one piece. Technical knowledge of AI alone will not bridge the gap between product and user adoption. Acquisition professionals must be trained in how to buy AI, how to deliver AI, and how to sustain modern technology in a way that brings true results and widespread adoption. In other words, all acquisition professionals need to be trained on technical, business, and contract domain knowledge. Technical knowledge alone will not bridge the valley of death and imbed new inventions into Government missions.

V. THE ASKS

And so I close with two major recommendations for consideration of the committee. The second recommendation is more important than the first.

My first recommendation is for more contractual authorities to be provided to more contracting offices across the Government. FAR Based contracts, Other Transactions, Partnership Intermediary Agreements, Commercial Solutions Openings, Public Private Partnerships… these are examples of tools that should be in every modern acquisition professional's arsenal. Each one responds to unique objectives or marketplace conditions. Using a less optimal contract authority could result in less optimal competition, and ultimately a completely different product. In preparation for this testimony I spoke with some Contracting and Agreements officers who were among the few that have access to all the tools. One told me that each authority has served a different purpose in AI Acquisition, and that it would have been very difficult and not quite as effective for her to have used the Federal Acquisition Regulation on every procurement action. Many offices are limited to the FAR, and this therefore limits the tools that contracting officers can use as an effective business advisor and civil servant in modern procurement.

However, even if every single contract office had access to all statutory contract authorities, my fear is that – without any accompanying workforce development – very few offices would take advantage of these additional authorities. Or worse – due to substandard hiring and workforce development, contract offices managers would not trust their contracting officers to make sound business and contract decisions, as some tools require more expertise than others.

And so, the more important recommendation I have is one that I have echoed throughout this testimony. If we are to believe that AI is a revolutionary technology that will impact every single government mission – we must make AI Acquisition training mandatory across the entire federal acquisition workforce. The trainings must be robust and focused on the three core competency skills: (1) technical

knowledge of artificial intelligence, its architecture and risks; (2) business knowledge of the marketplace, and more effective market research strategies; and (3) contract knowledge as applied to AI, to include IP, responsible use of AI, and agile contracting techniques. Again, this training must be robust and substantive, captivating and inspirational. Not dull and basic. The emergence of technology makes this style of training that important.

Our bridge builders need to get smart about this. They need to be *trained to be trusted.* Anything less would strip them of the flexibility to successfully negotiate a successful public private meeting of the minds.

VI. CLOSING

There is an urgency to this, but we are not too late if we start now. There are already existing commercial AI capabilities that could improve public welfare and defense with greater efficiency and less taxpayer expense. Some of these existing narrow AI solutions are not just low hanging fruit. The fruit has fallen from the tree and sits in the grass, perfectly ripe and ready to eat. And because no one picks it up, it rots into obsolescence. So many technologies that could have provided high value impact. So who are those forces who can walk over and pick this fruit up? Again, it is my belief that this can be done by a talented and diverse acquisition team, to include the AI procurement expert. People with the right knowledge, skills, and abilities are scarce and in high demand. We need to expand the pool of people with the knowledge of what to buy, how to buy it, and then how to deliver it.

Thank you for giving me this opportunity to speak about the importance of AI Procurement. I have dedicated years of hard work as a civil servant trying to make this better. My eyes have been opened to the importance of AI and the critical need to transform our current procurement processes and culture to fully take advantage of it. But, unfortunately – my eyes were opened because I fit a very niche role in the Government as the Acquisition chief of DoD AI center. My understanding and dedication, along with the small percentage of hard-working civil servants in this space should not be niche. Every acquisition professional in the Government should come to understand what I have come to understand.

Only then will AI stop serving as a specialized field for a small percentage of Government procurement professionals, but rather a core competency – mandatory for all incoming interns. Only then will we get serious about buying, adopting, and using this already existing and powerful technology. And only then, will we finally start to realize the major surges in efficiency and savings in cost that this technological revolution will absolutely provide for the welfare and defense of our nation. Thank you.

Will's AI Acquisition Book List

Books on Tech Industry

The Code, Margaret O'Mara

Genius Makers, Cade Metz

Competing in the Age of AI, Marco Iansiti, Karim R. Lakhani

AI Technology & Government

AI Superpowers, Kai-Fu Lee

T-Minus AI, Michael Kanaan

Army of None, Paul Scharre

Wired for War, P.W. Singer *** book club favorite*

AIQ, Nick Polson, James Scott

The Kill Chain, Christian Brose ***book club favorite*

Government Contracting and Technological History

Freedom's Forge, Arthur Herman *** book club favorite*

Boyd, Robert Coram

Never Mind We'll Do It Ourselves, Alec Bierbauer and Col. Mark Cooter

Skunk Works, Ben Rich and Leo Janus

American Moonshot, Douglas Brinkley

Fiction

2034, Elliot Ackerman, Admiral James Stavridis

Ghost Fleet, P.W. Singer and August Cole

Change Management and Office Productivity

The Phoenix Project, Gene Kim, Kevin Behr, and George Spafford

The Unicorn Project, Gene Kim *** book club favorite*

Thinking Fast and Slow, Daniel Kahneman

Inside the Five-Sided Box, Ash Carter

AI For The People

Internet users, not the government or Big Tech, should control the AI to filter online content

Testimony by Michael Shellenberger

Before the Senate Committee on Homeland Security and Governmental Affairs

On the topic of:
"Governing AI Through Acquisition and Procurement"

September 14, 2023

Chairman Peters, ranking member Paul, and Committee members thank you for your stated concern with the implications of AI for our civil liberties and constitutional rights, and for requesting my testimony. I am honored to provide it.

The ability to create deep fakes and fake news through the use of AI is a major threat to democracy, say many experts. "AI-generated images and videos have triggered a panic among researchers, politicians and even some tech workers who warn that fabricated photos and videos could mislead voters, in what a U.N. AI adviser called in one interview the 'deepfake election,'" reported the Washington Post late last month. "The concerns have pushed regulators into action. Leading tech companies recently promised the White House they would develop tools to allow users to detect whether media is made by AI."[1]

But the threat of AI to elections today is as overblown as the threat of Russian disinformation to elections in 2020. Never before has the U.S. been better prepared to detect deep fakes and fake news than we are today. In truth, the U.S. Department of Defense has been developing such tools for decades. In 1999, Defense Advanced Research Applications (DARPA) described its funding for R&D as having the goal of "total situational awareness" through "data mining," "face recognition," and computer networks to evaluate "semantic content." The proposal anticipates the direction of the technology over the following 25 years.[2]

Before elaborating on this point, I want to emphasize that I view AI as a human, not a machine, problem, as well as dual-use technology with the potential for good and bad. My attitude toward AI is the same, fundamentally, as it is toward other powerful tools we have developed, from nuclear energy to biomedical research. With such powerful tools, democratic civilian control and transparent use of these technologies allow for their safe use, while secret, undemocratic, and military control increases the danger. The problem, in a nutshell, is not with the technology of computers attempting to emulate human thinking through algorithms, but rather who will control it and how.

---

[1] Cat Zakrzewski, "ChatGPT breaks its own rules on political messages," Washington Post, August 28, 2023.

[2] J. Brian Sharkey, "Charging Into the Next Millenium: Total Information Awareness," *Accessed via Internet Archive*, June 7-10, 1999.

There is a widespread belief that users already choose their own content on social media platforms. We choose who to follow, and see their posts on the Facebook, X, Instagram, Facebook, and YouTube feeds. In truth, social media platforms decide a significant portion of what users see. YouTube's recommendation algorithm, for example, determines 70% of what people watch on the platform, a share that did not change between 2018[3] and 2022.[4]

The amount of recommended content is lower on other platforms. Meta said last year that just 15% of total Facebook feed content is recommended content from non-followed accounts,[5] while 40 percent of Instagram's feed content is.[6]

But Meta CEO Mark Zuckerberg said last year that he expects Facebook will double the percentage of recommended content by the end of 2023. And users have little to no control over what is recommended to them. In fact, research published in late 2022 found that users have little control over the videos that YouTube feeds them.[7] On every other platform, the algorithms are hidden from users.

The heavy lifting of censorship or "content moderation" was by 2021 done overwhelmingly by AI. Zuckerberg said, "more than 95% of the hate speech that [Facebook] take[s] down is done by an AI [artificial intelligence] and not by a person. . . . And I think it's 98 or 99% of the terrorist content that we take down is identified by an AI and not a person."[8] Similarly, 99% of Twitter's content takedowns started with machine learning.[9]

The problem with AI technology today funded by the US government, whether DARPA or National Science Foundation (NSF), is fundamentally around the control of these technologies by small groups of individuals and institutions remarkably unaccountable to the citizens of the United States. While there is always a diversity of

---

[3] Ashley Rodriguez, "YouTube's algorithms drive 70% of what we watch," *QZ*, July 13, 2018.

[4] Hana Kiros, "Hated that video? YouTube's algorithm might push you another just like it," *MIT Tech Review*, September 20, 2022.

[5] Meta, Q2 2022 Earnings, July 27, 2022.

[6] Rachael Davies, "Nearly half of the posts you see on Instagram are from accounts you don't follow," *Evening Standard,* April 28, 2023.

[7] Hana Kiros, "Hated that video? YouTube's algorithm might push you another just like it," *MIT Tech Review*, September 20, 2022

[8] Feerst, Alex. "The Use of AI in Online Content Moderation" *Digital Governance Working Group*, Sept. 2022. (p. 2.)

[9]Kristen Ruby, "Twitter Artificial Intelligence," *Ruby Media Group*, December 26, 2022.

agendas and motivations behind what decision-makers in the AI space are doing, many U.S. government-funded individuals and institutions behind deep fake alarmism are, not coincidentally, demanding greater governmental or nongovernmental control over social media platforms and Internet companies.

Why is that? Why have elements within the US government promoted AI for online censorship? And can AI be used to advance free speech and free expression instead?

AI and the Censorship Industrial Complex

This Censorship Industrial Complex of government agencies and government contractors has its roots in the war on terrorism and the expansion of surveillance after 9/11. President George W. Bush that year authorized the National Security Agency to monitor Americans who were suspected of having a 'nexus to terrorism,' resulting in the Agency's now-infamous and illegal interception of information."[10] In 2003 DARPA told Congress that NSA was its "experimental partner" using [Total Information Awareness (TIA)] and AI to detect false information.[11] Ten years later, in 2013, a US military contractor named Edward Snowden revealed to reporters that the NSA was collecting telephone records of millions of Verizon customers,[12] and accessing Google and Facebook to secretly collect data.[13]

During the same period, the U.S. intelligence community (IC) and DOD alike recognized how essential AI would become to their operations overall. In 2013, a New York Times report on the NSA's use of AI foreshadowed how "counter-disinformation" experts would, nearly a decade later, describe fighting misinformation online.[14] "Computers could instantly sift through the mass of Internet

---

[10] Scott Shane, "Giving In to the Surveillance State," *New York Times*, August 22, 2012.

[11] DARPA, "Report to Congress Regarding the Terrorism Information Awareness Program," *DARPA Information Awareness Office*, May 20, 2003.

[12] Glenn Greenwald, "NSA collecting phone records of millions of Verizon customers daily," *The Guardian*, June 6, 2013.

[13] Barton Gellman &Laura Poitras, "U.S., British intelligence mining data from nine U.S. Internet companies in broad secret program," *The Washington Post*, June 7, 2013.

[14] James Risen & Eric Lichtblau, "How the U.S. Uses Technology to Mine More Data More Quickly," *The New York Times*, June 8, 2013.

communications data," reported the Times, "see patterns of suspicious online behavior and thus narrow the hunt for terrorists." In 2014, the DOD unveiled its "Third Offset Strategy," which emphasized that AI would change how the US prepared for cyberwar with China and Russia.[15]

In 2015, DARPA launched the funding track that directly resulted in the AI tools that leading Internet and social media companies use today. That fall, DARPA invited proposals for its MediFor program.[16] The goal? Develop a science and practice for "determining the authenticity and establishing the integrity of visual media."[17] DARPA funded universities to create the MediFor platform to automatically detect manipulations.[18]

DARPA's warning eight years ago is identical to the Washington Post's warning about deep fake last month. "Mirroring this rise in digital imagery is the associated ability for even relatively unskilled users to manipulate and distort the message of the visual media," warned DARPA. "While many manipulations are benign, performed for fun or for artistic value, others are for adversarial purposes, such as propaganda or misinformation campaigns."

The adoption of AI grew alongside alarmism about deep fakes and "misinformation," and "disinformation" more broadly. In 2016, Facebook reported it had developed AI to automatically censor offensive live videos.[19] In early January 6, 2017, outgoing Obama Administration DHS Secretary Jeh Johnson designated "election infrastructure" as "critical infrastructure," which would become the mandate of the Cybersecurity and Infrastructure Security Agency (CISA), which Congress created the following year to protect. In 2018, journalists revealed that Facebook was using AI to predict users' future actions for advertisers.[20]

---

[15] Gentile et al., A History of the Third Offset, 2014–2018, *Rand Corporation*, 2021.

[16] Dr. William Corvey, Media Forensics (MediFor) (Archived), *darpa.mil*, nd.

[17] Media Forensics (MediFor) Grant DARPA-BAA-15-58, *grants.gov*, September 29, 2015.

[18] Contractors included Notre Dame, Purdue University, Duke University, Ideal Innovations Inc, Schaefer Corporation, University of Siena, New York University, University of Southern California, Politecnico di Milano, Unicamp, NVIDIA, Columbia University, Dartmouth, University of Albany, UC Berkeley, and Kitware.

[19] Kristina Cooke, "Facebook developing artificial intelligence to flag offensive live videos," *Reuters*, December 1, 2016.

[20] Sam Biddle, "Facebook uses artificial intelligence to predict your future actions for advertisers, says confidential document," *The Intercept*, April 13, 2018.

In 2019, DARPA launched "Semantic Forensics," the successor to Medifor. SemaFor funded think-tanks, academic institutions, software companies, social media, and search engine organizations as part of a four-year project to develop AI meant to detect deep fakes, or synthetic or manipulated media.[21] It gave contracts to five primary organizations: Kitware, PAR Government, STR, Lockheed Martin, and SRI International, with this financing further divided amongst other universities and research institutes.

Commercial interests in both policing deep fake and advocating policies to censor synthetic media popped up during this period. Also in 2019, a new nongovernmental organization called The "DeepTrust Alliance" launched a series of events called the "Fix Fake Symposia."[22] The DeepTrust Alliance described itself as "the ecosystem to tackle disinformation," and its website invited audiences to "Join the global network actively driving policy and technology to confront the threat of malicious deep fakes and disinformation."[23]

The goal of Deep Trust appeared to be to advocate for policies aimed at criminalizing "digital harms," including forms of speech that hurt people. "If the behavior is malicious," said the group's CEO, Kathryn Harrison, in 2020, "that's a problem. Laws need to be extended to digital harms… There needs to be a standard set of practices" across social media platforms.[24] "I want to see society put more safeguards in place," she said. "This is like cars, right? When you first had cars, you didn't have seat belts…. We're in a very similar situation in the media ecosystem and can save information at light speed but no safety net. That's what we need to build."

It was also in 2020 that DHS' CISA created an "Election Integrity Partnership" to censor election skepticism. It partnered with four groups: Graphika, the University of Washington, the Atlantic Council's DFR Lab, and the Stanford Internet Observatory. Graphikia and UW are DARPA's Semafor grantees. In Deep Trust's report, it names those four groups and progressive philanthropic donors, and other NGOs and government. EIP claims it classified 21,897,364 individual posts

---

[21] Semantic Forensics (SemaFor) Grant HR001119S0085, *sam.gov*, November 19, 2019.

[22] Aros Harrinson, "Deepfake, Cheapfake: The Internet's Next Earthquake?"*Fix Fake Symposium Proceedings Part 1*, 2020.

[23] DeepTrust Alliance, Homepage, *deeptrustalliance.org*, nd.

[24]Jon Prial & Kathryn Harrison, "Episode 133: Tackling Digital Disinformation with Kathryn Harrison," *Georgian Impact Podcast*, December 11, 2020.

comprising unique "misinformation incidents" from August 15, 2020, to December 12, 2020, from a larger 859 million set of tweets connected to"misinformation narratives."[25]

By January of 2021, CISA unilaterally broadened its scope "to promote more flexibility to focus on general" misinformation, disinformation, and malinformation. Where misinformation can be unintentional, disinformation is defined as deliberate, while malinformation can include accurate information that is "misleading." Two months later, DARPA announced that it had funded Accenture Federal Services (AFS), Google/Carahsoft, New York University (NYU), NVIDIA, and Systems & Technology Research (STR) to "develop automated tools that aid analysts as they tackle the looming rise of automated multimodal media manipulation," otherwise known as deep fakes or fake news.[26]

While social media platforms use AI to identify and censor content, the decisions of what to censor, and how remain in the hands of humans, specifically executives at social media platforms. And so those individuals and groups that wished to see greater censorship by social media platforms rolled out a major initiative in the spring of 2022 to establish a US government agency to do precisely that. In April, DHS announced that it had created a "Disinformation Governance Board," ostensibly to protect national security by fighting disinformation, misinformation, and malinformation on social media.[27] One week earlier, former U.S. President Barack Obama gave a speech at Stanford calling for government regulation of online speech with the same justification as Deep Trust's Kathryn Harrison: preventing harm and protecting democracy.

One month later, in May of 2022, DARPA launched its "Model Influence Pathways," or MIP, program to automate the process of discovering the origins and "pathways" of "misinformation, disinformation, and manipulated information."[28] The

---

[25] UW Center for an Informed Public, Digital Forensic Research Lab, Graphika, and Stanford Internet Observatory, "The Long Fuse: Misinformation and the 2020 Election," *Stanford Digital Repository: Election Integrity Partnership*, 2021.

[26]Matt Turek, "DARPA Announces Research Teams Selected to Semantic Forensics Program," *darpa.mil*, March 2, 2021.

[27] Amanda Seitz, "Disinformation board to tackle Russia, migrant smugglers," *AP*, April 28, 2022.

[28] Dr. Brian Kettler, Model Influence Pathways (MIP), *darpa.mil*, May 4, 2022.

goal of the program appears to be to develop tools so social media companies can reduce the virality or spread of disfavored social media posts. In that sense, it is within the vision of Stanford Internet Observatory's leader, Renee Diresta, who has long championed simply reducing the spread of disfavored views, rather than removing them from platforms outright. Preventing virality delivers most of the benefits of outright censorship with the benefit of not being noticed and thus not triggering the Streisand effect.[29]

The Federal Trade Commission in June of last year warned Congress about the dangers of using AI for censorship and urged "great caution." Good intentions weren't enough, said FTC, because "it turns out that even such well-intended AI uses can have some of the same problems — like bias, discrimination, and censorship — often discussed in connection with other uses of AI."[30] The FTC specifically pushed back against the idea, widely promoted by individuals and institutions within the Censorship Industrial Complex, that AI should be used to reduce harm. Noted the report authors, "while some harms refer to content that is plainly illegal, others involve speech protected by the First Amendment."

The FTC's warning was well-timed. Six months later, the Twitter Files would reveal Twitter executives over-ruling the determination by their own Trust and Safety team that President Donald J. Trump's tweets had not incited violence, but they deplatformed him anyway, under both external societal pressure and internal employee pressure. Shortly after, emails revealed White House staff demanding that Facebook executives censor "often-true" information about COVID-19 vaccine side effects under explicit or implicit financial threats, behaviors which the Fifth Circuit Court of Appeals last week ruled were unconstitutional.[31]

Both the Twitter and Facebook files exposed the large involvement, influence over, and infiltration by former government intelligence and security officials. "Facebook currently employs at least 115 people, in high-ranking positions, that formerly worked at FBI/CIA/NSA/DHS," noted an analyst. "17 CIA, 37 FBI, 23 NSA,

---

[29] Michael Shellenberger, "Why Renee Diresta Leads the Censorship Industry," *Public.Substack.com*, April 3, 2020.

[30] Federal Trade Commission, Combatting Online Harms Through Innovation, *Report to Congress*, June 16, 2022.

[31] Michael Shellenberger, "War on Free Speech War On Free Speech Means Social Media Users Must Be Free To Moderate Their Content," *Public*, September 9, 2023.

38 DHS."[32] This influence may carry over to today's people seeking to rescue, ostensibly independently, the legitimacy of the US government, which sits at the intersection of technology and foreign policy. Harrison, for example, worked in the French Ministry of Defense, received a graduate degree from Georgetown, and was a term member at the Council on Foreign Relations before working with IBM on AI and then founding Deep Trust.[33]

Why have elements within the US government promoted AI for online censorship? Part of the reason is a well-intentioned concern over real-world harm, and undermining of liberal democracy. But another part of it appears to stem from an inappropriate and exaggerated sense of entitlement by DARPA contractors to work with social media platforms to censor disfavored voices.

User-Based Content Moderation

The Fifth Circuit Court ruling showed the limits of the First Amendment to protect free speech online. The judges ruled that the Cybersecurity and Infrastructure Security Agency (CISA) of the Department of Homeland Security had likely not violated the First Amendment in creating an elaborate system for "flagging" content for Facebook, Twitter, and other social media platforms to censor. The court suggested that such mass flagging operations may be constitutionally protected free speech, at least if done right.

I believe that the way CISA used AI to mass-flag so-called "Covid misinformation" in 2021, through its partnership with "The Virality Project," created by Stanford Internet Observatory (SIO) and others, was a government infringement on freedom of speech. Through such mass flagging, CISA indirectly demanded that Twitter and Facebook censors "often true" information about vaccine side effects. We believe that, with Biden simultaneously threatening the Section 230 legal status of the social media platforms, having CISA's partners make their demands constituted coercion.

---

[32] @nameredacted, twitter.com/NameRedacted247/status/1604641866342756352?s=20, X, December 18, 2022, 4:56 PM.

[33] Kathryn Ann Harrison *Experiences*. LinkedIn. Retrieved September 11, 2023.

But I also recognize that the Fifth Circuit court is saying that such AI-supported mass flagging by "government partners" like SIO could be constitutionally protected if it did not involve coercion or, on the flip side, any incentive to cooperate. The First Amendment prevents the government from "abridging" or limiting speech. It doesn't prevent government officials from telling publishers, whether of books, news articles, or social media posts, that, in their opinion, they shouldn't be publishing those books, articles, or social media posts. The line the Circuit Court wants to draw is on relatively direct and obvious coercion, not jawboning.

Whether or not the Supreme Court decides to hear the case and draw the line somewhere else, the ruling points to the need for Congress to take action to protect freedom of speech by defunding government contractors that advocate widened censorship by social media platforms, and exercising greater oversight over contractors developing AI tools.

The threat to our civil liberties comes not from AI but from the people who want to control it and use it to censor, rather than let users control, information. The obvious solution is for Congress to require that social media companies allow users to moderate their own content in exchange for Section 230's sweeping liability protections, which allow them to exist. This specific suggestion is something another committee will need to consider.

What this committee can consider is a related FTC recommendation, which is using the power of procurement to put AI tools in the hands of users, not the hands of big tech companies. "Filters that enable people, at their discretion, to block certain kinds of sensitive or harmful content are one example of such user tools," FTC notes. The way these tools work should be transparent; users should have a right to know how these tools work. Giving users control over what content they see and don't see is the solution most consistent with the American tradition of free of speech.

Users should be able to decide for themselves whether or not to use these filters and other tools, not Internet companies, the government, a nongovernmental organization, or anyone else. Some tools are already becoming available. Microsoft a Video Authenticator in 2020, while Adobe's Content Credentials allow users to detect whether the content is likely to be authentic and unaltered. Requiring people to affirmatively choose their filters will require more reflective and slow thinking about their content choices.

FTC errs in suggesting that Congress give government-certified researchers, rather than users, access to the algorithms and content moderating filters. A longstanding goal of its leaders is to allow US government-certified researchers to gain access to the data of social media platforms so they can then demand censorship of disfavored views behind closed doors. This is what the "Platform Accountability and Transparency Act," which Obama endorsed, would do. It would allow "researchers" to act as de facto censors. Such activities may be constitutional, but they are antithetical to the values of transparency, privacy, and free speech.

Finally, this committee should seek to encourage or even mandate that DARPA contractors be required to share their research in a more visible way, and stand for questions from the general public. Of the roughly 60 organizations, many if not most of which have been funded by the US government to fight "mis- and dis-information," that my colleagues and I emailed in the spring, none agreed to stand for an interview.[34] The refusal to speak to the public is an odd behavior from those whose livelihoods depend on the goodwill of the public. Congress should consider some provision whereby contractor recipients of taxpayer money must expose themselves to scrutiny.

At the same time, deep fakes and other forms of synthetic media are new, deception, disinformation, and misinformation are not. One of the oft-repeated claims of those advocating expanded online censorship is that, by allowing falsehoods to go viral and undetected, the Internet poses a heretofore unanticipated threat. But the same thing was said about the Gutenberg printing press, the radio, and television. The solution today, as then, is for users to correct misinformation with good information, for themselves, not other people.

None of the above information is likely to put an end to the alarmism about the threat to democracy from deep fakes and AI. But it may help expose much of it as coming from individuals and institutions with an interest in exploiting the alarmism for personal or political gain.

---

[34] Matt Taibbi, "Report on the Censorship-Industrial Complex," *Racket News,* April 25, 2023.

Testimony to United States Senate
Committee on Homeland Security and Governmental Affairs
"Governing AI Through Acquisition and Procurement"

Biography of Tim Cooke

CEO and Owner, ASI Government LLC

My name is Tim Cooke. I am an Economist and Management Consultant with experience in federal acquisition and Information Technology dating from the late 1980's. I earned a PHD in Political Economy at Johns Hopkins University and became an Assistant Professor at Rice University before joining the Center for Naval Analysis FFRDC. I then helped to lead the growth of SRA International from $50M to $1.7B and its sale to private equity. When it was time to leave that phase of my career, I joined ASI Government LLC in 2013, two years later becoming the CEO, and then two years later, in 2017, the owner.

In those roles, I witnessed firsthand the power of government contracting officers to build bridges from the needs of government to the world's most consistently innovative economy that prevailed during the Cold War years. I have seen first-hand the continuing American leadership in many fields of technology including internet technology and now the potential of AI. In the field of AI, I led the SRA teams that deployed natural language processing and structured data AI technology for many Wall Street firms in the late 1990s. My interest in IT was inspired by my father's experience helping to build flight simulators as a contractor for the Navy in the late 1960s.

With that inspiration and experience, I earned a Hammer Award from Secretary Perry and his team at DoD for my work in acquisition reform in 1996 while at SRA. I subsequently earned a Fed 100 award, and the ACT-IAC award for Collaboration in 2020 for leading the development of the OFPP "Periodic Table of Acquisition Innovations," to discover and share contracting solutions, as well as leading the ASI team that developed the training for the Digital IT Acquisition Professional certification in 2015. More recently, I published an article for the U Penn Law *Regulation Review* that highlighted the key role of federal contracting officers in assuring the acquisition of responsible, ethical AI for government use cases.

I dedicate much of my time to volunteer service for better partnership between industry and government in pursuing better outcomes for taxpayers. I serve on the Board of Directors of the Professional Services Council, the Executive Committee of ACT-IAC where I am Vice Chair for Finance, the Board of Advisors of the National Contract Management Association (NCMA), and the Acquisition Committee of the Intelligence and National Security Alliance (INSA).

Testimony to United States Senate
Committee on Homeland Security and Governmental Affairs
"Governing AI Through Acquisition and Procurement"

I. Introduction

In the nascent Age of AI, helping the nation realize the value of responsible AI to meet the complex and competitive challenges of the coming decades is critically dependent on the knowledge, skills and abilities of federal contracting officers and their acquisition team collaborators. Because it is so vital to our nation's future, I allocate much of my time and company investment to that area, including hiring Will Roberts, who is testifying before you today.

Will is here to testify to the importance and urgency of equipping federal contracting officers and acquisition teams with the knowledge and practice of buying in this rapidly evolving market. He recommends requiring that the entire acquisition workforce be trained in the technology, business, and contracting knowledge, skills, and abilities sufficient to help create and effectively deploy the solutions needed by government missions. I wholeheartedly support his recommendations. As Will said in his testimony, much of our success in this new chapter of U.S History rests in the hands of federal acquisition teams orchestrated by Contracting Officers...who require a special level of *talent* to meet the global competition and the national needs.

II. Training Talented IT Acquisition Problem Solvers

ASI's founders were experts in government acquisition of information technology for the U.S. Air Force during the time of Clinger-Cohen revisions of IT acquisition regulations. Today's innovators are making their own waves from where they sit, some with the advice and assistance of my ASI Team. As Will pointed out, the fundamental question for government AI is not "how do I develop my AI solution, but *rather what do I buy and how do I buy it*?"

When ASI answered the challenge posed by OMB/OFPP and US Digital Service to develop the experiential Digital IT Acquisition Professional (DITAP) training program for IT Acquisition Specialists we saw the opportunity to modernize the federal talent pool of buyers to keep up with the rapid pace of technical progress in IT. Since the program's inception ASI has taught roughly 1,000 buyers of federal IT to build their skills in acquiring emerging technology. As the IT industry has evolved, we have expanded the original scope to include the realm of rapidly evolving AI technologies.

As an example of the success of that program, one of our Contracting Officer DITAP graduates designed and implemented the Pilot IRS program with support from IRS Procurement leadership. Utilizing only FAR authorities, IRS is able award, test, terminate, or pivot their use of AI and related technologies to achieve mission and management objectives. Contracts are awarded within 30 days under a modular contract approach in which each short phase is reassessed in periods of 1-3 months during the initial period of performance, as opposed to a year or more in standard contracting practice. It is an example of "buying like a VC" within the rules of the FAR. Pilot IRS is one element of the many appearing in the Periodic Table of Acquisition Innovations (PTAI) to share the best ideas in solving contracting problems. Fundamental to this success was the approval and support of IRS Procurement leadership at the time.

The OMB DITAP program is the kind of training that the Acquisition workforce needs and wants. DITAP-like programs need to increase in scope, breadth, and relevance. As IT and AI become embedded in more

mission solutions, the training needed to buy those solutions effectively and efficiently needs to be broadly acquired as it is needed by the workforce. On a government-wide basis, the training should ideally be consistent with the Agile Acquisition Framework in DoD and the Government-wide Category Management requirements. It should further be integrated in the the 2023 FAC-C and the Defense Acquisition Workforce (DAWDF) programs.

To be effective buyers within a category management discipline, federal acquisition teams must be smart about the marketplace, and expert in selecting and using the contracting authorities of the government to achieve the goals of the Federal Acquisition system. The degree of understanding Program Managers and Contracting Officers need to be aligned to mission goals and outcomes and incentivized to achieve them while earning returns for taxpayers in the form of value for money or ROI.

Flexibility and trust are intrinsic attributes of the profession necessary for achieving good outcomes from the marketplace. Of course, verification of that trust and management latitude (FAR Part 1) should be part of a value-added and efficient oversight of the acquisition system. Knowing the evolving art of the possible in acquisition is the first step.

My observation is that government Acquisition teams are consistently running faster to stand still as their workload grows and complexity while the workforce shifts to a less experienced mix. *Despite* additional buying authorities that can dramatically reduce required red tape, the work is only shifted to higher value and potentially more beneficial activities that were not possible due to the manpower required and the increasing workload. *Despite* numerous regulation revisions to simplify the administrative burden and create a more lean and supportive system, the churning of work continues. *Despite* current continuous training requirements, much less the needed training proposed here, the relevance of training to the workforce continues to lag the technology and business model OODA loop (Observe, Orient, Decide, Act) of relevance. *Despite* the urgency of critical buying activity, the behavior of a loss-averse, compliance-driven workforce spends precious time preparing to be second guessed or protested, limiting their ability to provide timely acquisitions. *Despite* good intentions, the workforce often must resort to work arounds to attract the best of industry to serve the defense and economic security needs of the nation. All of this makes it more challenging to develop and deploy the needed training to be effective buyers of AI and complex solutions for government missions.

III. Accessing AI Ecosystems through Cloud Service Providers

As buyers of AI services, government acquisition teams are finding that Cloud Service Providers (CSPs) are often partners of choice because of the effectiveness and efficiency of the capabilities they offer. There is an emerging convergence in the markets for AI, Cloud, and Cyber technologies. CSPs are the gateway service providers for all those technologies and the pathway to solutions for government use cases. For mission solutions of the future, CSP offerings will become more important. Business relationships with CSPs will continue to grow in importance.

Government constraints such as the anti-deficiency rule do not align with the consumption-based utility business model of CSPs. That fact has caused much unnecessary consternation by Contracting Officers and government procurement lawyers. To realize the potential of AI, the government needs to become expert in new fields and skills such as the commercial best practice for managing cloud investments and

expenditures known as "Cloud Fin Ops". Several federal departments are testing FinOps Cloud acquisition and management methods now, including Energy, Veterans Affairs, Office of Personnel Management, and Army. GSA and OMB are assisting in These experiments are being monitored by GSA and OMB for potential near-term adoption on a broader scale. I recommend that they be instituted across government as a standard practice for consumption-based cloud services.

The leading CSPs like AWS, Azure, and Google, offer hundreds of AI and ML solutions that combined with DevOps, Data, and Security offerings yield a daunting array of complex potential solutions for government buyers, who need to understand the choices and risks involved to make good buying decisions. And that is just the technical dimension of the choice, not including the market intelligence or the knowledge and expertise to appropriately use many government contracting authorities.

Mapping those cloud market capabilities and solutions to the wide array of government use cases is an overriding challenge of the first order (ACT-IAC is developing a National Use Case Library that may be helpful). The expertise of government program managers partnered with the other members of the acquisition team needs to mature in order to effectively solve the combinatorial capability matching problem for each use case. The learning needed by program managers is just as challenging as for Contracting Officers, and ideally, they learn together and practice together with other members of the acquisition team. Any of the specialists on an acquisition team can stall its progress if they do not share perspectives and understanding of the mission problem and solution opportunities.

Fortunately, the US economy has developed the world's leading CSPs and commercial AI solutions. The critical job of the federal acquisition workforce will be to develop, deliver and manage contracts that provide the government access to globally competitive, responsible applications of AI. That is why they need to understand what they are buying in terms of the technology, the marketplace in terms of who the leading providers are, and the contract mechanisms available to bring the suppliers of capabilities to the needs of government.

IV. The AI Ecosystem Complements Human Expertise

The current AI is more than the sum of its parts as we have learned from the evolving recent revolution in Large Language Models (LLMs) and Generative AI. Its effectiveness depends largely upon the knowledge of the application area. Functional knowledge of government business and improvement opportunities is essential to get started.

It is very early in this new age of AI. The use cases that I'm familiar with are largely point solutions for problems that have been solved in other ways, including human (natural) intelligence-based manual solutions, coding logical relationships between high quality data vectors based on known relationships, and data analytics and statistical models. All these methods require structured rather than unstructured data and analytics. Human (natural) intelligence is largely expressed through written or spoken unstructured communication. The success of the new AI is built on Large Language Models (LLMs) using massive models trained with or without human inputs regarding the entities and relationships between the language-based elements. These models became possible in part because of the continuing miniaturization (Moore's Law) of the physical elements underlying parallel processing units in computer

technology. In the coming decades, as the cumulative power of AI increases, we will discover functional and systematic uses of AI that the acquisition workforce needs to know what these solutions are and how to buy them – as-a-service, open source, tailored be-spoke models, or other solution type. Your legislation should provide reasonable, efficient guidance for continuing education consistent with local discretion and global oversight as new, more capable solutions become available.

Machine learning is capable of distinguishing patterns that humans have not discovered. This is the basis of the hope and fear in building and verifying AI-based information processing machines whose results are not explainable to users and stakeholders. It is also the source of angst about unintended and undesirable behaviors of AI. We need to capitalize on the power of AI while balancing its risks. From the point of view of government users and buyers of the technology, it is the acquisition team that defines how the risk is to be mitigated. But Government acquisition teams are not yet familiar with the risks, or the opportunities for avoiding or mitigating those risks. They need to learn how to think about the risks and understand the contractual tools they can use to create responsible use. Organizations throughout government, including NIST and DoD's DIU, have already created smart principles and approaches for responsibly buying AI capabilities. Learning what they have done and being able to practice the associated methods needs to be part of AI training for the acquisition workforce.

V. Upskilling and Reskilling the Acquisition Workforce for the Age of AI

Commercial firms are making substantial investments in upskilling and reskilling their workforces *for the Age of AI, which is likely to be as disruptive to knowledge work as the industrialization of workflows was in manufacturing.* The government needs to do the same thing. More than upskilling, reskilling is to learn completely new kinds of work. If upskilling means to build new skills upon the foundation that you already have, then reskilling means starting anew without using your existing skills as a foundation. With the advent of Generative AI, a wholly new set of skills is needed by the acquisition workforce on both sides of the market. To get a sense of the difference, we must look back to today from a future world in which AI is pervasively embedded in the work to be done by buyers and sellers of products and services needed for future government missions.

To foresee and imagine the future workplace, like Edison or Ford, to see what potential exists for technologies like AI and what they portend for work in the future requires some imagination to extrapolate from the future of technology to the probable impacts on the workforce and its emerging skill needs. By its nature, AI disrupts knowledge workers first and foremost, and is now creating an urgent need for government to update its approaches to preparing the acquisition workforce for the future.

Access to and use of AI-powered tools are accelerating, and is already affecting the federal acquisition workforce, which is largely learning new skills on their own initiative, especially on the industry side of the market. Organized reskilling efforts lag the rapid rise of easy-to-use technology at many companies and the federal government. Like the DoD, Chief Digital and Artificial Intelligence Office (CDAO) initiative called *TryAI*, or the IRS's *Pilot IRS*, experiential learning enables ad hoc upskilling of current workers to use new AI tools in their daily practice. Some have called this upskilling-by-default. Upskilling can be a

cost-effective way to maintain the benefits of an experienced workforce that understands the business. There is also the potential to reskill people into occupations that have excess demand caused by a skill shortage, which may be the case in several federal occupations including the contracting profession.

More formally, employees should be able to choose from a set of instructional methods, from experiential learning to self-directed online courses that offer structured, still just-in-time learning with employees able to tailor their own paths to success. The new FAC-C revision by OMB paved the way for a new education paradigm for the GS1102 community that is based on a self-designed curriculum of courses pertinent to one's career path. One size does not fit all needs.

VI. Recommendations

Providing for the education of Contracting Officers, Program Managers, General Counsel, Agency Executives, and other members of Acquisition Teams is critical to government's access to AI solutions via private sector capabilities through the marketplace. This puts Acquisition teams in a lead role, with orchestration by the Contracting Officer. Preparing the Contracting Officers of the future for the new age of AI needs to start now. Your earlier legislation on AI Training for the acquisition workforce was timely, though I know of little progress at OMB on this complex topic.

1. Encourage upskilling by default by providing guidance on appropriate use of Generative AI and LLMs to protect confidential information. It may be sufficient to provide a statement of objectives that Departments and Agencies would be required to tailor for their specific needs.
2. *Statements of objectives* for upskilling and reskilling the acquisition workforce for AI for the needs of the acquisition workforce covering the areas recommended by Will Roberts:
    a. High-level technical training on the methods and sources of AI capabilities
    b. Market intelligence on suppliers and their rapidly evolving offerings, including contract terms and conditions to improve access to needed current and future capabilities
    c. Expanded Federal contracting tools and authorities and guidance on their appropriate use, including expanding Other Transactions contracting tools across government.
3. Incorporate specific training material on contact terms and conditions covering:
    a. intellectual property rights;
    b. responsible use of AI;
    c. the incorporation of agile performance language in contracts.
4. Expand the DITAP program to all government buyers of IT capabilities.
    a. Provide funding to maintain DITAP's relevance and applicability to the rapidly changing world of Digital IT, including AI.
5. Provide guidance officials of the government allowing a "utility" consumption-based approach to contract pricing and encourage good government methods like the commercial best practice Cloud FinOps approach to provide more value for money and prevent waste.

Advancing the education of the acquisition workforce, especially the contracting officer cadre, to understand not just the technology of AI, but also its business aspects, the nature of the market, the types of contracting authorities available and how to use them will go further to speeding up adoption of

AI in government than any other initiative. Government does not produce its own AI tools and techniques, rather it buys them from commercial entities using the government's contracting authorities. Like the challenges faced in past wars and tech races, it will be the success of those contracts that will determine the success of the United States.

Statement for the Record
Submitted by Scale AI
In Response to the U.S. Senate Committee on Homeland Security and Governmental Affairs Hearing Entitled
"Governing AI through Acquisition and Procurement"

Scale AI (Scale) appreciates the opportunity to provide a Statement for the Record in response to the U.S. Senate Committee on Homeland Security and Governmental Affairs' (Committee) hearing entitled "Governing AI through Acquisition and Procurement."[1] We applaud the Committee's leadership in ensuring that the United States continues to maintain its global leadership in the adoption of AI.

Since our earliest days as a company, Scale has worked across industries to help accelerate the development of AI. Today, we work with the leading frontier model developers to fine-tune, red team, and test and evaluate their large language models. This work also provides us a unique vantage point to understand, create, and implement leading practices and to understand a thoughtful balance between a regulatory approach that maximizes innovation, while putting in place the proper guardrails.

In addition to our work with the leading tech and automotive companies, Scale has worked with the U.S. government since 2020, helping to ensure the United States leads the world in AI adoption. The Department of Defense (DoD) has led the deployment of AI in the federal government to date, and there are no shortage of use cases for every agency to better enhance their day-to-day operations through AI. Our collective years of working with leading commercial companies and the DoD has uniquely positioned Scale to understand what works and provide recommendations for the federal government to efficiently and fully embrace AI.

U.S. leadership in AI requires efficient government AI adoption

Artificial Intelligence is likely to be the most impactful technological innovation since the invention of the internet. Due to the promise of AI, nations around the world are heavily investing in leading the world in the adoption of it to reap the economic, national security, and societal benefits that will accompany it. To date, the United States has positioned itself as the global leader in AI adoption, both in the commercial sector and for government use cases like national security. However, this leadership cannot be taken for granted because other nations like China are actively vying to supplant the United States in this regard. Maintaining our leadership requires the right governance

[1] See, https://www.hsgac.senate.gov/hearings/governing-ai-through-acquisition-and-procurement-2/

framework to maximize innovation while ensuring the proper guardrails. To do so, it is critical that industry and government work together to best understand where the existing laws and regulations apply, where updates may be needed to modernize these to include AI, and finally where new ones are needed. The initial step is crucial in establishing a robust framework..

Beyond understanding the governance, our government must also safely and efficiently adopt AI. It is not a stretch to say that there are limitless use cases to help drive efficiencies across the federal government. For example, a Large Language Model (LLM) at the Department of Energy could rapidly speed up the time currently required to better understand the energy grid limitations in cities around the United States for next generation electric and hydrogen vehicle charging.

Because this technology is promising, it is critical that the federal government takes the proactive steps necessary to prepare for, fund, and acquire AI. If this does not happen, our nation risks losing its global leadership in the sector.

AI-ready data is the key to unlocking the power of AI

Both the commercial industry and the federal government have learned the importance of AI-ready data that is labeled and annotated to truly unlock the power of AI. For this reason, the leading commercial entities spend billions per year to ensure they have the right data strategy, management procedures, and retention policies in place.

The DoD also learned this lesson in its early work on AI, most notably with Project Maven–which was launched in 2017 to speed up AI adoption in the military–and now the DoD fully recognizes that AI ready data is critical for the Department's own data strategy.

Today, the Chief Digital and AI Office (CDAO) has been established with the responsibility of creating a Centralized Data Repository to build the DoD's "AI Scaffolding" or data infrastructure to power AI. While this top-down leadership is promising, it is also imperative that efforts take place at every level of the DoD to ensure that harnessing the power of AI-ready data is prioritized. Additionally, the Intelligence Community has also released a long-term data strategy that highlights the need for AI-ready data. While these efforts are maturing, they can serve as a blueprint for the rest of the federal government.

Every individual agency has numerous use cases for AI to improve their day-to-day efficiencies.This vision only becomes reality if raw and siloed data becomes AI-ready. For that reason, it is clear that long-term plans and funding is necessary within each agency to accelerate the development and prioritization of AI-ready data.

Scale recommends that all government agencies develop AI scaffolding and create AI-ready data strategies that will truly unlock the potential of AI across the myriad of government use cases.

Test and Evaluation is critical to protect U.S. investment and to ensure high-quality AI systems

Over the next year, the federal government is likely to spend millions of dollars acquiring AI systems, which are critical to maintain our place as the global leader in AI adoption. Understanding whether AI is safe to deploy is one of the most important questions the federal government must answer as AI is inevitably adopted broadly.

Scale has worked for years across the leading AI developers and this experience has demonstrated that the best way to ensure responsible AI is through a risk-based approach to Test and Evaluation with human oversight. This not only protects taxpayer resources by ensuring that the government acquires high-quality AI systems, but also is one of the strongest methods to ensure accuracy, limit bias, and uphold the Responsible AI Principles, such as those outlined in the Biden Administration's AI Bill of Rights.

Test and Evaluation has long been a key part of the product development cycle for responsibly bringing consumer-facing technologies to market and military technologies into production. This is essential for AI applications because they are rapidly developing and constantly iterating and therefore constantly presenting new opportunities and risks to the end user. A risk-based approach to Test and Evaluation will ensure that AI is factual, accurate, and explainable regardless of the underlying model or data being used. If the product–including the data infrastructure and underlying model–does not meet these requirements, we risk sacrificing user trust in the technology. These standards should be set by the specific agency based on the intended use case.

While red teaming is necessary to understand the unanticipated vulnerabilities associated with an AI system, a comprehensive Test and Evaluation framework is still imperative to ensure that those vulnerabilities are addressed. In practice, this means that the safety bar is aligned to the risk of the activity the AI is supporting. For example, an LLM used for mission critical intelligence analysis will need to meet a higher bar prior to deployment than an LLM used to compile routine reports.

The DoD has already begun to understand the best approaches to Test and Evaluation for LLMs with the idea to require this prior to acquisition. Industry is also working towards the development of commercial Test and Evaluation platforms and standards, that should be a critical part of the development process, prior to bringing a product to market. All federal agencies should learn from the DoD's years of experience working on this topic and update their acquisition guidelines accordingly.

Scale recommends that the federal government mandate a risk-based approach to Test and Evaluation with the deployment of any LLM on a government network.

Innovative and expedited paths to acquisition will allow the United States to maintain global competitiveness

The federal acquisition process was designed for legacy platforms, such as aircraft carriers and weapons systems, and can take years to complete. Many tweaks have been considered and adopted to help accelerate the process. However, the acquisition process has not yet adapted to reflect the rapid development of leading technologies today. In addition to the legacy platforms, the government is now acquiring software and other technologies that are iterating rapidly and may be outdated by the time a single acquisition process takes place.

For this reason, industry has long highlighted the "Valley of Death" concept that has hampered the federal government's ability to acquire innovative technology. To address this issue, the DoD developed Tradewinds Solutions Marketplace, a rapid acquisition vehicle that seeks to accelerate the procurement and adoption of AI and Machine Learning technology. Scale recently was awarded a contract through Tradewinds that occurred in near-record time for a federal government acquisition.

Beyond the DoD, the rest of the federal government should be encouraged to expedite the acquisition process by leveraging all available contracting authorities that could make a significant difference for any agency seeking emerging commercial technologies. A slow acquisition process can hamper the government's ability to use innovative technologies, but with AI, it will directly impact our nation's global leadership because the technology is proving vital to national and economic security. Our adversaries, like China, are not letting bureaucratic processes stand in their way of working towards global leadership.

When acquiring emerging technologies like AI, Scale strongly recommends that all federal agencies embrace, to the maximum extent possible, the full range of acquisition authorities to use frameworks such as Tradewinds or similar tools that allow for the timely acquisition of the best-in class commercial technologies. Doing so will directly help the U.S. lead the world in the adoption of AI.

Conclusion

Thank you again for convening this hearing today and calling attention to this critical topic. Scale looks forward to continuing to work with the Committee, Congress and the Biden-Harris Administration to put in place the right regulatory framework to ensure that the United States continues to lead the world in the adoption of AI.

Written Statement of Professor Anjana Susarla

Professor of Responsible AI, Michigan State University

Submitted to the U.S. Senate Committee on Homeland Security and Governmental Affairs

For a Hearing on Governing AI Through Acquisition and Procurement

July 25, 2023

I thank the Senate committee for investigating the role of procurement and acquisition of artificial intelligence (AI) in the federal government. I submit this statement to highlight to members of this Committee, and other members of Congress, that oversight and responsible use of Artificial Intelligence is critical especially in governmental uses of AI, and that federal agencies need to be cognizant of potential algorithmic harms posed by AI systems in the procurement of AI.

I am a professor of Responsible AI at the Eli Broad College of Business at Michigan State University. I have more than two decades of experience researching various aspects of sourcing (an important part of acquisition and procurement), artificial intelligence (AI), and social media. I am representing myself in this statement, and the views I express are my own.

AI can be defined[1] as "the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings". Automated decision making can be defined[2] as "the process of making a decision by automated means without any human involvement." Recent advances in AI are primarily enabled by machine learning, a prediction technology. The deployment of such predictive methods has not only resulted in labor substitution, but also resulted in the large-scale transformation of tasks formerly associated with cognition and reasoning[3].

In this statement, I want to highlight how the ubiquity of algorithmic decision making create huge invisible costs (which are disparate in their impact for different groups of people), which makes it imperative for policymakers to understand the impact of how governments should approach the procurement of such technological tools and services.

1. Several governmental functions become digitized and transformed using predictive AI methods, which makes it imperative to understand how such AI-based services are procured by the federal government.

Facial recognition software, commonly used in predictive policing[4] and national security[5], has been shown to exhibit biases against people of color[6] and a threat to civil liberties[7]. Facial surveillance is

---

[1] https://www.britannica.com/technology/artificial-intelligence

[2] https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/individual-rights/automated-decision-making-and-profiling/what-is-automated-individual-decision-making-and-profiling/

[3] Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2019. "Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction." *Journal of Economic Perspectives*, 33 (2): 31-50.

[4] https://www.forbes.com/sites/nikitamalik/2018/10/29/the-problems-with-using-artificial-intelligence-and-facial-recognition-in-policing/

[5] https://www.theverge.com/2019/12/20/21031255/facial-recognition-algorithm-bias-gender-race-age-federal-nest-investigation-analysis-amazon

[6] https://www.cnet.com/news/facial-recognition-has-always-troubled-people-of-color-everyone-should-listen

[7] https://theconversation.com/high-tech-surveillance-amplifies-police-bias-and-overreach-140225

only one facet of a broader set of challenges in employing algorithms and algorithmic decision making. The widespread use of predictive methods without adequate guardrails can result in substantial consumer harms[8].

Machine learning methods apply inductive logic in generalizing patterns from training data. A machine learning based resume screening tool was unwittingly found to be biased against women[9], since the training data reflected past practices where most resumes were submitted by men. Millions of applications for healthcare, food stamps, and cash benefits can be jeopardized when predictive methods label routine mistakes[10] as a "failure to cooperate." An audit[11] of an automated decision-making system that was responsible for allocating quality of care found that the algorithm was less likely to refer black people than white people who were equally sick to programs that are meant to improve quality of care for patients with complex medical needs. This is because the way risk scores were assigned to patients was based on total health-care costs accrued in one year. Taken in the aggregate, the average black person in the data set had a similar profile of overall health-care costs to the average white person. However, the average black person also was likely to have a greater prevalence of conditions such as diabetes, anemia, kidney failure and high blood pressure. In other words, looking at the health care costs accrued in a year could significantly understate the true health risks faced by the patient, and this was substantially more so in the case of black patients. This created a huge racial disparity in the way healthcare was administered. Such algorithmic biases end up reinforcing[12] a vicious cycle of discrimination without appropriate checks and balances.

2. Standard frameworks for evaluating procurement of technological systems may focus on performance, cost, and quality considerations but not necessarily how to evaluate equity, fairness, and transparency

The key issue in procurement of an AI system is to consider how algorithmic governance poses differential costs to different groups of citizens. In other words, using AI in a function performed by the government might lower the overall costs of performing that function, but in the process, does AI create a hidden set of costs that place undue burdens on different groups of individuals?

The asymmetric information in any procurement transaction is that one party (the party that is offering the contract) does not know the behavior of the other party (the party that accepts a contract). Moral hazard occurs in an economic transaction when a party changes[13] its behavior to the detriment to the other. Protections against liability might encourage risky behavior, and it could be those individuals that are more risk-taking may be the ones who would opt for the more comprehensive coverage contracts (and willing to pay a premium for it). Buying insurance is more attractive for more risk-taking individuals (or in the case of health insurance, sicker individuals would benefit). Knowing they are protected against accidents might induce people to engage in more risky behavior. In the absence of actual information about borrower behavior or insurer behavior, buyers need to rely on behavioral proxies for risk. Measures such as deductibles and co-payments

[8] https://www.brookings.edu/research/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/

[9] https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G

[10] Eubanks, Virginia. *Automating Inequality: How High-tech Tools Profile, Police, and Punish the Poor.* First edition. New York, NY, St. Martin's Press, 2018.

[11] Obermeyer, Z., Powers, B., Vogeli, C. & Mullainathan, S. *Science* 336, 447–453 (2019).

[12] O'Neil, Cathy. 2017. Weapons of Math Destruction. Harlow, England: Penguin Books.

[13] https://en.wikipedia.org/wiki/Moral_hazard

are meant to protect insurers against some types of moral hazard. Such asymmetric information whereby each party to a contract lacks knowledge of the other results in a failure of efficient markets[14]. When buyers lack knowledge about the complex world of healthcare and given the difficult of assessing the quality of healthcare, we have a world where private contracts could fail. This could also be the case for procurement of complex artefacts such as AI-based systems. Proxies to assess the quality and functioning of AI and other complex technical systems do not consider the algorithmic harms and burdens on the disadvantaged. What is especially insidious is the use of black box training methods in building algorithms where we do not pay attention to how our methods of training end up reinforcing or amplifying societal biases. We may then be in a situation where procurement contracts for complex AI systems will pose negative externalities on the eventual users of these systems, i.e., citizens who are using services that rely on opaque AI systems.

3. We need comprehensive adoption of algorithmic auditing methods and frameworks to mitigate such biases.

In the absence of strong algorithmic accountability practices[15], we could have an appearance of scrutiny without genuine accountability[16]. We need frameworks to recognize harms of predictive processes[17]. Algorithmic auditing would require credentialing[18], standards of practice and extensive training. We also require comprehensive risk mitigation practices such as adopting[19] institutional review boards for AI and transparent disclosures highlighting the provenance of training data. The National Institute of Standards and Technology (NIST) has outlined a comprehensive AI risk management framework[20].

In summary, this committee is undertaking a critically important review of how the federal government acquires AI system. Legislators can support private and public adoption of the NIST risk management framework. That would have the effect of imposing accountability, similar to other regulations that mandate transparency and open governance.

---

[14] Arrow, Kenneth J. "Uncertainty and the Welfare Economics of Medical Care." *The American Economic Review* 53, no. 5 (1963): 941–73.

[15] Sasha Costanza-Chock, Inioluwa Deborah Raji, and Joy Buolamwini. 2022. Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22). Association for Computing Machinery, New York, NY, USA, 1571–1583

[16] Elizabeth Anne Watkins, Emanuel Moss, Jacob Metcalf, Ranjit Singh, and Madeleine Clare Elish. 2021. Governing Algorithmic Systems with Impact Assessments: Six Observations. In Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21). Association for Computing Machinery, New York, NY, USA, 1010–1022. https://doi.org/10.1145/3461702.3462580

[17] A few examples include Bryan Casey, Ashkon Farhangi, and Roland Vogl, "Rethinking Explainable Machines: The GDPR's 'Right to Explanation' Debate and the Rise of Algorithmic Audits in Enterprise," *Berkeley Technology Law Journal* 34, no. 1 (2019): 143–188; Timnit Gebru et al., "Datasheets for Datasets," *arXiv* 1803.09010 (2018), https://arxiv.org/abs/1803.09010; Margaret Mitchell et al., "Model Cards for Model Reporting," *Proceedings of the Conference on Fairness, Accountability, and Transparency* (2019): 220–29, http://doi.org/10.1145/3287560.3287596; Emanuel Moss et al., "Governing with Algorithmic Impact Assessments: Six Observations," *AAAI / ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)* (2020), https://dx.doi.org/10.2139/ssrn.3584818

[18] https://hbr.org/2018/11/why-we-need-to-audit-algorithms

[19] https://hbr.org/2021/04/if-your-company-uses-ai-it-needs-an-institutional-review-board

[20] https://www.nist.gov/itl/ai-risk-management-framework