

S. HRG. 115-649

DIGITAL DECISION-MAKING: THE BUILDING BLOCKS OF MACHINE LEARNING AND ARTIFICIAL INTELLIGENCE

HEARING

BEFORE THE

SUBCOMMITTEE ON COMMUNICATIONS,
TECHNOLOGY, INNOVATION, AND THE INTERNET
OF THE

COMMITTEE ON COMMERCE,
SCIENCE, AND TRANSPORTATION
UNITED STATES SENATE

ONE HUNDRED FIFTEENTH CONGRESS

FIRST SESSION

DECEMBER 12, 2017

Printed for the use of the Committee on Commerce, Science, and Transportation



Available online: <http://www.govinfo.gov>

U.S. GOVERNMENT PUBLISHING OFFICE

37-295 PDF

SENATE COMMITTEE ON COMMERCE, SCIENCE, AND TRANSPORTATION

ONE HUNDRED FIFTEENTH CONGRESS

FIRST SESSION

JOHN THUNE, South Dakota, *Chairman*

ROGER F. WICKER, Mississippi	BILL NELSON, Florida, <i>Ranking</i>
ROY BLUNT, Missouri	MARIA CANTWELL, Washington
TED CRUZ, Texas	AMY KLOBUCHAR, Minnesota
DEB FISCHER, Nebraska	RICHARD BLUMENTHAL, Connecticut
JERRY MORAN, Kansas	BRIAN SCHATZ, Hawaii
DAN SULLIVAN, Alaska	EDWARD MARKEY, Massachusetts
DEAN HELLER, Nevada	CORY BOOKER, New Jersey
JAMES INHOFE, Oklahoma	TOM UDALL, New Mexico
MIKE LEE, Utah	GARY PETERS, Michigan
RON JOHNSON, Wisconsin	TAMMY BALDWIN, Wisconsin
SHELLEY MOORE CAPITO, West Virginia	TAMMY DUCKWORTH, Illinois
CORY GARDNER, Colorado	MAGGIE HASSAN, New Hampshire
TODD YOUNG, Indiana	CATHERINE CORTEZ MASTO, Nevada

NICK ROSSI, *Staff Director*

ADRIAN ARNAKIS, *Deputy Staff Director*
JASON VAN BEEK, *General Counsel*
KIM LIPSKY, *Democratic Staff Director*
CHRIS DAY, *Democratic Deputy Staff Director*
RENAE BLACK, *Senior Counsel*

SUBCOMMITTEE ON COMMUNICATIONS, TECHNOLOGY, INNOVATION,
AND THE INTERNET

ROGER F. WICKER, Mississippi, <i>Chairman</i>	BRIAN SCHATZ, Hawaii, <i>Ranking</i>
ROY BLUNT, Missouri	MARIA CANTWELL, Washington
TED CRUZ, Texas	AMY KLOBUCHAR, Minnesota
DEB FISCHER, Nebraska	RICHARD BLUMENTHAL, Connecticut
JERRY MORAN, Kansas	EDWARD MARKEY, Massachusetts
DAN SULLIVAN, Alaska	CORY BOOKER, New Jersey
DEAN HELLER, Nevada	TOM UDALL, New Mexico
JAMES INHOFE, Oklahoma	GARY PETERS, Michigan
MIKE LEE, Utah	TAMMY BALDWIN, Wisconsin
RON JOHNSON, Wisconsin	TAMMY DUCKWORTH, Illinois
SHELLEY CAPITO, West Virginia	MAGGIE HASSAN, New Hampshire
CORY GARDNER, Colorado	CATHERINE CORTEZ MASTO, Nevada
TODD YOUNG, Indiana	

C O N T E N T S

	Page
Hearing held on December 12, 2017	1
Statement of Senator Wicker	1
Letter dated December 11, 2017 to Hon. Roger Wicker and Hon. Brian Schatz from Dean Garfield, President and CEO, Information Technology Industry Council (ITI)	83
Letter dated December 12, 2017 to Hon. John Thune and Hon. Bill Nelson from Marc Rotenberg, President, EPIC; Caitriona Fitzgerald, Policy Director, EPIC; and Christine Bannan, Policy Fellow, EPIC	84
Statement of Senator Schatz	2
Statement of Senator Moran	41
Statement of Senator Peters	44
Statement of Senator Udall	46
Statement of Senator Young	48
Statement of Senator Cantwell	50
Statement of Senator Markey	55
Statement of Senator Cruz	56
Statement of Senator Cortez Masto	58
Statement of Senator Blumenthal	60
 WITNESSES 	
Dr. Cindy L. Bethel, Associate Professor, Department of Computer Science and Engineering, Mississippi State University	4
Prepared statement	5
Daniel Castro, Vice President, Information Technology and Innovation Foundation (ITIF)	8
Prepared statement	10
Victoria Espinel, President and CEO, BSA The Software Alliance	17
Prepared statement	18
Report entitled “The \$1 Trillion Economic Impact of Software	63
Dr. Dario Gil, Ph.D., Vice President, AI and IBM Q	26
Prepared statement	27
Dr. Edward W. Felten, Ph.D., Robert E. Kahn Professor of Computer Science and Public Affairs, Princeton University	32
Prepared statement	34
 APPENDIX 	
Response to written questions submitted to Dr. Cindy L. Bethel by:	
Hon. Amy Klobuchar	87
Hon. Tom Udall	87
Hon. Maggie Hassan	88
Response to written questions submitted to Daniel Castro by:	
Hon. Tom Udall	90
Hon. Gary Peters	90
Hon. Maggie Hassan	91
Response to written questions submitted to Victoria Espinel by:	
Hon. Gary Peters	92
Hon. Maggie Hassan	94
Response to written questions submitted to Dr. Dario Gil, Ph.D. by:	
Hon. Amy Klobuchar	96
Hon. Tom Udall	97
Hon. Gary Peters	98
Hon. Maggie Hassan	99

IV

	Page
Response to written questions submitted to Dr. Edward W. Felten, Ph.D. by:	
Hon. Amy Klobuchar	102
Hon. Tom Udall	103
Hon. Gary Peters	103
Hon. Maggie Hassan	104

DIGITAL DECISION-MAKING: THE BUILDING BLOCKS OF MACHINE LEARNING AND ARTIFICIAL INTELLIGENCE

TUESDAY, DECEMBER 12, 2017

U.S. SENATE,
SUBCOMMITTEE ON COMMUNICATIONS, TECHNOLOGY,
INNOVATION, AND THE INTERNET,
COMMITTEE ON COMMERCE, SCIENCE, AND TRANSPORTATION,
Washington, DC.

The Subcommittee met, pursuant to notice, at 10 a.m. in room SR-253, Russell Senate Office Building, Hon. Roger Wicker, Chairman of the Subcommittee, presiding.

Present: Senators Wicker [presiding], Schatz, Blunt, Cruz, Fischer, Moran, Sullivan, Heller, Inhofe, Capito, Young, Cantwell, Klobuchar, Blumenthal, Markey, Booker, Udall, Peters, Hassan, and Cortez Masto.

OPENING STATEMENT OF HON. ROGER F. WICKER, U.S. SENATOR FROM MISSISSIPPI

Senator WICKER. This hearing will come to order. Senator Schatz will be here in a few moments and has sent word that we should go ahead and proceed. Today's Subcommittee meets to examine the commercial applications of artificial intelligence and machine learning for the U.S. economy. We are also gathered to discuss how the responsible design and deployment of intelligent systems can foster innovation and investment, propelling the United States as a leader in artificial intelligence.

I'm glad to convene this hearing, and as I mentioned, my colleague and friend Senator Schatz will be here in a moment.

Artificial intelligence refers to technology that is capable of taking on human-like intelligence. Through data inputs and algorithms, AI systems have the potential to learn, to reason, to plan, perceive, process, make decisions, and even act for themselves.

Although AI applications have been around for decades, recent advancements, particularly in machine learning, have accelerated in their capabilities because of the massive growth in data gathered from billions of connected devices and the digitization of everything. Developments in computer processing technologies and better algorithms are also enabling AI systems to become smarter and perform more unique tasks.

Every day consumers use technology that employ some degree of AI, smartphone mapping apps that suggest faster driving routes, for example. Online search platforms are learning from past que-

ries to generate increasingly customized results for users. We all know that when we click on a site on the Internet, news suggestions and advertisements on social media, and semi-autonomous vehicles are just a few examples of how machines and computer programs are taking on increasingly cognitive tasks.

The excitement surrounding this technology is deserved. AI has the potential to transform our economy, and so let's talk about that today. AI's ability to process and sort through troves of data can greatly inform human decisionmaking and processes across industries, including agriculture, health care, and transportation. In turn, businesses can be more productive, profitable, and efficient in their operations.

As AI systems mature and become more accurate in their descriptive, predictive, and prescriptive capabilities, there are issues that should be addressed to ensure the responsible development and use of this technology. Some of these issues include: understanding how data is gathered, what data is provided for an intelligent machine to analyze, and how algorithms are programmed by humans to make certain predictions. Moreover, understanding how the human end user interacts with or responds to the digital decision, and how humans interpret or explain decisions of the AI system over time will also need to be addressed.

These are important considerations to ensure that the decisions made by AI systems are based on representative data that does not unintentionally harm vulnerable populations or act in an unsafe, anticompetitive, or biased way. So there's a lot to think about.

In addition to these issues, other considerations, such as data privacy and cybersecurity, AI's impact on the workforce, and human control and oversight over intelligence systems should also be addressed as the technology develops. Fundamental to the success of machine learning and AI in enhancing U.S. productivity and empowering human decisionmaking is consumer confidence and trust in these systems. To build consumer confidence and trust, it is critical that the integration of AI into our commercial and government processes be done responsibly.

To that end, I look forward to learning from today's witnesses about how AI is advancing in our economy and what best practices our industry and AI researchers are considering to achieve all the promised economic and societal benefits of this technology.

Senator Schatz, do you have anything to add to my comprehensive opening statement?

**STATEMENT OF HON. BRIAN SCHATZ,
U.S. SENATOR FROM HAWAII**

Senator SCHATZ. I think everything has been said, but not everybody has said it. So good morning. Thank you very much, Mr. Chairman.

AI is advancing fast. Each year, processing power gets better, hardware gets cheaper, and algorithms are easier to train thanks to bigger and better datasets. These advances are, for companies and economies, great opportunities.

Technologists, historians, and economists say that we're at the cusp of the next industrial revolution, but there are concerns. We've seen that AI can be a black box. It can make decisions and

come to conclusions without showing its reasoning. There are also known cases of algorithms that discriminate against minority groups, and when you start to apply these systems to criminal justice, health care, or defense, the lack of transparency and accountability is worrisome.

Given the many concerns in a field that's advancing so quickly and is so revolutionary, it's hard to believe that there is no AI policy at the Federal level, and that needs to change. To start, the government should not purchase or use AI systems if we can't explain what it does, especially if these systems are making decisions about our citizens' lives.

There also needs to be more transparency and consumer control on data collection. Too many consumers still do not know what data is being collected, how it's being collected, or who owns it. Some of our current laws and regulations work, but some of them are too old and outdated to be used as a strong foundation for AI.

For example, companies often use data scraping to build their AI models. This falls under the Computer Fraud and Abuse Act, a 1986 law that was written before the Web was really in operation. AI is now used to write news articles, edit photographs, artificially reconstruct movies, all actions that fall under the Digital Millennium Copyright Act, which was passed in 1998, 10 years before the iPhone.

Our laws used to apply to actions in the physical world, but now they apply to software systems that ultimately do the same thing. I'm glad to see that industry and academia are being proactive by coming up with policy, principles, and professional ethics codes, but with this kind of patchwork, the system is only as strong as its weakest link.

From the private sector to academia to government, everyone has to wrestle with the ethical and policy questions that AI raises. And that's why I intend to introduce a bill creating an independent Federal commission to ensure that AI is adopted in the best interest of the public. If created, the commission would serve as a resource and coordinating body to the 16-plus Federal agencies that govern the use of AI. Otherwise, we risk bureaucratic turf wars and outdated inconsistent rules.

The commission would also be tasked with asking the tough ethical questions around AI. For instance, employers may not legally ask interviewees about their religion, marital status, or race, but they use software that mines social media data that may make the same inferences. Judges cannot base their sentencing decisions on whether the defendant has family members in jail, yet these facts contribute to a risk score calculated by machine learning algorithms.

In these cases, it's not clear yet where we draw the line on what is legal and what is not. In some instances, existing statutes will suffice. A lot of our laws actually work just fine when it comes to AI, and great laws can survive the test of time. But there are a few things that need to be wrestled with today in Congress and with our agencies, and that's why this hearing is so important.

Thank you very much. I look forward to hearing the testimony.
Senator WICKER. Thank you, Senator Schatz.

Our witnesses today are Dr. Cindy Bethel, Associate Professor, Department of Computer Science and Engineering at Mississippi State University; Mr. Daniel Castro, Vice President, Information Technology and Innovation Foundation; Ms. Victoria Espinel, Chief Executive Officer, BSA-The Software Alliance; Dr. Dario Gil, Vice President, IBM Research, AI, and IBM Q; and Dr. Edward Felten, Robert E. Kahn Professor of Computer Science and Public Affairs, Princeton University.

Friends, we will begin on my left with Dr. Bethel and proceed with 5-minute opening statements down the table. Thank you very much.

Dr. Bethel.

STATEMENT OF DR. CINDY L. BETHEL, ASSOCIATE PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING, MISSISSIPPI STATE UNIVERSITY

Dr. BETHEL. Good morning, Chairman Wicker, Ranking Member Schatz, and members of the Committee. Thank you for the opportunity to appear before you today. I'm Dr. Cindy Bethel. I'm Associate Professor of Computer Science and Engineering at Mississippi State University, or MSU.

It is an honor to speak with the Committee today about digital decisionmaking associated with artificial intelligence, known as AI, from the academic perspective, and about applications of AI being developed at MSU.

Today I will address three primary areas: first, a brief introduction to AI; next, three AI projects being developed at MSU; and last I will discuss some key points associated with AI.

A critical aspect associated with the advancement of science and technology is the development of algorithms with AI including machine learning for digital decisionmaking. In order for a system to make a decision, it must first acquire information, process and learn from that information, and then use that information to make decisions.

The gold standard would be the ability for a system to make a decision in a manner similar to a human expert in that area. This is a relatively new field of science that will be—provide ongoing research opportunities, including efforts to enhance existing algorithms and develop new and more efficient and effective algorithms. This is critical to the overall advancement of science in various disciplines, such as robotics, medicine, economics, and others.

There are countless application areas for which AI is beneficial and may make a significant impact on society. At MSU, we are on the forefront of AI developments with our research efforts, and today I will focus on three of those—excuse me—projects.

First is the integration of robots into high-risk law enforcement operations, such as SWAT teams, who I've trained with monthly for the last 6 years. This is an application for highly dynamic situations that require officers to make life-critical decisions. The more information they have prior to entering these dangerous and unknown environments, the better decisions the officers can make.

We are developing algorithms using robots into an environment prior to entry to provide critical audio and video intelligence to allow officers to make more informed decisions. This information

continues the dynamics of how they make entry or process a scene. Our research through intelligent interfaces will help inform how this information is delivered to the officers and maximize safety and performance.

Second, Mississippi State University is researching the development of autonomous cargo transport systems used in the fast-paced dynamic environment of a top 100 logistics company. This involves choosing the proper sensors to ensure that data is able to make—help make informed decisions on the vehicle's path.

Further, we are researching a variety of human factors because controlled vehicles will transfer from autonomous to human driver control and back again. Our work shows that if a human is not actively engaged in driving, they may—there may be insufficient situational awareness to control—take control on short notice. So how the human driver will be notified of the transfer of control is another critical aspect of our work.

To ensure maximum safety, researchers at MSU are exploring situations in which humans are operating in close proximity to the autonomous vehicles, such as a vehicle that's docking to deliver cargo, and a worker is unloading and loading this cargo.

Finally, another one of MSU's AI research projects involves a robotic therapy support system, known as Therabot that is in the form of a stuffed robotic dog. Therabot is an alternative to animal-assisted therapy for people who may be allergic to animals or may not be able to care for a live animal. Therabot will be used to provide support during clinical therapy sessions and home therapy practice with children or adults who are dealing with post-traumatic stress disorders or other mental health concerns. The algorithms being developed modify the autonomous behaviors of the robot based on the interactions with the human to accommodate his or her preferences and in response to different levels of stress detected by the robot, and to provide improved comfort and support.

AI will only be as good as the data the system receives and its ability to process that information to make a decision. The better the quality and quantity of information available to the system, the better the results will be from the machine learning process, which results in better final decisions from the system. Otherwise, decisionmaking capabilities can be limited or inaccurate.

The potential applications of AI are almost limitless. The United States can and should remain at the forefront of AI research, application development, and innovation provided that government proceeds with a light regulatory touch that doesn't stifle this potential.

Thank you so much for the opportunity to testify today on these important topics. I appreciate your time and attention to the advancement and impact of digital decisionmaking and AI.

[The prepared statement of Dr. Bethel follows:]

PREPARED STATEMENT OF DR. CINDY L. BETHEL, ASSOCIATE PROFESSOR OF COMPUTER SCIENCE AND ENGINEERING; DIRECTOR OF THE SOCIAL, THERAPEUTIC, AND ROBOTIC SYSTEMS (STARS) LAB; BILLIE J. BALL ENDOWED PROFESSOR IN ENGINEERING; MISSISSIPPI STATE UNIVERSITY

Chairman Wicker, Ranking Member Schatz, and Members of the Committee, thank you for the opportunity to appear before you today. I am Dr. Cindy Bethel, an Associate Professor of Computer Science and Engineering, the Billie J. Ball En-

dowed Professor in Engineering, and the Director of the Social, Therapeutic, and Robotic Systems (STaRS) Lab at Mississippi State University. It is an honor to speak with the committee today about digital decision-making associated with artificial intelligence (AI) from the academic perspective and about applications of AI being developed at Mississippi State University.

A critical aspect associated with the advancement of science and technology is the development of algorithms associated with AI including machine learning for digital decision-making. In order for a system to make a decision it must acquire information, have a means of processing and learning from that information, and have the ability to use that information to make an informed decision. The gold standard would be the ability for a machine or system to make a decision in a manner similar to that of a human who is an expert in that area. We have made considerable progress toward this goal since the inception of what is considered artificial intelligence, which began in 1943 with the research performed by McCulloch and Pitts. Today many machines and systems rely upon the use of artificial intelligence for digital decision-making. This is a relatively new field of science that will provide many lifetimes of research including efforts to enhance existing algorithms and develop new, more efficient, and effective algorithms. This is critical to the overall advancements of science in many disciplines, such as robotics, medicine, economics, and many others. There are often disagreements in the field as to what is considered AI and what algorithms and techniques used for learning are considered the best.

There are many application areas for which artificial intelligence is beneficial and may make a significant impact on society. At Mississippi State University, we are actively conducting AI research with several research projects that use AI and machine learning techniques, but I will focus on three primary projects.

The first project is the integration of robots into law enforcement operations, especially high risk, life critical incident responses such as those used with special weapons and tactics (SWAT) teams. I have been training monthly with rural regional SWAT team members since 2011. This is an example of an application in which high risk, dynamic situations are encountered, that require often officers to make life-critical decisions. The more information or intelligence they have prior to entering these dangerous and unknown environments, the better decisions the officers can make. We are investigating and developing algorithms related to computer vision, sensor fusion, and scene understanding to send a robot in prior to entry to provide audio and video feedback to officers during a response, highlighting what is critical information for them to attend to so that they are not overwhelmed with information when under high stress. The algorithms identify what is important to the officers in the environment, such as children, weapons, and other possible threats. This information can change the dynamics of how they make entry or process the scene. We are also researching in what ways this information needs to be provided to the officers, so that they can use it to their advantage to keep them and others safer in the performance of their duties. For example if officers are conducting a slow and methodical search of a building for a suspect, and the environment is quiet, dark and threat is nearby, they would not want to receive the information in a openly visual manner such as a video stream on a mobile phone that would highlight them in the environment and make them more likely to be the target of harm. In this case, they may want a verbal description of the scene that comes across on their radio earpiece. If they are in a gunfight or if there is an alarm that is sounding in the environment, but they are in a relatively “safe” location, then they may want to receive this information in a visual form, because audio transmission would be difficult to hear. We are researching the development of intelligent interface switching in which the manner that information is delivered to the officers may change depending on what is happening in the environment they are operating in. The officers are excited and ready to start deploying some of these artificial intelligence and machine learning applications in real-world responses.

A second project that we are working on at MSU is the development of autonomous cargo transport systems to be used in the fast-paced, dynamic environment of a top 100 logistics company. A primary factor that needs to be considered is what sensors need to be used to make informed decisions on the path the vehicles must travel. We also need to consider humans, who are sitting in the driver's seat of these vehicles, because control of the vehicles will change between fully autonomous to human driver operated. Research has shown that if the human is not actively involved in the activity of driving, that there may not be adequate situation awareness to be able to take back control of the vehicle with relatively short notice if needed. This has potential for life critical decision-making. We are investigating how the system needs to be able to alert the driver that control is being transferred, either to the vehicle or back to the human driver.

It is also important to consider the types of notifications that need to occur to ensure safety and situational awareness of what is happening around the vehicle. Also there are situations in which humans are operating in close proximity to the autonomous vehicle. As human drivers, we observe consciously or unconsciously behaviors of other drivers to infer what nearby vehicles will do next, but if there are not those cues, then how do humans in the environment understand what the vehicle or system will do next? This is a major issue of concern. This occurs also when the vehicle is docking to deliver cargo and a human is involved in unloading and loading this cargo. We are exploring methods of notification to the person about what the vehicle will do next. This is of significant concern, and if the incorrect decision is made, the vehicle could cause harm to the human.

The third project involves a robotic therapy support system, known as Therabot™ that is in the form of a stuffed robotic dog. Therabot™ is an alternative to animal-assisted therapy, for people who may be allergic to animals or may not be able to care for a live animal. Therabot™ will be used to provide support during clinical therapy sessions and for home therapy practice with children or adults who are dealing with post-traumatic stress disorders and other mental health concerns. The algorithms being developed for this project modify the autonomous behaviors and responses of the robot based on the interactions with the human to accommodate his or her preferences and in response to different levels of stress detected. Machine learning is being used to understand what behaviors the user prefers and to provide better support. It allows Therabot™ to be customizable to each individual user. It will learn and respond to each user as a dog would each person it encounters. Currently, Therabot™ can detect different types of touch, such as petting, patting, and hugging and will respond in different ways. If a person is under stress during the interaction he or she may squeeze the robot and the robot will adapt its behaviors to provide more comfort and support.

Algorithms used in the research and development of systems that are capable of digital decision-making are being developed and enhanced by researchers all across the world. Mississippi State University is at the forefront of these research developments and is continually contributing through publications and sharing of knowledge, algorithms, and software developments.

Research as a whole involves exploring what others have performed and then determining if there are modifications that can be made to improve upon those algorithms or the development of new algorithms to meet the needs of the application of use. For example, algorithms are being developed to learn a user's preference for how close a robot stands to them and still feel comfortable, or which friend they like most on a social media account. There are many methods for solving a problem. I typically tell my students who are interested in pursuing research that "you need to be first or best at something and it is always best to be first."

Artificial intelligence will only be as good as the data the system receives and its ability to process that information to make a decision. This is a critical aspect to the advancement of this field that can impact almost any other discipline. The system or machine must have the ability to perceive information and that typically comes from sensors and other forms of data. There are many types of sensing systems such as a camera, streaming video, thermal images based on the heat signature of items in the environment, radar, and many others. There are also methods of gathering information from sources such as social media, mobile device location history, purchase history, product preferences, websites visited, etc., that can assist in the decision-making process. The better the quality and quantity of information available to the system, the better the results will be from the machine learning process, which results in a better final decision from the system.

Algorithms are programmed to receive data as an input, process that data, learn from large amounts of data, and then use that information to make a digital decision. If there is not sufficient amounts of data available to train the machine learning algorithms or enough diversity in the data to allow the learning algorithms to adapt to different aspects, then the decision-making capabilities can be limited or inaccurate.

Another major issue of concern is the processing power necessary to handle large amounts of information received and come to a decision. This is especially a concern on smaller systems such as robots that have limited onboard processing capabilities. Many of the AI problems that are being addressed in the research community are performed on high powered computing resources and simulations are performed to validate the results. This is fine for many scientific applications, but in order for AI and machine learning to be beneficial in real-world applications, it will be necessary to perform the decision-making processes in real-time. The results need to be made available in an instant and not have to wait for processing time to provide

a result. This is improving, but sensing and processing are currently significant limitations to the application and use of AI in digital decision-making.

The level of human engagement necessary for digital decision-making depends on the state of the AI system. There are different levels of autonomous decision making. There is full autonomy, where the system receives the input and then processes the information from that data and makes a decision with no input from a human. There is supervised autonomy, in which the system receives information, processes the data, and comes up with possible results, and the human may have the ability to override or make the final decision. The more common level of autonomy is supervised autonomy. The level of human engagement also needs to consider the ramifications of the decision-making process. If it is a life-critical decision, then most people are more comfortable with a human remaining involved in the process to ensure an ethical and "good" decision is the final result. There are many ethical hurdles that will need to be decided at some point as to who is responsible if an AI system makes an incorrect decision, especially if the decision could result in harm to humans. There have been discussions in the field among researchers regarding who is responsible for this decision, such as the programmer, the company that made the system, or others. The current state often requires a human to be involved at some level of the final decision-making process unless it is low risk or well validated that the system will always make a "right" decision.

The fields of artificial intelligence and machine learning are at such an early stage of scientific development that standards and best practices are discussed among researchers; however, there is not a single standard or set of best practices that I am aware of that all researchers and scientists follow at this point. The biggest concern is providing the best possible answers that will not result in harm and provides benefits to the users.

The algorithms developed for machine learning and artificial intelligence can be used in almost any area of research, development, and discipline. It can be used to improve the decision-making process of humans. The processing of some information by a computer can be faster than what can be achieved by a human brain. Almost any aspect of society can benefit from the use of high quality artificial intelligence capabilities.

A critical aspect in the development of artificial intelligence using machine learning and other techniques is the impact on the humans that are involved. A current limitation to the advancement of artificial intelligence is the quality and cost effectiveness of sensing capabilities to provide high quality information or data to the system to make those digital decisions. Another critical limitation to current artificial intelligence capabilities is onboard processing capabilities and the cost effectiveness of those systems. We have come a long way in the advancement of artificial intelligence; however we still have a long way to go! The potential applications of AI are almost limitless. The United States can and should remain at the forefront of AI research, application development, and innovation provided that the government proceeds with a light regulatory touch that doesn't stifle this potential.

Thank you so much for the opportunity to testify today on these important topics. I appreciate your time and attention to the advancement and impacts of digital decision-making and artificial intelligence.

Senator WICKER. Thank you, Dr. Bethel. Precisely 5 minutes.

Dr. BETHEL. Thank you.

[Laughter.]

Senator WICKER. Mr. Castro, we're delighted to have you.

**STATEMENT OF DANIEL CASTRO, VICE PRESIDENT,
INFORMATION TECHNOLOGY AND INNOVATION FOUNDATION
(ITIF)**

Mr. CASTRO. Thank you. Chairman Wicker, Ranking Member Schatz, and members of the Committee, I appreciate the invitation to be here today.

AI has the potential to create a substantial and lasting impact on the economy by increasing the level of automation in virtually every sector, leading to more efficient processes and higher quality outputs, and boosting productivity in per capita incomes.

In the coming years, AI is expected to generate trillions of dollars of economic value and help businesses make smarter decisions, de-

velop innovative products and services, and boost productivity. For example, manufacturers are using AI to invent new metal alloys for 3D printing, pharmaceutical companies are using AI to discover new lifesaving drugs, and agricultural businesses are using AI to increase automation on farms.

Companies that use AI will have an enormous advantage compared to their peers that do not; therefore, the United States should prioritize policy initiatives that promote AI adoption in its traded sectors where U.S. firms will face international competition.

Many other countries already see the strategic importance of becoming lead adopters of AI, and they have begun implementing policies to pursue this goal. For example, this past March, Canada launched the Pan-Canadian AI Strategy, which is intended to help Canada become an international leader in AI research. The U.K.'s new budget, which was published last month, includes several provisions that have the goal of making the U.K. a world leader in AI, including establishing a new research center and funding about 500 Ph.D.'s. Japan has created an AI technology strategy designed to develop and commercialize AI in a number of fields, including transportation and health care. And China has declared its intent to be the world's premier AI innovation center by 2030.

However, to date, the U.S. Government has not declared its intent to be a global leader in this field, nor has it begun the even harder task of developing a strategy to achieve that vision. Moreover, China, which has launched this ambitious program to dominate the field, has already surpassed the United States in terms of the total number of papers published and cited in some AI disciplines, such as deep learning.

The U.S. should not cede its existing advantages in AI. Instead, it should pursue a multipronged national strategy to remain competitive in this field. First, the Federal Government should continue to expand its funding to support strategic areas of AI, especially in areas industry is unlikely to invest in, as well as better plan and coordinate Federal funding for AI R&D across different agencies.

Second, the Federal Government should support educational efforts to ensure a strong pipeline of talent to create the next generation of AI researchers and developers, including the retraining and diversity programs as well as pursue integration policies that allow U.S. businesses to recruit and retain highly skilled computer scientists.

Third, Federal and state regulators should conduct regulatory reviews to identify regulatory barriers to commercial use of AI in various industries, such as transportation, health care, education, and finance.

Fourth, the Federal Government should continue to supply high-value datasets that enable advances in AI, such as providing open access to standardized reference datasets for text analysis and facial recognition. Federal agencies should also facilitate data-sharing between industry stakeholders just as the Department of Transportation has done on safety for autonomous vehicles.

And, fifth, the Federal Government should assess what types of economic data it needs to gather from businesses to monitor and

evaluate AI adoption, much like it tracked rural electrification or broadband connectivity as key economic indicators.

Now, as with any technology, there will be some risks and challenges associated with AI that require government oversight, but the U.S. should not replicate the European approach to AI, where rules creating a right to explanation and a right to human review for automated decisions risk severely curtailing the uses of AI.

Instead, the U.S. should create its own innovation-friendly approach to providing oversight of the emerging algorithmic economy just as it has for the Internet economy. Such an approach should prioritize sector-specific policies over comprehensive regulation, outcomes over transparency, and enforcement actions against firms that cause tangible harm over those that merely make missteps without injury.

In many cases, regulators will not need to intervene because the private sector will address problems about AI, such as bias or discrimination, on its own. Moreover, given that U.S. companies are at the forefront of efforts to build AI that is safe and ethical, maintaining U.S. leadership in this field will be important to ensure these values remain embedded in this technology.

AI is a transformational technology that has the potential to significantly increase efficiency and innovation across the U.S. economy, creating higher living standards and improve quality of life. But while the United States has an early advantage in AI, many other countries are trying to be number one—they’re trying to be number one. We need more leadership on this issue. And I look forward to working with any member of the Committee on their proposed legislation and new ideas in this space. And I commend you all for holding this hearing.

Thank you for the opportunity to be here today. And I look forward to the questions.

[The prepared statement of Mr. Castro follows:]

PREPARED STATEMENT OF DANIEL CASTRO, VICE PRESIDENT,
INFORMATION TECHNOLOGY AND INNOVATION FOUNDATION (ITIF)

Introduction

Chairman Wicker, Ranking Member Schatz and members of the subcommittee, I appreciate the opportunity to appear before you to discuss the importance of artificial intelligence (AI) to the U.S. economy and how best to govern this important technology. My name is Daniel Castro, and I am vice president of the Information Technology and Innovation Foundation (ITIF), a non-profit, nonpartisan think tank whose mission is to formulate and promote public policies to advance technological innovation and productivity, and director of ITIF’s Center for Data Innovation.

What is Artificial Intelligence?

AI is a field of computer science devoted to creating computer systems that perform tasks much like a human would, particularly tasks involving learning and decision-making.¹ AI has many functions, including, but not limited to:

- Learning, which includes several approaches such as deep learning (for perceptual tasks), transfer learning, reinforcement learning, and combinations thereof;
- Understanding, or deep knowledge representation required for domain-specific tasks, such as medicine, accounting, and law;
- Reasoning, which comes in several varieties, such as deductive, inductive, temporal, probabilistic, and quantitative; and

¹ Daniel Castro and Joshua New, “The Promise of Artificial Intelligence,” Center for Data Innovation, October 2016, <http://www2.datainnovation.org/2016-promise-of-ai.pdf>.

- Interacting, with people or other machines to collaboratively perform tasks, and for interacting with the environment.

The cause of many misconceptions about AI, particularly its potential harms, is that some people conflate two very distinct types of AI: narrow AI and strong AI. Narrow AI describes computer systems adept at performing specific tasks, but only those specific types of tasks—somewhat like a technological savant.² For example, Apple's Siri virtual assistant is capable of interpreting voice commands, but the algorithms that power Siri cannot drive a car, predict weather patterns, or analyze medical records. While other algorithms exist that can accomplish those tasks, they too are narrowly constrained—the AI used for an autonomous vehicle will not be able predict a hurricane's trajectory or help doctors diagnose a patient with cancer.

In contrast, strong AI, also referred to as artificial general intelligence (AGI), is a hypothetical type of AI that can meet or exceed human-level intelligence and apply this problem-solving ability to any type of problem, just as the human brain can easily learn how to drive a car, cook food, and write code.³ Many of the dystopian fears about AI—that it will eliminate most jobs or go out of control and wipe out humanity, for example—stem from the notion that AGI is feasible, imminent, and uncontrollable.⁴ However, at least for the foreseeable future, computer systems that can fully mimic the human brain are only going to be found in scripts in Hollywood, and not labs in Silicon Valley.

The application of AI has seen a surge in recent years because of the development of machine learning—a branch of AI that focuses on designing algorithms that can automatically and iteratively build analytical models from data without needing a human to explicitly program the solution. Before machine learning, computer scientists had to manually code a wide array of functions into a system for it to mimic intelligence. But now developers can achieve the same, or better, results more quickly and at a lower cost using machine learning techniques. For example, Google uses machine learning to automatically translate content into different languages based on translated documents found online, a technique that has proven to be much more effective than prior attempts at language translation.⁵

What Are the Potential Benefits of AI?

AI will have a substantial and lasting impact on the economy by increasing the level of automation in virtually every sector, leading to more efficient processes and higher-quality outputs, and boosting productivity and per-capita incomes. For example, the McKinsey Global Institute estimates that by 2025 automating knowledge work with AI will generate between \$5.2 trillion and \$6.7 trillion of global economic value, advanced robotics relying on AI will generate between \$1.7 trillion and \$4.5 trillion, and autonomous and semi-autonomous vehicles will generate between \$0.2 trillion and \$1.9 trillion.⁶ Deloitte estimates that the Federal Government could save as much as \$41.1 billion annually by using AI to automate tasks.⁷ And Accenture predicts that by 2035, AI could increase the annual growth rate of the U.S. economy by 2 percentage points, the Japanese economy by 1.9, and the German economy by 1.6.⁸ The report also found that, for the 12 countries surveyed, AI would boost labor productivity rates by 11 to 37 percent.⁹

There are a vast and diverse array of uses for AI, and many U.S. businesses are already using the technology today. Manufacturers are using AI to invent new metal alloys for 3D printing; pharmaceutical companies are using AI to discover new life-saving drugs; mining companies are using AI to predict the location of mineral deposits; and agricultural businesses are using AI to increase automation on farms.

² Irving Wladawsky-Berger, “Soft Artificial Intelligence Is Suddenly Everywhere,” *The Wall Street Journal*, January 16, 2016, <http://blogs.wsj.com/cio/2015/01/16/soft-artificial-intelligence-is-suddenly-everywhere/>.

³Ibid.

⁴ Robert D. Atkinson, “It’s Going to Kill Us!” and Other Myths About the Future of Artificial Intelligence,” (Information Technology and Innovation Foundation, June 2016), http://www2.itif.org/2016-myths-machine-learning.pdf?_ga=1.201838291.334601971.1460947053.

⁵ Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (New York: Basic Books, 2015).

⁶ James Manyika et al., Disruptive Technologies: Advances That Will Transform Life, Business, and the Global Economy, (McKinsey Global Institute, May 2013), <http://www.mckinsey.com/business-functions/business-technology/our-insights/disruptive-technologies>.

⁷ Peter Viechnicki and William D. Eggers, “How much time and money can AI save government?” (Deloitte, April 26, 2017), <https://dupress.deloitte.com/dup-us-en/focus/cognitive-technologies/artificial-intelligence-government-analysis.html>.

⁸ Mark Purdy and Paul Daugherty, “Why Artificial Intelligence Is the Future of Growth,” (Accenture, September 28, 2016), https://www.accenture.com/us-en/_acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth.pdf.

⁹Ibid.

The International Data Corporation (IDC) estimates that the market for AI technologies that analyze unstructured data will reach \$40 billion by 2020.¹⁰ And AI startups have attracted significant investment, with U.S. investors putting \$757 million in venture capital in AI start-ups in 2013, \$2.18 billion in 2014, and \$2.39 billion in 2015.¹¹

In some cases, the principle benefit of AI is that it automates work that would otherwise need to be performed by a human, thereby boosting efficiency. Sometimes AI can complete tasks that it is not always worth paying a human to do but still creates value, such as writing newspaper articles to summarize Little League games.¹² In other cases, AI adds a layer of analytics that uncovers insights human workers would be incapable of providing on their own, thereby boosting quality. In some cases, it does both. For example, researchers at Stanford have used machine learning techniques to develop software that can analyze lung tissue biopsies with significantly more accuracy than a top human pathologist and at a much faster rate.¹³ By analyzing large volumes of data, researchers can train their computer models to reliably recognize known indicators of specific cancer types as well as discover new predictors.

AI is also delivering valuable social benefits, such as by helping authorities rapidly analyze the deep web to crack down on human trafficking, fighting bullying and harassment online, helping development organizations better target impoverished areas, reducing the influence of gender bias in hiring decisions, and more.¹⁴ Just as AI can help businesses make smarter decisions, develop innovative new products and services, and boost productivity to drive economic value, it can achieve similar results for organizations generating social value, and many of these solutions have the potential to scale globally.

Finally, AI will be an increasingly important technology for defense and national security. AI can address many goals, such as improving logistics, detecting and responding to cybersecurity incidents, and analyzing the enormous volume of data produced on the battlefield. Moreover, AI will be a core enabler of the Pentagon's "Third Offset Strategy," a policy designed to keep the United States ahead of adversaries, especially ones capable of fielding numerically superior forces, through technological superiority.¹⁵ Indeed, one top Pentagon general has suggested that the Defense Department should never buy another weapons system that does not have AI built into it.¹⁶

How Should Policymakers Support the Adoption and Use of AI?

Given the potential economic impact of AI in raising productivity, policymakers should develop a national strategy to support the development and adoption of AI in U.S. businesses. In particular, given the enormous advantage that AI-enabled firms will have compared to their non-AI-enabled peers, the United States should focus on AI adoption in its traded sectors where U.S. firms will face international competition. Many other countries see the strategic importance of becoming lead adopters of AI, and they have begun implementing policies to pursue this goal. These include:

¹⁰ "Cognitive Systems Accelerate Competitive Advantage," IDC, accessed September 29, 2016, <http://www.idc.com/promo/thirdplatform/innovationaccelerators/cognitive>.

¹¹ "Artificial Intelligence Exploses: New Deal Activity Record for AI Startups," CB Insights, June 20, 2016, <https://www.cbinsights.com/blog/artificial-intelligence-funding-trends/>.

¹² Steven Levy, "Can an Algorithm Write a Better News Story Than a Human Reporter?" *Wired*, April 24, 2012, <https://www.wired.com/2012/04/can-an-algorithm-write-a-better-news-story-than-a-human-reporter/>.

¹³ Kun-Hsing Yu et al., "Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features," *Nature*, August 16, 2017, <https://www.nature.com/articles/neuroms12474>.

¹⁴ Larry Greenemeier, "Human Traffickers Caught on Hidden Internet," *Scientific American*, February 8, 2015 <http://www.scientificamerican.com/article/human-traffickers-caught-on-hidden-internet/>; Davey Alba, "Weeding Out Online Bullying Is Tough, So Let Machines Do It," *Wired*, July 10, 2015, <https://www.wired.com/2015/07/weeding-online-bullying-tough-let-machines/>; Michelle Horton, "Stanford Scientists Combine Satellite Data, Machine Learning to Map Poverty," *Stanford News*, August 18, 2016 <http://news.stanford.edu/2016/08/18/combining-satellite-data-machine-learning-to-map-poverty/>; Sean Captain, "How Artificial Intelligence is Finding Gender Bias at Work," *Fast Company*, October 10, 2015, <https://www.fastcompany.com/3052053/elasticity/how-artificial-intelligence-is-finding-gender-bias-at-work>.

¹⁵ Sydney Freedberg, "Faster Than Thought: DARPA, Artificial Intelligence, & The Third Offset Strategy," *Breaking Defense*, February 11, 2016, <https://breakingdefense.com/2016/02/faster-than-thought-darpa-artificial-intelligence-the-third-offset-strategy/>.

¹⁶ Jack Corrigan, "Three-Star General Wants Artificial Intelligence in Every New Weapon System," *Nextgov*, November 2, 2017, <http://www.nextgov.com/cio-briefing/2017/11/three-star-general-wants-artificial-intelligence-every-new-weapon-system/142225/>.

- Canada: In March 2017, Canada launched the Pan-Canadian Artificial Intelligence Strategy which sets a goal of establishing Canada as an international leader in AI research. The strategy has four goals, which include increasing the number of AI researchers and graduates; establishing three major AI research centers; developing global thought leadership on the economic, ethical, policy and legal implications of advances in AI; and supporting the national AI research community.¹⁷
- China: China's State Council issued a development plan for AI in July 2017 with the goal of making China a leader in the field by 2030. China's goal is to be equal to countries currently leading in AI by 2020. Then, over the subsequent five years, China will focus on developing breakthroughs in areas of AI that will be a "a key impetus for economic transformation."¹⁸ Finally, by 2030, China intends to be the world's "premier artificial intelligence innovation center."¹⁹ China's plan also signals its intent to require high school students to take classes in AI, which is one of the most ambitious efforts to develop human capital for the AI economy of any nation.
- Japan: Prime Minister Abe launched the Artificial Intelligence Technology Strategy Council in April 2016 to develop a roadmap for the development and commercialization of AI.²⁰ Published in May 2017, the roadmap outlines priority areas for research and development (R&D), focusing on the themes of productivity, mobility, and health. The strategy also encourages collaboration between industry, government, and academia to advance AI research, as well as stresses the need for Japan to develop the necessary human capital to work with AI. Japan also launched its Japan Revitalization Strategy 2017, which details how the government will work to support growth in certain areas of the economy. The 2017 strategy includes a push to promote the development of AI for telemedicine and self-driving vehicles to address the shortage of workers in Japan.
- UK: The United Kingdom has taken several steps to promote AI. The UK Digital Strategy, published in March 2017, recognizes AI as a key field that can help grow the United Kingdom's digital economy.²¹ The UK's new budget, published in November 2017, includes several provisions that have the goal of establishing the UK as a world leader in AI, such as by establishing a "Centre for Data Ethics and Innovation" to promote the growth of AI, facilitating data access for AI through "data trusts," and funding 450 PhD researchers working on AI.²²

While the U.S. Government has put significant funding behind AI R&D—approximately \$1.1 billion in 2015—it has not done enough to maintain U.S. leadership.²³ The most ambitious AI program comes from China, which as of 2014 surpassed the United States in terms of total number of papers published and cited in AI fields, such as deep learning.²⁴ For both economic and national security reasons, the United States cannot afford to cede its existing advantages in AI, and should instead look to capitalize on its head start by developing a strategy to support AI development and adoption. Such a strategy should include policies to address the following:

- Funding: The government should continue to expand its funding to support the "National Artificial Intelligence Research and Development Strategic Plan," a set of R&D priorities identified by the Networking and Information Technology

¹⁷ "Pan-Canadian Artificial Intelligence Strategy Overview," Canadian Institute for Advanced Research, March 3, 2017, <https://www.cifar.ca/assets/pan-canadian-artificial-intelligence-strategy-overview/>.

¹⁸ Graham Webster *et al.*, "China's Plan to 'Lead' in AI: Purpose, Prospects, and Problems," New America Foundation, August 1, 2017, <https://www.newamerica.org/cybersecurity-initiative/blog/chinas-plan-lead-ai-purpose-prospects-and-problems/>.

¹⁹ Ibid.

²⁰ Josh New, "How Governments Are Preparing for Artificial Intelligence," August 8, 2017, <https://www.datainnovation.org/2017/08/how-governments-are-preparing-for-artificial-intelligence/>.

²¹ Department of Digital, Culture, Media, and Sport, *UK Digital Strategy*, (United Kingdom: Department for Digital, Culture, Media, and Sport, 2017), <https://www.gov.uk/government/publications/uk-digital-strategy>.

²² Her Majesty's Treasury (HM Treasury), *Autumn Budget 2017* (United Kingdom: HM Treasury, 2017), <https://www.gov.uk/government/publications/autumn-budget-2017-documents/autumn-budget-2017>.

²³ Ibid.

²⁴ "National Artificial Intelligence Research and Development Strategic Plan," (National Science and Technology Council, October 2016), https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf.

Research and Development (NITRD) program that addresses strategic areas of AI in which industry is unlikely to invest, as well as better plan and coordinate Federal funding for AI R&D across different agencies.²⁵

- Skills: The Federal Government should support educational efforts to ensure a strong pipeline of talent to create the next generation of AI researchers and developers, including through retraining and diversity programs, as well as pursue immigration policies that allow U.S. businesses to recruit and retain highly skilled computer scientists.
- AI-Friendly Regulations: Federal and state regulators should conduct regulatory reviews to identify regulatory barriers to commercial use of AI in various industries, such as transportation, health care, education, and finance.
- Data Sharing: Some advances in AI are made possible when large volumes of accurate and representative data are made part of a data commons. The government should continue to supply high-value datasets that enable advances in AI, such as its efforts to produce standardize reference datasets for text analysis and facial recognition. Similarly, Federal agencies should facilitate data sharing between industry stakeholders, such as the Department of Transportation's draft "Guiding Principles on Data Exchanges to Accelerate Safe Deployment of Automated Vehicles."²⁶
- Economic Indicators: Understanding the degree to which U.S. firms have automated processes using AI will be a key metric to assessing the effectiveness of various policies. The Census Bureau should assess what type of economic data it should gather from businesses to monitor and evaluate AI adoption, much like it has tracked rural electrification or broadband connectivity as key economic indicators.

How Should Policymakers Address Concerns About Workforce Disruption?

One of the most common fears about AI is that it will lead to significant disruptions in the workforce.²⁷ This fear is not new—concerns about technology-driven automation have been a perennial policy concern since at least the 1930s when Congress debated legislation that would direct the Secretary of Labor to make a list of all labor-saving devices and estimate how many people could be employed if these devices were eliminated.²⁸ This concern has been exacerbated by a frequently-cited study by two Oxford academics which predicted that 47 percent of U.S. jobs could be eliminated over the next 20 years.²⁹

This study's predictions are misleading and unlikely for at least three reasons. First, the estimate includes a number of occupations that have little chance of automation, such as fashion models and barbers. Second, while this rate of productivity seems high and even threatening, it is only slightly higher than rates enjoyed in the mid-1990s when U.S. job creation was robust and unemployment rates low. Third, it succumbs to what economists call the "lump of labor" fallacy which holds that once a job is gone, there are no other jobs to replace it. The reality is that AI-driven productivity enables organizations to either raise wages or reduce prices. These changes lead to increases in spending, which in turn creates more jobs. And given that consumers' wants are far from satisfied, there is no reason to believe that this dynamic will change anytime soon.

But while predictions about massive AI-driven unemployment are vastly overstated—indeed, by historical standards occupational churn, the rate at which some jobs expand while others contract, is at its lowest levels in 165 years—there will still be some worker displacement as AI creates higher levels of productivity.³⁰ So policymakers can and should do more to help workers make transitions between jobs and occupations, such as by providing strong social safety net programs, re-

²⁵ Ibid.

²⁶ "Draft U.S. DOT Guiding Principles on Voluntary Data Exchanges to Accelerate Safe Deployment of Automated Vehicles," (U.S. Department of Transportation, December 1, 2017) <https://www.transportation.gov/av/data>.

²⁷ For a thorough rebuttal of this concern, see Robert D. Atkinson, "It's Going to Kill Us!" And Other Myths of Artificial Intelligence," (Information Technology and Innovation Foundation, June 2016), <http://www2.itif.org/2016-myths-machine-learning.pdf>.

²⁸ John Scoville, "Technology and the Volume of Employment," Proceedings of the Academy of Political Science 18, no. 1 (May 1938): 84–99.

²⁹ Carl B. Frey and Michael A. Osbourne, "The Future of Employment: How Susceptible Are Jobs to Computerisation?" (Oxford Martin School, University of Oxford, Oxford, September 17, 2013), http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf.

³⁰ Robert D. Atkinson and John Wu, "False Alarmism: Technological Disruption and the U.S. Labor Market, 1850–2015," (Information Technology and Innovation Foundation, May 2017), <http://www2.itif.org/2017-false-alarmism-technological-disruption.pdf>.

forming unemployment insurance, and offering worker retraining. The failure to give workers training and assistance to move into new jobs or occupations not only contributes to higher structural unemployment, but also increases resistance to innovation and automation.³¹

How Should Policymakers Provide Oversight of AI?

When it comes to AI, the primary goal of the United States should be to accelerate the development and adoption of the technology. But as with any technology, there will be some risks and challenges that require government oversight. The presence of risk, however, does not mean that the United States should embrace the precautionary principle, which holds that new technology must first be proven safe before it can be used. Instead, policymakers should rely on the innovation principle, which says that policymakers should address risks as they arise, or allow market forces to address them, and not hold back progress because of speculative concerns. The innovation principle is especially useful when fears about a new technology exceed public awareness and understanding about how the technology works and how potential problems will be mitigated.³²

To understand why this is important, consider the differences between the United States and the European Union in the Internet economy. Compared to Europe, the United States has had more success in the Internet economy, at least in part, because of its vastly more simplified data protection regulations. Yet even as the United States continues to produce the majority of the major global Internet companies, the European Union has decided to double down on its onerous data protection rules in the forthcoming General Data Protection Regulation (GDPR), a far-reaching set of policies that will substantially raise the costs, and in some cases, limit the feasibility of using AI in Europe. For example, the GDPR creates both a right to explanation and a right to human review for automated decisions, two requirements that will make it difficult for companies to construct business models that rely extensively on complex algorithms to automate consumer-facing decisions. The GDPR also requires organizations to only use data for the purposes for which they originally collected it, a rule that strictly limits the application of AI to existing data.³³ If the United States wants to compete for global leadership in AI, it should be careful not to follow Europe down this path.

While the United States should not replicate the European model, it should create its own innovation-friendly approach to providing oversight of the emerging algorithmic economy just as it has for the Internet economy. Such an approach should prioritize sector-specific policies over comprehensive regulations, outcomes over transparency, and enforcement actions against firms that cause tangible harm over those that merely make missteps without injury. For example, rather than industry-wide rules requiring “algorithmic transparency” or “AI ethics”—proposals that focus on means, rather than ends—policymakers should look to address specific problems, such as ensuring financial regulators have the skills necessary to provide oversight of fintech companies relying heavily on AI to make lending decisions or provide automated financial advisors.

In many cases, regulators will not need to intervene because the private sector will address problems about AI, such as bias or discrimination, on its own—even if to outsiders an algorithm appears to be a “black box.” After all, one company’s hidden biases are another company’s business opportunities. For example, if certain lenders were to use algorithms that consistently denied loans to ethnic or religious minorities who have good credit, then their competitors would have an incentive to target these individuals to gain new customers.

Moreover, the private sector is actively seeking out solutions to eliminate problems like unintentional bias in AI that may skew its results.³⁴ For example, a group of leading AI companies in the United States have formed an association to develop and share best practices to ensure that AI is fair, safe, and reliable, while another technology trade association has publicly committed itself to ensuring that the pri-

³¹ See forthcoming report: “Technological Innovation, Employment, and Workforce Adjustment Policies,” (Information Technology and Innovation Foundation, January 2018).

³² Daniel Castro and Alan McQuinn, “The Privacy Panic Cycle: A Guide to Public Fears About New Technologies,” (Information Technology and Innovation Foundation, September 2015), <http://www2.itif.org/2015-privacy-panic.pdf>.

³³ Nick Wallace, “UK Regulations Need an Update to Make Way for Medical AI,” Center for Data Innovation, August 12, 2017, <http://datainnovation.org/2017/08/uk-regulations-need-an-update-to-make-way-for-medical-ai/>.

³⁴ Cliff Kuang, “Can A.I. Be Taught to Explain Itself?” *New York Times*, November 21, 2017, <https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html>.

vate sector designs and uses AI responsibly.³⁵ Indeed, given that U.S. companies are at the forefront of efforts to build AI that is safe and ethical, maintaining U.S. leadership in this field will be important to ensure these values remain embedded in the technology.

But policymakers should be careful not to misclassify certain concerns as “AI problems” that would be best dealt with on a technology-neutral basis. For example, discrimination in areas such as access to financial services and housing are best addressed through existing legal mechanisms. No new laws and regulations are needed simply because a company uses AI, instead of human workers, to make certain decisions.³⁶ Companies cannot use AI to circumvent laws outlawing discrimination.

Finally, certain problems, such as sexism in hiring practices, are not necessarily made worse by AI. On the contrary, using AI can actually reduce human biases. For example, companies can use AI to police undesirable behaviors, like automatically flagging job advertisements that use gender-specific terminology, such as “waitress” instead of “wait staff,” or stereotypical images, such as a female nurse.³⁷ And unlike human processes, where it may take years or decades to change social norms and company culture, businesses can refine and tweak code over a period of days or weeks. For example, Google changes its search engine 500 to 600 times per year.³⁸ Thus companies will likely have more success eliminating bias when it appears in AI, than when it appears elsewhere in society.

Conclusion

AI is a transformational technology that has the potential to significantly increase efficiency and innovation across the U.S. economy, creating higher living standards and improved quality of life. But while the United States has an early advantage in AI given its top talent in computer science and deep bench of companies, large and small, investing in the field, many other countries are actively vying to challenge U.S. leadership in this domain. In particular, China, with its highly skilled computer science workforce and significant funding for AI R&D, could easily catch and surpass the United States, leading to it gaining economic and military advantages.

Unfortunately, U.S. policy debates about AI too often overemphasize the potential impact on worker displacement from automation or bias from algorithms and ignore the much more pressing concern about the potential loss of competitiveness and defense superiority if the United States falls behind in developing and adopting this key technology.

Yet, when it comes to AI, successfully integrating this technology into U.S. industries should be the primary goal of policymakers, and given the rapid pace at which other countries are pursuing this goal, the United States cannot afford to rest on its laurels. To date, the U.S. Government has not declared its intent to remain globally dominant in this field, nor has it begun the even harder task of developing a strategy to achieve that vision. Some may think this is unnecessary, believing that the United States will automatically prevail in this technology simply because it has a unique culture of innovation and has prevailed on past technologies.³⁹ Such views are naïve and dangerous, and if followed, likely will lead to the United States being surpassed as the global leader in AI with significant negative consequences for the U.S. economy and society. However, it is not too late to begin to ensure continued U.S. leadership, and I commend you for holding this hearing so that we can have this conversation.

Senator WICKER. Thank you, Mr. Castro.
Ms. Espinel.

³⁵ See “Partnership on AI,” <https://www.partnershiponai.org/> and “AI Policy Principles,” Information Technology Industry Council, <https://www.itic.org/resources/AI-Policy-Principles-FullReport2.pdf>.

³⁶ Travis Korte and Daniel Castro, “Disparate Impact Analysis is Key to Ensuring Fairness in the Age of the Algorithm,” Center for Data Innovation, January 20, 2015, <http://datainnovation.org/2015/01/disparate-impact-analysis-is-key-to-ensuring-fairness-in-the-age-of-the-algorithm/>.

³⁷ Amber Laxton, “Critics of ‘Sexist Algorithms’ Mistake Symptoms for Illness,” Center for Data Innovation, August 3, 2015, <http://datainnovation.org/2015/08/critics-of-sexist-algorithms-mistake-symptoms-for-illness/>.

³⁸ Daniel Castro, “Data Detractors Are Wrong: The Rise of Algorithms Is a Cause for Hope and Optimism,” Center for Data Innovation, October 25, 2016, <http://datainnovation.org/2016/10/data-detectors-are-wrong-the-rise-of-algorithms-is-a-cause-for-hope-and-optimism/>.

³⁹ Patrick Tucker, “What the CIA’s Tech Director Wants from AI,” Defense One, September 6, 2017, <http://www.defenseone.com/technology/2017/09/cia-technology-director-artificial-intelligence/140801/>.

**STATEMENT OF VICTORIA ESPINEL, PRESIDENT AND CEO,
BSA | THE SOFTWARE ALLIANCE**

Ms. ESPINEL. Good morning, Chairman Wicker, Ranking Member Schatz, and members of the Subcommittee. My name is Victoria Espinel, and I am the President and CEO of BSA | The Software Alliance.

BSA is the advocate for the global software industry in the United States and around the world. Our members are at the forefront of developing artificial intelligence and related software services. I commend the Subcommittee for a hearing on this important topic, and I thank you for the opportunity to testify.

At the outset, I think it's important to answer a key question: What is AI? So let me provide you with a brief anecdote. A 60-year-old woman was initially diagnosed with a conventional form of leukemia. She went through chemotherapy to treat the disease, but her recovery was unusually slow. Conventional tests failed to reveal a problem, but her doctor suspected that something was still wrong. After several frustrating months, they turned to an AI-powered, cloud-based system capable of cross-referencing the patient's genetic data with insights gleaned from tens of millions of studies from around the world. Within minutes, the doctors learned that the patient might be suffering from an extremely rare form of leukemia that required a unique course of treatment. The doctors were able to quickly update her treatment plan and watch her condition improve significantly.

This is AI: it's innovative; it's powerful; it's lifesaving. AI is not the image that we see in science fiction movies of robots demolishing tall buildings; instead, the AI provided by BSA members today is a tool that uses data to help people solve complex problems, simplify our daily lives, improve business operations, and enhance government services.

AI is powered by software, which is itself a major engine of economic growth. The software industry contributed more than \$1.14 trillion to the U.S. GDP in 2016—a \$70 billion increase in just 2 years. The software industry is a powerful job creator supporting over 10.5 million jobs with a significant positive impact on jobs and economic growth in every one of the 50 states.

For example, in Mississippi, software is contributing over \$800 million to its GDP and over 7,000 jobs, a 25 percent increase in jobs in just 2 years. Over 4,000 miles away in Hawaii, software is contributing over \$1 billion to its GDP and over 16,000 jobs. Across every single state in the country, the economic impact of software is on the rise.

AI is helping all industry sectors. Whether it is securing networks, improving health, or helping American farmers save money, the impact of AI is already visible in every industry, in every state, and across the globe.

We should also be prepared to address important issues that may arise as AI-enabled services are used. Let me focus on two. First, AI will change the skill sets needed for certain jobs, and while new AI-related jobs will be created, there will be shifts in the economy. BSA members are already helping launching groundbreaking initiatives to provide free training, including to youth and military veterans, to ensure that both the current workforce and the next

generation are prepared for the future. We are dedicated to this work, and we look forward to collaborating with all of you on this effort.

Second, we are mindful of the need to ensure that AI is both trained and used fairly and responsibly. At the same time, we recognize the potential of AI to make human decisions more accurate and less biased and the need to push toward that outcome.

As our companies seek to ensure responsible AI deployment, there are several steps that Congress and the administration can take. First, as I highlighted earlier, AI depends on data, so we urge Congress to pass the Open Government Data Act, which would make non-sensitive government data more open, more available, and more usable for the general public.

Ranking Member Schatz, thank you for your great work as sponsor of the Open Government Data Act. We hope that Congress will act soon to send it to the President's desk. We also encourage Congress and the administration to be leaders on digital trade to encourage global data flows.

Second, we encourage increased investment in government research, including on how AI can contribute to both positive economic and social outcomes and policies that incentivize private sector research and development.

And, third, we need to prioritize education and workforce development so that our young people and our current workforce are prepared for the future.

As part of all of this, we need to have a meaningful dialogue with all stakeholders about how to address any challenges that lie ahead. The legislation introduced by Senators Cantwell, Young, and Markey is a good step, and we thank you for that. Thanks to Senator Schatz as well for the legislation that you are currently working on.

In closing, we look forward to working with all of you towards a clear understanding of AI and to address the challenges and embrace the opportunities ahead. BSA members are part of the solution to these challenges, and we are eager to work with you as we chart a responsible path forward.

Thank you, and I look forward to your questions.
[The prepared statement of Ms. Espinel follows:]

PREPARED STATEMENT OF VICTORIA ESPINEL, PRESIDENT AND CEO,
BSA—THE SOFTWARE ALLIANCE

Good morning Chairman Wicker, Ranking Member Schatz, and members of the Subcommittee. My name is Victoria Espinel, and I am the President and CEO of BSA | The Software Alliance.

BSA is the leading advocate for the global software industry in the United States and around the world.¹ Our members are at the forefront of developing cutting-edge artificial intelligence (AI) and related software-enabled technologies and services that are having a significant impact on the U.S. and global economy. I commend the Subcommittee for holding a hearing on this important topic, and I thank you for the opportunity to testify on behalf of BSA.

¹ BSA's members include: Adobe, ANSYS, Apple, Autodesk, Bentley Systems, CA Technologies, CNC/Mastercam, DataStax, DocuSign, IBM, Microsoft, Oracle, salesforce.com, SAS Institute, Siemens PLM Software, Splunk, Symantec, Trimble Solutions Corporation, The MathWorks, Trend Micro and Workday.

I. AI: Defining the Landscape

The term “artificial intelligence” often conjures images of all-knowing robots with physical and cognitive abilities far superior to those of their human creators. The actual AI services that are in the market today—and that BSA members provide—bear no resemblance to the sinister images of the future that consumers often see in the movies, with robots taking over big cities and small towns.

Instead, they increasingly are becoming a foundational technology that drives many products and services that people use every day. Whether it is a personal digital assistant that helps consumers locate the nearest restaurant, a fraud detection monitoring service that prevents criminals from placing charges on credit cards, or a tool that helps teachers identify students with additional needs and develop personalized lesson plans, we increasingly rely on a diverse range of AI-enabled services every day.

But what is “AI”?

Although definitions of AI vary, one common description of AI is that it refers to machines that act intelligently in pursuit of human-defined objectives. At its core, AI is simply a tool. It includes a broad range of technologies, but the AI systems that BSA members largely provide assist in the analysis of enormous volumes of data to find connections that improve the quality and accuracy of human decision-making. Although some AI systems have a limited degree of autonomy, such as submarines that map the ocean bed and measure ocean currents, and others are minutely supervised, such as robot surgical tools assisting doctors with hip replacement surgeries, the vast majority provide advice and recommendations to humans rather than acting independently. AI makes possible important tasks that would otherwise be economically or physically infeasible, such as inspecting wind turbine blades or the interior of oil pipelines.

AI systems, like other software systems, use sophisticated algorithms. An algorithm is a set of instructions that processes various inputs and provides an output in a systematized way. The algorithms used in AI are particularly well-suited to analyzing massive volumes of data from many different sources, and in identifying patterns across the enormous number of variables in such data that may interact in complex and unexpected ways. Through this analysis, AI systems can enhance perception, learning, reasoning, and decision-making, and improve the ability of people to solve complex and challenging problems.

The use of systems, including software, to help people solve complex problems is not new. Research into AI dates back many decades, but we have witnessed tremendous advances in AI capabilities over the past five to ten years. These advances have been fueled by a number of related developments, including the proliferation of technologies that generate vast amounts of data, the affordability of data storage, and ever-growing data processing capabilities.

BSA members have made significant investments in enhancing these data-driven technologies to develop innovative AI solutions for use across a broad range of applications in a wide variety of contexts.

II. AI Services Provide Substantial Benefits

Advances in AI and software-enabled data analytics are fueling job and economic growth in the United States and around the world, improving how businesses in every sector operate, and producing real societal gains. We must recognize that AI will change the skill sets needed for certain jobs. And while new, AI-related jobs will be created, there will be shifts in the labor market. And although we should be mindful of the need to ensure that AI is deployed fairly and responsibly, we should also recognize the potential of AI to make human decisions more accurate and less biased, and thereby to promote fairness and inclusiveness across all segments of society.

A. AI and Related Software Services Are Creating Jobs and Economic Growth

In high-tech and low-tech industries alike, the analysis of data has made businesses more agile, responsive, and competitive, boosting the underlying productivity of many key pillars of our economy.

The economic implications of the data revolution—and AI and related software solutions that leverage that data—are enormous. Economists predict that making better use of data could lead to a “data dividend” of \$1.6 trillion in the next four years, and that data-enabled efficiency gains could add almost \$15 trillion to global GDP

by 2030.² In addition, experts predict that applications of AI technologies could grow the global economy by \$7.1 to \$13.17 trillion over the next eight years.³

AI systems are powered by software, which itself is a major engine of economic growth. In September, Software.org: the BSA Foundation released a study with data from the Economist Intelligence Unit (EIU) showing that the software industry alone contributed more than \$1.14 trillion to U.S. GDP in 2016—a \$70 billion increase in just two years.⁴ The study also showed that the software industry is a powerful job creator, supporting over 10.5 million jobs, with a significant impact on job and economic growth in each of the 50 states.⁵

B. AI and Related Software Services Are Improving Every Industry

The benefits of AI are not limited to the software sector. In fact, AI innovation is stimulating growth across all industry sectors as businesses, big and small, use AI and related software services to improve supply chains, secure their networks, and evaluate how to improve their products and services. There are numerous examples of this positive impact across a wide swath of industries, for instance:

- *Cybersecurity.* AI tools are revolutionizing how we monitor network security, helping analysts parse through hundreds of thousands of security incidents per day to weed out false positives and identify threats that warrant further attention by network administrators. By automating responses to routine incidents and enabling security professionals to focus on truly significant threats, AI-enabled cyber tools are helping enterprises stay ahead of their malicious adversaries.⁶
- *Financial Services.* AI is improving fraud detection by providing companies with real-time information that helps them identify and investigate different types of fraud, reducing the losses attributed to fraudsters by billions of dollars. In a matter of seconds, machine learning algorithms can generate a risk score for a transaction by parsing through large volumes of data about the vendor and the purchaser to determine the likelihood of fraud.⁷ These tools are protecting consumers from the risk of fraudulent charges and from the frustration associated with “false declines.”

²See BSA, *What's the Big Deal With Data?* 14 (Oct. 2015), available at http://data.bsa.org/wp-content/uploads/2015/12/bsadatastudy_en.pdf. The potential of digital data to improve the healthcare system is substantial: some estimates predict that if the healthcare sector were to use data more effectively to drive efficiency and quality, the sector could save more than \$300 billion every year. See James Manyika *et al.*, *Big Data: The Next Frontier for Innovation, Competition, and Productivity*, McKinsey Global Institute (May 2011), available at http://www.mckinsey.com/insights/business_technology/big_data/the_next_frontier_for_innovation.

³See *Disruptive technologies: Advances that will transform life, business, and the global economy*, McKinsey Global Institute (May 2013), available at <http://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/disruptive-technologies>.

⁴Software.org: The BSA Foundation, *The Growing \$1 Trillion Economic Impact of Software* 5 (Sept. 2017), available at https://software.org/wp-content/uploads/2017_Software_Economic_Impact_Report.pdf.

⁵*Id.*

⁶For example, IBM's Watson for Cyber Security is a cybersecurity tool that can analyze 15,000 security documents per day—a rate essentially impossible for any individual to achieve. Watson's data processing capabilities enable analysts to more quickly identify incidents that require human attention. See IBM, *IBM Delivers Watson for Cyber Security to Power Cognitive Security Operations Centers* (Feb. 13, 2017), <https://www-03.ibm.com/press/us/en/pressrelease/51577.wss>; Jason Corbin, *Bringing the Power of Watson and Cognitive Computing to the Security Operations Center*, Security Intelligence (Feb. 13, 2017), https://securityintelligence.com/bringing-the-power-of-watson-and-cognitive-into-the-security-operations-center/?cm_mc_uid=70595459933115020631816&cm_mc_sid=50200000=1503364089&cm_mc_sid=52640000=1503365578. Splunk uses a similar model, with machine-learning algorithms conducting real-time analysis and processing of massive volumes of data from all sensors on a network to identify anomalies, feeding visualization tools that help network administrators efficiently triage security incidents. See David Braue, *Machine learning key to building a proactive security response: Splunk*, CSO Online (Aug. 20, 2015), <https://www.cso.com.au/article/582483/machine-learning-key-building-proactive-security-response-splunk/>. Microsoft's Windows 10 Anniversary Edition introduced AI-driven capabilities for automatically isolating suspicious network traffic pending adjudication by network administrators. See Chris Hallum, *Defense Windows clients from modern threats and attacks with Windows 10*, Channel 9 video content (Oct. 6, 2016), available at <https://channel9.msdn.com/events/Ignite/2016/BRK2135-TS>; “Intelligent Security: Using Machine Learning to Help Detect Advanced Cyber Attacks,” <https://www.microsoft.com/en-us/security/intelligence>.

⁷See generally Pablo Hernandez, *CA Technologies Uses AI Tech to Combat Online Fraud*, eSecurityPlanet, May 4, 2017, available at <https://www.esecurityplanet.com/network-security/ca-technologies-uses-ai-tech-to-combat-online-fraud.html>.

- *Agriculture.* AI is helping farmers tackle some of the biggest issues they face, including declining crop yields and changing weather patterns, through precision farming, better data analysis, and improved operational efficiency. For instance, tools like computer vision and deep-learning algorithms are enabling farmers more effectively to process data for purposes of monitoring crop and soil health.⁸
- *Manufacturing.* AI-enabled tools are also helping factory owners streamline their manufacturing processes and resolve problems common to most factories, such as inaccurate demand forecasting and capacity planning, unexpected equipment failures and downtimes, and supply chain bottlenecks. Predictive maintenance, for instance, allows manufacturers to achieve 60 percent or more reduction in unscheduled system downtime. Cameras powered by computer vision algorithms can fix product defects immediately and identify root causes of failure. AI thus enables manufacturers to reduce waste, shorten production periods, increase yields on production inputs, and improve both revenue and workplace safety.⁹
- *Healthcare.* AI technologies are already providing solutions that help save lives. A 2016 Frost & Sullivan report predicts that AI has the potential to improve health outcomes by 30 to 40 percent.¹⁰ AI is helping fuel these improved health outcomes not by replacing the decision-making of healthcare professionals, but by giving these professionals new insights and new ways of analyzing and understanding the health data to which they have access. For example, AI tools are powering machine-assisted diagnosis and surgical applications are being used to improve treatment options and outcomes. Image recognition algorithms are helping pathologists more effectively interpret patient data, thereby helping physicians form a better picture of patients' prognosis.¹¹ The ability of AI to process and find patterns in vast amounts of data from disparate sources is also driving important progress in biomedical and epidemiological research.¹²
- *Education.* AI technologies offer tools for students, teachers, and administrators to help students learn more effectively both within and outside of the classroom. AI programs can, for example, analyze a student's performance in a particular skill across subjects over the course of a year and automatically provide new content or specified learning parameters, offering students continual, individualized practice and feedback. They can also help teachers better understand student performance, quickly identify students that need particular attention, and develop lesson plans that customize instruction, content, pace, and testing to individual students' strengths and interests.¹³ AI solutions also are helping administrators track attendance patterns and gain insights on student performance more broadly.¹⁴

⁸ See Kumba Sennaar, *AI in Agriculture—Present Applications and Impact*, techemergence (Nov. 17, 2017), <https://www.techemergence.com/ai-agriculture-present-applications-impact/>.

⁹ See Mariya Yao, *Factories Of The Future Need AI To Survive And Compete*, Forbes.com (Aug. 8, 2017), <https://www.forbes.com/sites/mariyayao/2017/08/08/industrial-ai-factories-of-future/#2d7ab2fd128e>.

¹⁰ See *From \$600 M to \$6 Billion, Artificial Intelligence Systems Poised for Dramatic Market Expansion in Healthcare*, Frost & Sullivan (Jan. 5, 2016), <https://ww2.frost.com/news/press-releases/600-m-6-billion-artificial-intelligence-systems-poised-dramatic-market-expansion-healthcare>.

¹¹ See e.g., Meg Tirrell, *From coding to cancer: How AI is changing medicine*, cnbc.com (May 11, 2017), <https://www.cnbc.com/2017/05/11/from-coding-to-cancer-how-ai-is-changing-medicine.html>.

¹² For instance, AI is helping biologists who are aiming to treat 100 molecular genetic diseases by 2025. See Splunk, *Machine Learning Helps Recursion Pharmaceuticals Treat Genetic Diseases* (Nov. 7, 2017), https://www.splunk.com/en_us/newsroom/press-releases/2017/splunk-machine-learning-helps-recursion-pharmaceuticals-treat-genetic-diseases.html. In another example, Microsoft researchers are also using AI and related technologies to better understand the behavior of cells and their interaction, which could ultimately help "debug" an individual's specific form of cancer and allow doctors to provide personalized cancer treatment. See generally, Microsoft, *Biological Computation*, <https://www.microsoft.com/en-us/research/group/biological-computation/>.

¹³ See Software.org: The BSA Foundation, *The Growing \$1 Trillion Economic Impact of Software*, *supra* note 4, at 7; see also Daniel Faggella, *Examples of Artificial Intelligence in Education*, TechEmergence (Mar. 7, 2017), <https://www.techemergence.com/examples-of-artificial-intelligence-in-education/>.

¹⁴ Benjamin Herold, *Are schools ready for the power and problems of big data?*, Education Week (Jan. 11, 2016), available at <http://www.edweek.org/ew/articles/2016/01/13/the-future-of-big-data-and-analytics.html>.

C. AI Services Provide Tremendous Societal Benefits

The range of potential societal benefits from the use of AI services is equally vast. For example, AI solutions are at the heart of new devices and applications that improve the lives of people with disabilities, including helping people with vision-related impairments interpret and understand photos and other visual content.¹⁵ This technology opens new possibilities for people with vision impairments to navigate their physical surroundings, giving them increased independence and greater ability to engage with their communities.

AI is also helping governments improve constituent services in ways that save time, money, and lives. For example, cities are optimizing medical emergency response processes using AI-based systems, enabling them to more strategically position personnel and reduce both response times and the overall number of emergency trips.¹⁶ AI is also helping to leverage data to improve disaster response and relief efforts, including after the 2015 earthquake in Nepal.¹⁷

* * * *

Whether it is detecting financial fraud, improving health outcomes, making American farmers more competitive, or enhancing government and emergency services, the impact of AI and related software services is already visible in every industry, in every state, and across the globe.

III. Fostering Consumer Trust in AI

Even as society gains from the substantial benefits that AI offers, we also recognize that there may be legitimate concerns about how AI systems are deployed in practice, which may also affect trust and confidence in AI. In particular, as people increasingly apply AI services in new contexts, questions may arise about how they operate, whether they treat people fairly and are free from improper bias, and their impact on jobs. Like many technologies, AI has an almost infinite range of beneficial uses, but we should also take appropriate steps to ensure that it is deployed responsibly. We recognize that responsible deployment of AI should instill consumer confidence that these important issues will be appropriately addressed.

A. Enhancing Understanding of AI Systems

Building trust and confidence in AI-enabled systems is an important priority. In some instances, the complexity of these technologies, which are designed to identify patterns and connections that humans could not easily identify on their own, can make it challenging to explain how certain aspects of AI systems work. BSA members understand that, in order to promote trust, companies that build and deploy AI systems will need to provide meaningful information to enhance understanding of how these systems operate.

Indeed, ensuring that AI systems operate as intended and treat people fairly is an important priority. We are eager to participate in meaningful dialogues with other stakeholders about how best to accomplish that goal, and we welcome opportunities such as this one to help advance that dialogue. Currently, relevant technical tools and operational processes that could improve understanding and confidence in AI systems are still being developed, and it is an area of robust research. Although more work needs to be done, it is already clear that expectations are highly context-specific—and demands will vary based on this context. As we seek to address these important issues, we will aim to ensure that we remain sufficiently flexible to respond to concerns, and to adapt to the changing landscape as these emerging technologies, and potential solutions to new challenges, continue to evolve.

B. Preparing the Workforce for the Jobs of the Future

As AI services improve every industry, they will likely have a multi-dimensional impact on employment. The deployment of AI in the workplace will enable employees to focus on tasks that are best suited to uniquely human skillsets, such as creativity, empathy, foresight, judgment, and other social skills. Although there ap-

¹⁵ For instance, Microsoft recently released an intelligent camera app that uses a smartphone's built-in camera functionality to describe to low-vision individuals the objects that are around them. See Microsoft, *Seeing AI*, <https://www.microsoft.com/en-us/seeing-ai/>.

¹⁶ See Kevin C. Desouza, Rashmi Krishnamurthy, and Gregory S. Dawson, *Learning from public sector experimentation with artificial intelligence*, Brookings Institution (June 23, 2017), <https://www.brookings.edu/blog/techtank/2017/06/23/learning-from-public-sector-experimentation-with-artificial-intelligence/>.

¹⁷ See Patrick Meier, *Virtual Aid to Nepal: Using Artificial Intelligence in Disaster Relief*, Foreign Affairs (June 1, 2015), available at <https://www.foreignaffairs.com/articles/nepal/2015-06-01/virtual-aid-nepal>.

pears to be no consensus on the precise impact AI will have on employment, there is broad recognition that widespread deployment of these technologies will create demand for new types of jobs, and that these jobs often will require skills that many workers today do not yet have.

Current estimates indicate the United States will not have enough workers to meet the predicted high demand for computer science-related jobs. For example, by 2020, the U.S. Bureau of Labor Statistics predicts that there will be 1.4 million computing jobs, but just 400,000 computer science students with the skills necessary to fill those jobs.¹⁸ It is imperative that the United States takes steps now to ensure that we have a sufficient pipeline of workers with the skills needed to perform these new, high-quality jobs.

Yet even these estimates do not take into account the extent to which the use of AI may require new skills. Because AI services will likely be integrated across all sectors of the economy, the new jobs AI creates, and the new skills that will be needed, will reach beyond the tech sector, and will also likely extend to workers in both urban and rural areas. Indeed, many of these jobs will “look nothing like those that exist today,” and will include “entire categories of new, uniquely human jobs” that will require “skills and training that have no precedents.”¹⁹ As a result, one key challenge that lies ahead is determining how to ensure that the U.S. workforce has the skills necessary for the future.

BSA members are working hard to help address this challenge. BSA recognizes that this will require a multi-faceted solution, including cooperation with public and private stakeholders. We seek to identify opportunities and partnerships that focus on retraining the workforce with new skills, creating a pipeline of workers with skills to fill the next generation of jobs, increasing access to those jobs for skilled workers, and increasing deployment of cloud services, which facilitate employment and collaboration in different geographic regions.

Notably, BSA members already have begun helping workers and youth acquire new skills that will enable them to leverage AI systems.²⁰ BSA members offer several high-tech and business training programs, including at the high school level. Some programs target populations not traditionally associated with tech jobs, such as military veterans.²¹ These initiatives illustrate just some of the ways in which AI-based employment concerns can be meaningfully addressed.

¹⁸ See Allie Bidwell, *Tech Companies Work to Combat Computer Science Education Gap*, U.S. NEWS & WORLD REPORT, Dec. 27, 2013, available at <https://www.usnews.com/news/articles/2013/12/27/tech-companies-work-to-combat-computer-science-education-gap>.

¹⁹ H. James Wilson, Paul R. Daugherty, Nicola Morini-Bianzino, *The Jobs that Artificial Intelligence will Create*, MIT Sloan Management Review (Mar. 23, 2017), available at <https://sloanreview.mit.edu/article/will-ai-create-as-many-jobs-as-it-eliminates/>.

²⁰ See, e.g., Allen Blue, *How LinkedIn is Helping Create Economic Opportunity in Colorado and Phoenix* (Mar. 17, 2016), <https://blog.linkedin.com/2016/03/17/how-linkedin-is-helping-create-economic-opportunity-in-colorado-and-phoenix>; Markle Foundation, *Why Microsoft and the Markle Foundation are Working Together to Connect Workers with New Opportunities in the Digital Economy*, <https://www.markle.org/microsoft>. IBM, for instance, has established Pathways in Technology Early College High Schools (P-TECH Schools). P-TECH schools are innovative public schools that offer students the opportunity to earn a no-cost associates degree within six years in fields such as applied science and engineering—and to acquire the skills and knowledge necessary to pursue further educational opportunities or to step easily into well paying, high-potential informational technology jobs. IBM designed the P-TECH model to be both widely replicable and sustainable as part of an effort to reform career and technical education. See IBM, *IBM and P-TECH*, <https://www-03.ibm.com/press/us/en/presskit/42300.wss>. Likewise, Salesforce offers free high-tech and business skills training through Trailhead, its online learning platform, with the goal of preparing them for the estimated 3.3 million jobs created by the Salesforce economy worldwide from 2016 to 2022, nearly 1 million of which are forecasted to be in the United States. See International Data Corporation, *The Salesforce Economy Forecast: 3.3 Million New Jobs and \$859 Billion New Business Revenue to Be Created from 2016 to 2022* (Oct. 2017), available at <http://www.salesforce.com/assets/pdf/misc/idc-study-salesforce-economy.pdf>; see also Gavin Mee, *How the Salesforce Economy is Driving Growth and Creating Jobs*, Oct. 24, 2017, available at <https://www.salesforce.com/uk/blog/2017/10/idc-how-the-salesforce-economy-is-driving-growth-and-creating-jobs>; Gavin Mee, *Guest Blog: Gavin Mee, Salesforce—Evolving tech means change in digital skills*, TechUK (Apr. 26, 2017), at <https://www.techuk.org/insights/opinions/item/10695-guest-blog-gavin-mee-salesforce-evolving-tech-means-change-in-digital-skills>.

²¹ For example, the Splunk4Good initiative, which partners with non-profits, is helping military veterans and their families, along with youth, train for careers in technology, providing free access to Splunk licenses and its extensive education resources to help them attain marketable skillsets. See Splunk, *Splunk Trains Workforce of Tomorrow With Amazon Web Services, NPower, Wounded Warrior Project and Year Up*, (Sept. 26, 2017) https://www.splunk.com/en_us/newsroom/press-releases/2017/splunk-trains-workforce-of-tomorrow-with-amazon-web-services-npower-wounded-warrior-project-and-year-up.html.

IV. Opportunities for Congress and the Administration to Facilitate AI Innovation

As innovation in AI and related software services increasingly fuels growth in the global economy, countries around the world are taking steps to invest in education, research, and technological development to become a hub for AI innovation. For example, the UK government recently released an Industrial Strategy, which identifies putting the UK at the forefront of the AI and data revolution as one of four key strategies that will secure its economic future.²² In the EU, the European Parliament recently issued a report on civil law rules regarding robotics, which highlights the opportunities robotics and AI offer and encourages investment in such technology so Europe can maintain leadership in this space.²³ Likewise, in Japan, the government recently issued a new strategy designed to strengthen collaboration between industry, the government, and academia on matters related to robotics, and also issued a report offering the first systematic review of AI networking issues in Japan.²⁴ In China, the government has issued a “Next Generation Artificial Intelligence Development Plan,” which lays out objectives for AI development in China for the next 13 years and calls on China to become a global AI innovation center by 2030.²⁵

In the United States, a flexible policy framework that facilitates responsible AI deployment and increased investment will be key to preserving U.S. global economic competitiveness. An essential part of that effort will be ensuring the ability to access data, and to transfer that data seamlessly across borders, which are vital for AI to flourish. It also will be important to support investment in AI-related education, workforce development, and research. To that end, there are several steps that Congress and the Administration could take to spur AI innovation and continued economic growth.

A. Pass OPEN Government Data Act

First, Congress should pass the OPEN Government Data Act. This legislation, which the House recently passed as Title II of the Foundations for Evidence-Based Policymaking Act, recognizes that government-generated data is a national resource that can serve as a powerful engine for creating new jobs and a catalyst for economic growth. To that end, the OPEN Government Data Act would require agencies to make non-sensitive government data more open, available, and usable for the general public. Making such data more readily available will improve government transparency, promote government efficiency, and foster innovation of data-driven technologies such as artificial intelligence.

We would like to thank Ranking Member Schatz for his tireless work as an original sponsor of the OPEN Government Data Act. We are hopeful that the Senate will act soon to secure its final passage into law.

B. Support Efforts to Promote Digital Trade and Facilitate Data Flows

We also urge Congress and the Administration to continue supporting efforts to expand digital trade. Indeed, the new digital data economy, which increasingly relies on AI and related software services, will benefit from a globally recognized system for digital trade that facilitates cross-border data flows and establishes clear rules, rights, and protections. There are several opportunities for Congress and the Administration to lead in this area.

First, the ongoing NAFTA discussions provide an important opportunity to modernize the trade agreement, which was initially negotiated when digital services were in their infancy. We are encouraged that the Administration has made it an objective to seek to prohibit market access barriers to digital trade, including restrictions on data transfers, data localization mandates, and technology transfer requirements.

²² See UK Secretary of State for Business, Energy and Industrial Strategy, *Industrial Strategy: Building a Britain fit for the future* (Nov. 2017), available at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/662541/industrial-strategy-white-paper-print-version.pdf.

²³ See European Parliament 2014–2019, *Resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics*, Eur. Parl. Doc. P8 TA (2017)0051, <http://www.europarl.europa.eu/sides/getDoc.do?pubRef=/EP//NONSGML+TA+P8-TA-2017-0051+DOC+PDF+V0//EN>.

²⁴ See Fumio Shimpo, *Japan’s Role in Establishing Standards for Artificial Intelligence Development*, Carnegie Endowment for International Peace (Jan. 12, 2017), <http://carnegieendowment.org/2017/01/12/japan-s-role-in-establishing-standards-for-artificial-intelligence-development-pub-68311>.

²⁵ See Elsa Kania, *China’s Artificial Intelligence Revolution*, The Diplomat (Jul. 27, 2017), available at <https://thediplomat.com/2017/07/chinas-artificial-intelligence-revolution/>.

Second, another key priority is ensuring that transatlantic trade continues to thrive. In particular, we appreciate Congress's and the Administration's leadership on issues relating to the EU-U.S. Privacy Shield, which both protects privacy and facilitates data transfers between the EU and United States. We encourage your continued support as the Administration proceeds with its ongoing successful implementation of the framework.

Third, as other countries seek to modernize their trade policies, the Administration should engage key global partners to ensure that new trade initiatives facilitate data-driven innovation and protect against market access barriers for e-commerce and digital trade.

C. Invest in AI research, education, and workforce development

Unlocking the full promise of AI technologies also requires a long-term strategy of investing in education, workforce development, and research. Because human beings ultimately drive the success of AI, supporting education, training, and research is essential to extracting the maximum level of benefit that AI technologies offer.

As an initial matter, Congress and the Administration should ensure that education programs are developing human talent more effectively. Broadly speaking, this means that Congress and the Administration should support science, technology, engineering, and mathematics (STEM) education at all levels. It also means creating and supporting programs that help educate researchers and engineers with expertise in AI, as well as specialists who apply AI methods for specific applications and users who operate those applications in specific settings.²⁶ For researchers and engineers, these programs should include training in computer science, statistics, mathematical logic, and information theory, and for specialists, they should focus on software engineering and related applications.²⁷

Congress and the Administration should also support the development of new and innovative ways to ensure the U.S. workforce is prepared for the jobs of the future. Because AI will generate new jobs in categories both known and unforeseen, we need to develop thoughtful and effective approaches to equip the U.S. workforce with the skills necessary to seize the opportunities these new technologies create and to optimize the role of AI in modern life.

Continued scientific research is essential to fully tapping the potential of AI technology. Congress and the Administration should therefore also promote both public and private sector research to help ensure that the United States remains a leader in this space. The U.S. Government should invest in the types of "long-term, high-risk research initiatives" in which the commercial sector may be reluctant to invest. In the past, such R&D investments have led to "revolutionary technological advances... [such as] the Internet, GPS, smartphone speech recognition, heart monitors, solar panels, advanced batteries, cancer therapies, and much, much more."²⁸ Congress and the Administration should also adopt policies that incentivize private-sector R&D, including by expanding access to financing.

Passing the OPEN Government Data Act, supporting efforts to promote digital trade and facilitate cross-border data flows, and investing in AI research, education, and workforce development will be critical to maximizing the opportunities AI presents and helping to ensure that the United States maintains leadership in AI innovation and deployment, even as other nations increase their own efforts to take advantage of the possibilities that AI offers.

* * *

We appreciate Congress's leadership on the important issue of facilitating AI innovation and its responsible deployment. Thank you and I look forward to your questions.

Senator WICKER. Thank you very much.
Dr. Gil.

²⁶See U.S. Executive Office of the President, *Preparing for the Future of Artificial Intelligence*, National Science and Technology Council Committee on Technology 26 (Oct. 2016), available at https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.

²⁷See *id.*

²⁸See The National Artificial Intelligence Research and Development Strategic Plan (Oct. 2016), available at https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf.

**STATEMENT OF DR. DARIO GIL, PH.D., VICE PRESIDENT,
AI AND IBM Q**

Dr. GIL Chairman Wicker, Ranking Member Schatz, members of the Subcommittee, thank you for inviting me here today. My name is Dario Gil, and I am the Vice President of AI and quantum computing at IBM.

The idea of creating a thinking machine is not new, and precedes modern computing. Calculating machines were built in antiquity and improved throughout history by many mathematicians. The term “artificial intelligence” was first introduced 61 years ago in 1956, and AI, as an academic discipline, took off. Three years later, IBM scientist Arthur Samuel coined the term “machine learning” to refer to computer algorithms that learn from and make predictions on data by building a model from sample inputs without following a set of static instructions.

One type of machine learning and AI algorithm that has gained tremendous attention over the past several years is an artificial neural network, notably, deep learning. These networks are inspired by the architecture of the human brain, with neurons organized as layers, and different layers may perform different kinds of operations on their inputs. When presented with sample data, neural net can be trained to perform a specific task, such as recognizing speech or images. Over the last decade, the explosion of digital data and the growth in processing speed and power have made it possible to use neural nets in real-world solutions.

While many tend to focus on the automation features of AI, we believe its true impact will be felt in humans carrying out complex tasks which they cannot do on their own. My prepared testimony provides detailed examples on the many ways in which IBM’s AI platform for enterprise business, Watson, is being used to augment human abilities across many industries, from strengthening cybersecurity to enhancing the customer experience to improving agriculture and optimizing supply chains. AI is playing a bigger and bigger role in all realms of commerce, and its uses will only grow.

Now, there is no question the advent of AI will impact the nature of jobs, yet history suggests that even in the face of technological transformation, employment continues to growth with economic expansion and the creation of entirely new jobs despite the disappearance of some occupations.

Jobs are made out of tasks. Those tasks that cannot be automated, that cannot be automated by AI, are those in which workers will provide the greatest value, commanding higher wages and incomes as a result.

Now, the creation of AI itself will require new job categories associated with how we design and train them, how we secure them, and verify that they work as planned, and how we integrate them into our workflows. The application of AI will change our professions, opening up new categories of work and increasing demand for some existing professions as we combine the capabilities of these AI systems with our own expertise. For example, many more cybersecurity professionals will be needed to engage with AI systems and act decisively upon the threat to information they provide.

We must address the problem of shortage of workers with the skills needed for these and many other roles. A useful example in this regard is software programming, which is taught as a critical skill in many high schools and colleges. We should promote a similar movement for AI techniques, such as machine learning.

To enjoy the full benefits of AI, we'll also need to have confidence in the recommendations, judgments, and uses of AI systems. In some cases, users will need to justify why an AI system produced its recommendations. For example, doctors and clinicians using AI systems to support medical decisionmaking may require to provide specific explanations for a diagnosis or course of treatment, both for regulatory and liability reasons.

IBM is actively innovating in this field. We're deriving best practices for how, when, and where AI algorithms should be used. We're creating new AI algorithms that can be trusted, are explainable, and are more accurate.

Of course, no single company can guarantee the safe and responsible use of such a pervasive technology. For this reason, IBM is a founding member of the Partnership on AI, a collaboration with other key industry leaders and many scientific and nonprofit organizations. It's focused on the formulation of best practices, on AI technologies, and advancing the public's understanding of AI.

In addition, a recently created MIT-IBM Watson AI Lab has one of its research pillars advancing shared prosperity with AI, exploring how AI can deliver economic and societal benefits to a broader range of people, nations, and enterprises.

In a similar way, we look forward to working closely with the Members of Congress to ensure the responsible, ethical, and equitable use of AI as this technology continues to evolve.

[The prepared statement of Dr. Gil follows:]

PREPARED STATEMENT OF DARIO GIL, VICE PRESIDENT, AI AND IBM Q

Introduction

Chairman Wicker, Ranking Member Schatz, members of the Subcommittee. Thank you for inviting me here today. My name is Dario Gil and I am Vice President, AI and quantum computing at IBM.

We have arrived at a remarkable moment in the history of information technology. An explosion of data and computation has given us access to massive amounts of digitized knowledge. With it, we have enormous intelligence and power to see patterns and solve problems we never could have before.

Increasingly, the engine we use to tap into this knowledge is artificial intelligence. We can now train algorithms with emerging Artificial Intelligence (AI) technologies to learn directly from data by example. Moreover, we can do this at scale cost through cloud networks to create machines that help humans think. For these reasons, AI is the most important technology in the world today.

The rise of AI has been accompanied by both boundless enthusiasm about its ability to transform our lives, and fears it could potentially harm or displace us. At IBM, we take a different approach. We are guided by the use of artificial intelligence to augment human intelligence. We focus on building practical AI applications that assist people with well-defined tasks. We believe people working collaboratively with these learning systems is the future of expertise.

In my testimony, I'll provide an overview of AI and describe how work in this field has evolved. Then I'll offer some examples of the rapidly growing commercial applications of IBM Watson, the best-known artificial intelligence platform for enterprise business today. I'll also look at how we're beginning to combine AI with other emerging technologies, such as blockchain, to optimize business and provide trust in transactions. I'll describe how AI will impact the nature of work leading to many new and improved job opportunities. Finally, I'll examine IBM's position on the responsible and ethical use of AI.

The Evolution of AI

The idea of creating a ‘thinking’ machine is not new and precedes modern computing. The study of formal reasoning dates to ancient philosophers such as Aristotle and Euclid. Calculating machines were built in antiquity and were improved throughout history by many mathematicians. In the 17th century Leibniz, Hobbes and Descartes explored the possibility that all rational thought could be made as systematic as algebra or geometry.

In 1950, Alan Turing, in his seminal paper *Computing Machinery and Intelligence*, laid out several criteria to assess whether a machine could be deemed intelligent. They have since become known as the “Turing test.” The term “artificial intelligence” was first introduced in 1956, sixty-one years ago, and AI as an academic discipline took off. Three years later, in 1959, IBM scientist Arthur Samuel coined the term “machine learning” to refer to computer algorithms that learn from and make predictions on data by building a model from sample inputs, without following a set of static instructions.

An algorithm is simply a set of rules to be followed in calculations or other problem-solving operations. It can be as basic as the steps involved in solving an addition problem or as complex as instructing a computer how to perform a specific task. One type of machine learning and AI algorithm that has gained tremendous attention over the past several years is an artificial neural network. It has been essential to the explosive growth of AI systems today.

Artificial neural networks are inspired by the architecture of the human brain. They contain many interconnected processing units, called artificial neurons, which are analogous to biological neurons in the brain. Typically, neurons are organized in layers. Different layers may perform different kinds of operations on their inputs. When presented with sample data, an artificial neural network can be trained to perform a specific task, such as recognize speech or images. For example, an algorithm can learn to identify images that contain cars by analyzing numerous images that have been manually labeled as “car” or “no car.” It can then use those results to identify cars in images that it has not seen before.

Even though neural networks and other machine learning algorithms were actively researched more than six decades ago, their practical use was hindered by the lack of digitized data from which to learn from and insufficient computational power. At the time, most data were in analog form and not easily available to a computer. Training of the neural network algorithm was and remains a computationally intensive process. Due to the limitations of processors, it could not be implemented effectively.

Over the last decade, the explosion of digital data, the growth in processing speed and power, and the availability of specialized processing devices such as graphical processing units (GPUs) have made it possible to use artificial neural networks in real-world solutions. Today, computation is carried out not only in the cloud and in data centers, but also at the edge of the network, in sensors, wearable devices, smart phones, embedded electronics, factory machines, home devices, or components in a vehicle.

These conditions have also allowed researchers and engineers to create incredibly complex neural networks, called *deep learning* networks. They perform in ways comparable to humans in many tasks. For certain tasks, such as speech and image recognition, game playing, and medical image classification, these networks can outperform people. Today, neural networks are used in a variety of applications, including computer vision, speech recognition, machine translation, social network analysis, playing board and video games, home assistants, conversational devices and chatbots, medical diagnostics, self-driving cars, and operating robots.

In addition to machine learning, AI systems deploy a variety of other algorithms and technologies that include knowledge representation, machine reasoning, planning and scheduling, machine perception (speech and vision), and natural language processing and understanding. At IBM, we are actively researching and advancing these and other technologies so that we can continue to enhance AI systems.

We are also envisioning and developing the next-generation infrastructure required for increasingly complex AI tasks and workloads. This is the physical hardware required to run AI algorithms: the processors, servers, databases, storage, data centers and cloud infrastructure. When all these pieces are aligned in a way that allows algorithms to analyze data with maximum efficiency, we refer to it as the “full stack.”

By successfully engineering the full stack, we can build AI-powered solutions that we can apply to a broad array of societal and industry challenges. While many tend to focus on the benefit of automation, we believe that AI’s true impact will be felt in assisting people’s daily lives, and by helping us carry out extremely complex tasks we cannot do on our own. That includes everything from forecasting the

weather, to predicting how traffic will flow, to understanding where crops will grow the best. AI will also help us research the best combinations of compounds for drug development, repurpose chemical structures for new medicines, and optimize vastly intricate supply chains. I'd like to illustrate this further with a look at how IBM Watson is already being used across a range of different industries.

AI applications to industries

My first example illustrates how AI can assist humans in reacting to a problem when there is very little time to react. The security of data and on-line transactions is fundamental to the growth of commerce. But the simple fact is that most organizations can't keep up with the threats. Security experts are too few and over-stretched. Sophisticated attacks, including those using AI tools, are coming at a rate that makes them extremely difficult to stop. Entire networks are compromised in the blink of an eye. Watson for cybersecurity allows us to turn the tables. It sifts through the insights contained within vast amounts of unstructured data, whether it's documented software vulnerabilities or the more than 70,000 security research papers and blogs published each year. It instantly alerts security experts to relevant information, scaling and magnifying human cognition. It also learns from each interaction that has an alert, and works proactively to stop continued intrusion. Security analysts, armed with this collective knowledge, can respond to threats with greater confidence and speed.

The second example shows how AI is enhancing customer experience. Tax preparation is an area ripe for AI solutions. H&R Block is using Watson to understand context, interpret intent and draw connections between clients' statements and relevant areas of their tax return. Watson is working alongside H&R Block Tax Pros as they take clients through the tax return process, suggesting key areas where they may qualify for deductions and credits. Clients can follow along and understand how their taxes are being computed and impacted by numerous aspects of the tax code. They can also see the many paths to file a return with the IRS, pinpointing the route that delivers a maximum refund.

A third example demonstrates AI's ability to personalize the client experience. 1-800-Flowers launched an AI-powered gift concierge powered by Watson Conversation. It interacts with online customers using natural language. The service can interpret questions, then ask qualifying questions about the occasion, sentiment and who the gift is for to ensure that suggestions are appropriate and tailored to each customer. In this way, the customer can get the right flower for the right occasion.

The next example highlights AI's role in enhancing agricultural productivity. A program led by our Research division called IBM PAIRS is bringing the power of AI to improve crop yields. It works by processing and analyzing vast amounts of geospatial data to generate vastly improved long term weather and irrigation forecasts. Using these methods, IBM Research and Gallo Winery co-developed a precision irrigation method and prototype system that allows Gallo to use 20 percent less water for each pound of grapes it produces.

A final example shows AI's ability to optimize supply chains. Traditional brick and mortar retailers are under tremendous pressure from e-commerce. They must find new, cost-effective and efficient ways to deliver goods to buyers in order to stay in business. That means offering customers a range of delivery options—pick up in store, ship from the nearest store, or move goods seamlessly between store and e-commerce. Our client, a major American retailer, had to coordinate this effort across a thousand stores in their fulfillment chain. Our predictive models enabled them to determine optimal distribution across their entire chain, factoring in dozens of different variables. Over an eight-day period including Black Friday and Cyber Monday, they processed more than 4 million orders—a company record—at a savings of 19 percent per order compared to the prior year. This led to an overall savings of \$7.5 million dollars.

From cybersecurity, to customer experience, to personalization, to productivity and optimization, AI is playing a bigger and bigger role in all realms of commerce. And its uses will only grow.

AI and blockchain

The potential for AI becomes even greater when combined with other emerging technologies such as blockchain. Blockchain stores data in shared, secure, distributed ledgers that allow every participant appropriate access to the entire history of a transaction using a “permissioned” network—one that is highly secure and can determine who can see what. Blockchain holds the promise to be the way we may do all transactions in the future.

A typical AI process could include data collection, use of the data to train a neural network and then deployment of the pre-trained model to power the application.

Blockchain supports the AI process by reducing the risk of data tampering and provides data in a form that can be used and audited. There's an old saying in the computer industry "garbage in, garbage out," and that applies to data and how you use it. The integrity of the data used as input to the AI model is a necessary criterion in ensuring the value and usefulness of the model.

Because it can process and analyze massive quantities of data, AI can use blockchain data to gain valuable insights and detect patterns in how supply chains work and processes behave. Over time, this will generate a valuable source of clean, trusted transactional data that cuts across industries to give us new insights. That includes both structured and unstructured data—everything from Internet of Things (IoT) information to compliance and geospatial data that's stored on a blockchain. AI can use this information to generate valuable insights and detect patterns in near-real time, driving new efficiencies across business operations.

For example, IBM Research is working with Everledger, a company that tracks and protects diamonds and other valuables. We're using AI to analyze digital information on one million diamonds Everledger has stored on a blockchain. We can cross-check that data against UN regulations to prevent the sale of conflict diamonds. We can verify time and date stamps. We can certify laser inscriptions in the girdle of the stone. We can perform these analytics directly on the blockchain, without the need to extract the data first. This minimizes opportunities for data tampering and fraud. While this is a specialized application, it shows some of the kinds of data we can collect and analyze at huge scale.

We have also partnered with Walmart to use blockchain and AI techniques to ensure food safety. Today's food supply chains are highly complex and involve multiple components, stakeholders, and activities. This complexity makes it difficult to identify sources of contamination, counterfeit substitutions, loss of refrigeration, or food transportation safety issues as products move from their sources to their consumption by consumers. Blockchain supports traceability by tracking the food products from origin to destination and by allowing certification of respective transactions and events along the way. AI-powered technologies are used to analyze this information to help ensure food that can be eaten safely.

AI and the Future of Work

Artificial intelligence will alter the way people work. This has been true of many new technologies that have benefited human populations over time because they dramatically improved industrial output. They have led to fewer grueling jobs. In the process, new types of jobs have emerged. However, such disruptive improvements have always called for a period of training and adjustment.

We need to openly understand and recognize this fact, so that we can create the right conditions to make this transition as successful as possible. As a nation, we need to be prepared to offer the appropriate education and support to manage this change well. There's no question the advent of artificial intelligence will impact jobs. Despite the fear, anxiety, and prediction of massive job loss, history suggests that, even in the face of technological transformation, employment continues to grow and very few occupations disappear.

Rather, it is the transformation of occupations that is very likely to be widespread that will impact most workers. Occupations are made up of tasks. It is the tasks that are automated and reorganized where the transformation occurs. Workers will need new skills for the new transformed tasks and occupations. But, it is the tasks that cannot or will not be automated where workers provide the greatest value, commanding higher wages and incomes as a result.

Some "new collar jobs" will emerge—jobs that require advanced technical skills but do not necessarily require a full undergraduate education. A study by Accenture of more than 1,000 large companies that are already using or testing AI and machine-learning systems identified the emergence of entire categories of new, uniquely human jobs with no precedents.

For example, "trainers" will be required to teach AI systems how they should perform. They may write scripts for chatbots, helping them to improve their understanding of human communication, or help provide labeled data needed to train the algorithm. They may teach AI systems how to personalize their recommendations, or show compassion when responding to certain problems that require it. Another category of "explainers" will be needed to help convey how AI systems have arrived at a decision or a recommendation. They'll monitor the operations of AI systems or perform forensic analyses on algorithms and make corrections to them if they generate an incorrect result. Earlier, I referenced the shortage of qualified cybersecurity professionals. In the future, we'll need far more of them to engage with AI systems, review the recommendations these systems offer and act decisively upon threats.

There are actions we must take now to ensure the workforce is prepared to embrace the era of AI and the ways it will augment our economy. To begin, we must address the shortage of workers with the skills needed to make advances in AI, create new solutions and work in partnership with AI systems. We need to match skills education and training with the actual skills that will be required in the emerging age of AI.

At IBM, we have an educational model called P-TECH to train new collar workers for a job in technology. P-TECH combines the best of high-school, community college, hands-on-skills training and professional mentoring, and provides public high school students in grades 9–14 a path to post-graduation opportunities in fields aligned with the skills American employers are looking for.

Our goal must be to create multiple pathways like this for more people to acquire the skills that will be in demand, as AI use becomes more commonplace. We can use the example of the adoption of software programming as a critical skill that is taught in many high school and colleges. Some colleges require that all students learn how to code since they consider it a necessary skill for success. Students becoming proficient in programming have a wider range of job opportunities.

In the future, we may promote and see a similar trend with students gaining understanding of and proficiency in AI techniques such as machine learning. Preparing more U.S. students and workers for success in these well-paying new collar jobs is essential if we want a workforce that is ready to capitalize fully on AI's economic promise.

Let me also say that as well-intentioned as it may seem to some, taxing automation will not serve the cause of fostering employment in the new AI economy. It will only penalize technological progress. We should not adopt measures like this one that will harm America's competitiveness.

Inevitably, people adapt best by finding higher value in new skills. Technologies that are easiest to integrate and integrate with will be those that improve human productivity. But they should not replace human judgment. IBM Watson was designed from the beginning to work in concert with human expertise. It will only be successful as long as there are people with the right skills to engage with it.

Building trust in AI

To enjoy the full benefits of AI, we will also need to have confidence in the recommendations, judgments and uses of AI systems. IBM is deeply committed to the responsible and ethical development of AI. Last year, we published one of the first corporate white papers on this subject. The paper, which was intended to help launch a global conversation, centered around the need for safe, ethical, and socially beneficial management of AI systems.

Trust in automated systems is not a new concept. We drive cars trusting the brakes will work when the pedal is pressed. We perform laser eye surgery trusting the system to make the right decisions. We have automated systems fly our airplanes trusting they will navigate the air correctly. In these cases, trust comes from confidence that the system will not make a mistake, leveraging system training, exhaustive testing, and experience. We will require similar levels of trust for AI systems, applying these methodologies.

In some cases, users of AI systems will need to justify why an AI system produced its recommendation. For example, doctors and clinicians using AI systems to support medical decision-making may be required to provide specific explanations for a diagnosis or course of treatment, both for regulatory and liability reasons. Thus, in these cases, the system will need to provide the reasoning and motivations behind the recommendation, in line with existing regulatory requirements specific to that industry. In the European Union, this will be a requirement for all automated decision-making AI systems as of May 2018.

These safeguards can also help to manage the potential for bias in the decision-making process, another important concern with AI. Bias can be introduced both in the datasets that are used to train an AI system and by the algorithms that process that data, and how people interpret and communicate the discerned insights. Our belief is that the data and algorithmic aspects can not only be managed, but also that AI systems themselves can help eliminate many of the biases that already exist in human decision-making models today.

At the beginning of this year, IBM issued principles for transparency and trust to guide our development and use of AI systems. In summary, they state the following:

- We believe AI's purpose is to augment human intelligence
- We will be transparent about when and where AI is being applied, and about the data and training that went into its recommendations.

- We believe our clients' data and insights are theirs.
- We are committed to helping students, workers, and citizens acquire the skills to engage safely, securely, and effectively with cognitive systems, and to do the new kinds of work that will emerge in an AI economy.

In the same way that we are committed to the responsible use of AI systems, we are committed to the responsible stewardship of the data they collect. We also believe that government data policies should be fair and equitable and prioritize openness.

IBM is actively innovating in this field. We are deriving best practices for how, when, and where AI algorithms should be used. We are creating new AI algorithms that are more explainable and more accurate. We are working on the algorithmic underpinnings of bias and AI, such as creating technologies that can identify and cleanse illegal biases from training datasets.

Of course, no single company can guarantee the safe and responsible use of such a pervasive technology. For this reason, IBM is a founding member of the Partnership on AI, a collaboration with other key industry leaders and many scientific and nonprofit organizations. Its goal is to share best practices on AI technologies, advance the public's understanding, and serve as an open platform for discussion and engagement about AI and its influences on people and society.

AI has enormous transformative power. Much has been said about its potential to transform sectors and industries. However, AI is also giving us a technological toolkit to address many societal challenges. At IBM we are committed to pioneering new solutions, and showcasing and promoting the opportunities to use AI in social good applications. Three years ago, we launched the AI for Social Good program and have executed a number of AI for Good projects, from using AI to understand patterns of opioid addiction, to prototyping recommendation systems that would aid low-income individuals and help them stay out of poverty, to applying machine learning to understand transmission mechanisms of the Zika virus.

Earlier this year, we announced the MIT-IBM Watson AI Lab, a partnership with Massachusetts Institute of Technology (MIT) to carry out fundamental AI research. One of the research areas for the lab is focused on *advancing shared prosperity through AI*—exploring how AI can deliver economic and societal benefits to a broader range of people, nations and enterprises.

Lastly, no discussion of the future of AI would be complete without acknowledging the critical role of government. Public investment and policy support have been the twin pillars of American global technological leadership for the past half-century. We hope and expect the same will be true in the coming age of AI. For this reason, we enthusiastically welcome the interest and support of the United States Senate as this technology continues to evolve. Together, we can ensure that AI serves people at every level of society and advances the common good.

Senator WICKER. Thank you very much.
Dr. Felten.

**STATEMENT OF DR. EDWARD W. FELTEN, PH.D.,
ROBERT E. KAHN PROFESSOR OF COMPUTER SCIENCE
AND PUBLIC AFFAIRS, PRINCETON UNIVERSITY**

Dr. FELTEN. Chairman Wicker, Ranking Member Schatz, and members of the Subcommittee, thank you for the opportunity to testify today.

Progress in AI has accelerated over the last decade. Machines have met and surpassed human performance on many cognitive tasks, and some longstanding grand challenge problems in AI have been conquered.

Recent experience during this time teach us some useful lessons for thinking about AI as a developing technology.

First, AI is not a single thing; it's different solutions for different tasks. Success has come in "narrow AI," which applies a toolbox of specific technical approaches to craft solutions for specific applications. There has been a lot less progress on general AI, which tries to create a single, all-purpose artificial brain, like we see in the movies.

Second, successful AI doesn't think like a human. If it is an intelligence, it is sort of an alien intelligence. AI and people have different strengths and weaknesses, so teaming up with AI is promising if we can figure out how to work with an intelligence different from our own.

And, third, more engineering effort or more data translates into better AI performance. Progress requires a lot of hard work by experts, and that's why our AI workforce is so important.

The strategic importance of AI to the United States goes beyond its economic impact to include cybersecurity, intelligence analysis, and military affairs as well.

The U.S. is currently the world leader in AI research and applications, but our lead is not insurmountable. Countries around the world are investing heavily in AI, so our industry researchers and workforce need support in their efforts to maintain American leadership in this area. American companies recognized the potential of AI early on and have been investing and moving aggressively to hire top talent.

Our lead in research and development is less secure. Federal funding for AI research has been relatively flat. Aggressive hiring by industry has thinned the ranks of the academics who train the next generation of researchers. Industry does a lot of valuable research, but the public research community also plays an important role in basic research and in training young researchers, so investments in policies to support and grow the public research community are important.

Policies to enhance access to high-quality education for all American children, especially in computing, lay the foundation for our future workforce. And America has always been a magnet for talent from around the world, and that has to continue if we are to retain our leadership.

The many benefits of AI are tempered by some challenges. AI systems may pose safety risks, they may introduce inadvertent bias into decisions, and they may have unforeseen consequences. Much of the criticism of AI has centered on the risk of inadvertent bias, and real-world examples of biased AI are well documented.

The good news is that there are technical ways to eliminate bias. Developers can improve datasets to be more representative of the population, they can use algorithms that are more resistant to bias. Promising results on debiasing both data and algorithms are emerging from the research community, and that research should continue to be supported because it points a way to deploying AI more widely with less concern about bias.

In considering the risks of AI, it's important to remember that the alternative to relying on AI is to rely on people, and people are also at risk of error and bias. In the long run, AI systems will devise complex data-driven strategies to pursue goals, but people will continue to decide which goals the system should pursue. To better hold AI systems accountable, we need new technologies and new practices to connect AI with the human institutions that will govern it.

Regarding regulation, there is no need to create special regulations just for AI at this time. In sectors that are already regulated, the existing regulations are already protecting the public, and reg-

ulators need only consider whether and how to adjust the existing regulations to account for changes in practices due to AI. For example, the Department of Transportation, in the previous administration and this one, has been adapting vehicle safety regulation to enable safe deployment of self-driving cars.

Government agencies have important roles to play beyond regulation. More expertise, advice, and coordination is needed across the government to help agencies decide how to adapt regulations and use AI in their operations. New structures and new policies to strengthen this expertise would be very beneficial.

With good policy choices and the continued hard work and investment of American companies, researchers, and workers, AI can improve the health and welfare of Americans, boost productivity and economic growth, and make us more secure.

Americans currently lead the world in AI. We should not step on the brakes; instead, we should reach for the accelerator and the steering wheel.

Thank you for the opportunity to testify today.
 [The prepared statement of Dr. Felten follows:]

PREPARED STATEMENT OF EDWARD W. FELTEN ROBERT E. KAHN PROFESSOR OF COMPUTER SCIENCE AND PUBLIC AFFAIRS, PRINCETON UNIVERSITY

Chairman Wicker, Ranking Member Schatz, and members of the Committee, thank you for inviting me to speak today about how best to realize the benefits of artificial intelligence.

Artificial Intelligence (AI) and Machine Learning (ML)

Artificial intelligence (AI) and machine learning (ML) have been studied since at least 1950. There has been an unexpected acceleration in technical progress over the last decade, due to three mutually reinforcing factors: the availability of *big data sets*, which are analyzed by *more powerful algorithms*, enabled by *faster computers*. In recent years, machines have met and surpassed human performance on many cognitive tasks, and some longstanding grand challenge problems in AI have been conquered.

Industry has recognized the rise of AI as a technical shift as important as the arrival of the Internet or mobile computing. Companies around the world have invested heavily in AI research and development, and leaders of major companies have described adoption of machine learning as a bet-the-company opportunity.

The strategic importance of AI/ML to the United States goes beyond its economic impact. These technologies will also profoundly affect the future of security issues such as cybersecurity, intelligence analysis, and military affairs.

Fortunately, the United States is currently the world leader in AI/ML research, development, and applications, in both the corporate and academic spheres. Our national lead is not insurmountable, however. Countries around the world are investing heavily in AI/ML, so our scientists, engineers, and companies need support in their efforts to maintain American leadership.

The Nature of AI/ML Today

The history of AI teaches some important lessons that are useful in considering policy choices.

AI is not a single thing—it is different solutions for different tasks. The greatest progress has been in “narrow AI,” which applies a toolbox of specific technical approaches to craft a solution specific to one application or a narrow range of applications. There has been less progress on “general AI,” which strives to create a single, all-purpose artificial brain that could address any cognitive challenge and would be as adaptive and flexible as human intelligence. Indeed, there is no clear technical path for achieving general AI, so it appears that for at least the next decade the policy focus should be on the implications of narrow AI.

In a world of narrow AI, there will not be a single moment at which machines surpass human intelligence. Instead, machines may surpass human performance at different times for different cognitive tasks; and humans might retain an advantage on some cognitive tasks for a long time. Even if machines surpass humans in the

lab for some task, additional time and effort would need to be invested to translate that advance into practical deployment in the economy.

Successful AI does not think like a human—if it is an intelligence, it is an alien intelligence. Because AI solutions are task-specific and do not directly mimic the human brain, AI systems tend to “think” differently than people. Even when successful, AI systems tend to exhibit a different problem-solving style than humans do. An AI system might handle some extremely complex situations well while failing on cases that seem easy to us. The profound difference between human thinking and AI operation could make human-AI teaming valuable, if the strengths of people and machines can complement each other. At the same time, these differences create challenges in human-AI teaming because the teammates can have trouble understanding each other and predicting their teammates’ behavior.

On many cognitive tasks, more engineering effort or more data translates into better AI performance. Many AI systems learn from data. Such systems can be improved by re-engineering them to learn more from the available data or by increasing the amount of data available for training. Either way, devoting more effort to engineering and operating an AI system can improve its performance. Machines are generally worse than humans at learning from experience, but a machine with a very large data set has much more “experience” from which to learn. Using the narrow AI approaches that have been successful so far, expert AI developers must invest significant effort in applying AI to each specific task.

Benefits of AI/ML

AI is already creating huge benefits, and its potential will only grow as the technology advances further.

For example, AI is a key enabler of precision medicine. AI systems can learn from data about a great many patients, their treatments, and outcomes to enable better choices about how to personalize treatment for the particular needs, history, and genetic makeup of each future patient.

AI is also enabling self-driving cars, which will eventually be much safer than human drivers, saving thousands of American lives every year. Self-driving vehicles will improve mobility for elderly and disabled people who cannot drive and will lower the cost and increase the convenience of transporting people and goods.

Given the tremendous benefits of AI in these and other areas and the likelihood that the technology will be developed elsewhere even if the United States does not lead in AI, it would be counterproductive to try to stop or substantially slow the development and use of AI. We should not ask the industry and researchers to slam on the brakes. Instead, we should ask them to use the steering wheel to guide the direction of AI development in ways that protect safety, fairness, and accountability.

Policies to Support AI Progress

America’s leadership in AI has been driven by three factors: our companies, our researchers, and our talented workforce.

American companies recognized the potential of AI early on and have been investing heavily in AI and moving aggressively to hire top talent. This is the area in which our national leadership in AI seems safest, at least in the short run. In the longer run, however, industry must be able to work with world-leading American researchers and workforce to sustain its advantage.

Our lead in research and development is less secure. Federal funding for AI research and development has been relatively flat, even as the importance of the field has dramatically increased. Aggressive hiring by industry has thinned the ranks of the academic researchers and teachers who are needed to train the next generation of leaders. Although industry has carried out and supported a great deal of research, it cannot and does not cover the full spectrum. The public research community plays an important role in basic research, in research areas such as safety and accountability, and in training young researchers, so investments and policies to support and grow that community are a key enabler of continued American leadership.

The foundations of the future workforce are laid in our K-12 schools. Policies to enhance access to high-quality education for all American children, especially in computing, can grow the number of students who enter higher education eager and able to pursue studies in technical fields such as AI.

The American AI workforce has also been boosted immeasurably over the years by the attractiveness of our universities and industry to the most talented people from around the world. America has been a magnet for talent in AI and other technical fields, and that must continue if we are to retain our leadership. Policies to ensure that America remains an attractive place for foreign-born experts to live,

study, work, and start companies are among the most important steps for the future health of our AI enterprise.

Risks and Challenges of AI/ML

The benefits of AI are tempered by some risks and challenges: AI systems may pose safety risks; they may introduce inadvertent bias into decisions; and they may suffer from the kinds of unforeseen consequences brought on by any novel, complex technology. These are very serious issues that require attention from policymakers, AI developers, and researchers.

Much of the criticism of AI/ML systems centers on the risk that adoption of AI/ML will lead inadvertently to biased decisions. There are several ways this could happen. If a system is trained to mimic past human decisions, and those decisions were biased, the system is likely to replicate that bias. If the data used to train a system is derived from one group of people more than another, the result may serve the overrepresented group to the detriment of the underrepresented group. Even with ideal data, statistical artifacts can advantage larger groups to the detriment of smaller ones. Real-world examples of these sorts of biases are well-documented.

The solution is not to stop pursuing AI, but rather to take steps to prevent and mitigate bias. Practitioners should work to improve their data, to ensure that datasets are representative of the population and do not rely on past biased decisions. They should also improve their algorithms by developing and using AI systems that are more resistant to bias, so that even if flaws remain in the data, the system can produce results that are more fair. In both areas, data improvement and algorithm improvement, the research community is producing promising early results that will improve the anti-bias toolkit available to practitioners. A robust national AI research effort should include studies of algorithmic bias and how to mitigate it.

In considering the risks of bias and accountability in AI, it is important to remember that in most cases the alternative to relying on AI is to rely on human decisions, which are themselves at risk of error, bias, and lack of accountability. In the long run, we will likely rely much more on algorithms to guide decisions, while retaining the human role of determining which goals and criteria should guide each decision.

Accountability, Transparency, and Explainability

The importance of the decisions now made or assisted by AI/ML systems requires that the systems and their operators are accountable to managers, overseers, regulators, and the public. Yet accountability has proven difficult at times due to the complexity of AI systems and current limitations in the theory underlying AI. Improving practical accountability should be an important goal for the AI community.

Transparency is one approach to improve accountability. Disclosing details of a system's code and data can enable outside analysts to study the system and evaluate its behavior and how well the system meets the goals and criteria it is supposed to achieve. Full transparency is often not possible, however. For example, a system's code might include valuable trade secrets that justify withholding aspects of its design, or the data might contain private information about customers or employees that cannot be disclosed.

Even where transparency is possible, it is far from perfect as an accountability mechanism. Outside analysts may have limited practical ability to understand or test a system that is highly complex and meant to operate at very large scale. Indeed, even the designers of a system may struggle to understand the nuances of its operation. Computer science theory says that examining a system beforehand cannot hope to reveal everything the system will do when it is exposed to real-world inputs. So transparency, though useful, is far from a complete solution to the accountability problem.

Another approach to accountability is inspired by the field of safety engineering. The approach is to state clearly which safety, fairness, or compliance properties a system is designed to provide, as well as the operating conditions under which the system is designed to provide those properties. This is backed up with detailed evidence that the system will have the claimed properties, based on a combination of design reviews, laboratory testing, automated analysis tools, and safety monitoring facilities in place during operation. Rather than revealing everything about how the system works, this approach focuses on specific safety, fairness, or compliance requirements and allows system developers to use the full range of technical tools that exist for ensuring reliable behavior, including the tools that the system developers will already be using internally for quality control.

Much needs to be done to make this approach feasible for routine use. Research can develop and test different approaches to proving behavioral properties of systems. Professionals can convene to develop and pilot best practices and standards.

The overarching challenge is to understand how to relate the technical process of engineering for reliable operation to the administrative processes of management, oversight, and compliance.

Regulation and the Role of Government Agencies

There is no need to create special regulations for AI. Where AI is used in sectors or activities that are already regulated, the existing regulations are already protecting the public and regulators need only consider whether and how to adjust the existing regulations to account for changes in practices due to AI.

For example, the Department of Transportation (DOT) and National Highway Traffic Safety Administration (NHTSA) have taken useful steps, under the previous and current Administrations, to clarify how existing safety regulations apply to self-driving vehicles and how Federal safety regulations relate to state vehicle laws. These changes will serve to smooth the adoption of self-driving vehicles which, once they are mature and widely adopted, will save many thousands of lives.

Similarly, the Federal Aviation Administration (FAA) has been striving to adapt aviation regulations to enable safe, commercial use of unmanned aerial systems (UAS, or “drones”), which have benefits in many sectors, such as agriculture. The FAA has taken some steps to increase the flexibility to use UAS commercially, but the interagency process on UAS has been moving slowly. Agencies should be urged to work with the FAA to advance this important process.

Government agencies have important roles to play beyond regulation. For example, the National Institute of Standards and Technology (NIST) and the Department of Commerce can contribute by setting technical standards, codifying best practices in consultation with the private sector, and convening multi-stakeholder discussions, much as they have done in the area of cybersecurity.

All agencies should consider how they might use AI to better accomplish their missions and serve the American people. AI can reduce costs, increase efficiency, and help agencies better target their use of taxpayer dollars and other limited resources. The National Science and Technology Council’s subcommittee on Machine Learning and AI can serve as a focal point for interagency coordination and sharing of ideas and best practices.

With good policy choices and the continued hard work and investment of American companies, researchers, and workers, AI can improve the health and welfare of Americans, boost productivity and economic growth, and make us more secure. Americans currently lead the world in AI. We should not step on the brakes. Instead, we should reach for the accelerator and the steering wheel.

Thank you for the opportunity to testify. I look forward to answering any questions.

Senator WICKER. Great. Thank you so much.

Let me start with Dr. Gil. Machine learning and artificial intelligence capabilities are accelerating, and I think Americans listening today or maybe insomniacs listening at 2 in the morning 2 weeks from now on C-SPAN, know that we use AI for social media and online search queries and smartphone apps. Ms. Espinel mentioned health diagnosis, and I think she grabbed our attention there.

What other industries, Dr. Gil, stand to benefit the most from what we’re talking about today?

Dr. GIL. I believe actually artificial intelligence is going to touch every profession and every industry, but just to give some concrete examples.

Senator WICKER. OK, good.

Dr. GIL. From security and cybersecurity, there’s a class of category of problems that have to do with responding with low latency. So the nature of the problem that one has to address is too complex, there are too many dimensions of it, and AI can assist a professional detect a threat and be able to assess the proper response. So sometimes it has to do with how much time one has to make a decision, and can you be assisted?

In other areas in which, for example, in the health care profession, even though we may have more time in some occasions to perform a diagnosis or select a treatment, just the sheer complexity of the number of documents, or in this case, it could be genomic information that comes into play, goes beyond the expertise that any given person can have. So in that instance, it can assist, for example, in the process of, you know, medical diagnosis, as an example.

In agriculture, here we're talking about being able to integrate sensory measurements from soil and weather data, satellite data, to be able to predict what kind of irrigation one may have to deploy to improve the productivity of it, to give another example.

So I think that, you know, in every situation where we're trying to integrate knowledge that comes from sometimes measurements from the physical world and also bodies of evidence that we've accumulated through our expertise, combined with our own expertise, we facilitate every worker and professional to be able to make those better decisions.

Senator WICKER. OK. Now, Ms. Espinel and to all members of the panel, Ms. Espinel says that our government needs to talk about three changes to policy: open data, more government research, and prioritizing education and workforce development. Also, Dr. Felten says there's no need to create special regulations for AI.

Who wants to comment on this? Does anybody want to take issue that we have or is everybody in agreement with all four of these statements? Is there some room for nuances and disagreements? Does anybody want to respond?

Yes, Mr. Castro.

Mr. CASTRO. Thanks. So I think one of the most important things to look at in this space is it's about technology adoption. When we're looking at AI, the big question for the United States in terms of competitiveness is, are we going to be the lead in terms of adopting this technology rapidly in our industries before other countries do it in their industries? Because that's going to determine U.S. competitiveness long term.

So when we're talking about policies in this space, there are two types. There are the policies that help accelerate deployment and adoption of the technology and our R&D in this space, and then there are the regulations that might slow down adoption or, you know, kind of skew or realign how we do this adoption.

So when we're comparing ourselves to Europe, for example, which is also pursuing this, we have to ask the question: Why aren't we doing what Europe is doing to accelerate adoption? And, two, are we having smart regulations that allow us to apply it in our industries kind of better, smarter, and faster than they're doing?

Senator WICKER. Do you sort of agree with Dr. Felten, that regulations may actually impede our development of AI?

Mr. CASTRO. I think in most cases, regulation that's focused on AI specifically is probably misguided. If there's a problem there, we need to look broader and say, Why is this problem happening? And is it caused in human situations as well?

Senator WICKER. Dr. Felten, have I mischaracterized anything you've said?

Dr. FELTEN. No, you have not. I would agree with Ms. Espinel's points. With regard to regulation, in addition to sectoral regulation, there's an important role for agencies sometimes to create new regulatory structures to allow more activity, as the FAA has been working to do with drones. The FAA has been working to create new rules which allow broader commercial use of drones in the United States. And so although that is a change to regulation, it's one that enables more activity by the commercial sector.

Senator WICKER. Thank you very much.

Senator Schatz.

Senator SCHATZ. Thank you, Mr. Chairman.

Thank you for all of your testimony.

The way I see this is there's a competitiveness side of the ledger, which I think is not easy to do, but relatively straightforward. And Ms. Espinel's testimony spells out some of the steps that we can take on a bipartisan basis to make sure that we win that race internationally. That part is, again, not easy, but relatively straightforward morally and as a matter of policy.

Where I think it does get difficult is that I think I quibble with you, Mr. Castro, in the sense that I don't think we're—that the endgame here is just that we race as fast as we can in all sectors without regard to consequences and view any regulatory effort as contradictory to our national goals. I think part of what we have to do is recognize that in areas like health care and agriculture, it's a pretty much unalloyed good to have more data, to save lives, to make agriculture more productive, for instance. In defense, it's a little tricky. In criminal justice and policing, it's extremely tricky.

And so I don't think anybody in the Senate is talking about European style regulation. I think what we are saying is that this is a complex area that's going to revolutionize society and the way we interact with each other, the way machines interact with each other, and with ourselves, and if we're not careful, we could enshrine some of our thorniest societal problems in datasets. And I'm not as persuaded, Dr. Felten, as maybe you are, that we can program our way out of that.

And one of the challenges that I would just maybe ask, if we can start with Dr. Felten and go down the line, I'm worried about diversity in the industry. I think that to the extent that you have software engineers and decisionmakers both at the line level writing the code, but all the way up to project management and the people who are wrestling with some of these moral questions, are mostly white men, and I think that's not a trivial thing because they're not thinking about biases in policing. They may be thinking differently about autonomous weapons.

And so I'm wondering how you view—and I don't think this is a place for regulation, but I do think this is a place for us, as a society, to grapple with if this is going to be transformational and change everything, is it fair, is it rational, to have only, or I should say predominantly, white men in charge of setting up these algorithms that most of the rest of society can't even access because it's all proprietary?

Dr. Felten.

Dr. FELTEN. Sure. With respect to the question of whether—of the role of technology in addressing these issues of bias, I think technology has an important role to play, but it can't be the entire solution, as you suggested, Senator. What we need is a combination of institutional oversight and governance with technology. Technology can provide the levers, but we still need institutions that are able to pull those levers to make sure that the results are conducive to what our institutions and our society wants.

With respect to the question about diversity in the workforce, this is certainly an issue. The AI workforce is even less diverse than the tech workforce generally. And it's important to take efforts to improve that so we can put our whole team on the field as a nation. And I commend groups like AI for All that are working on that to you.

Senator SCHATZ. Dr. Gil and then Ms. Espinel, we have a minute and 20 seconds left.

Dr. GIL. OK. Perhaps I could address the topic of bias associated with these models because bias can be introduced at the level of the dataset, as you properly pointed out, if the data that has been collected, you know, is not representative of the whole population, in this case, it's to make the right assessments. You can introduce then a bias in the—

Senator SCHATZ. Well, just to be clear, it can be empirically valid, right? I mean, the simplest example is here's where crimes have been committed in the past, right? It turns out Ferguson, right? Load that into the dataset, it's a predictive algorithm, works every time. You go over there, you find more and more crime, and it spins and spins. And then on top of it, enshrining that bias in an algorithm, I'm not sure that—it gets more permanent than it would be if it were up to the individual judgment of sheriffs, VAs, people.

Dr. GIL. Yes. So this is a very active field of research in which we're very active that actually has also to do with in the cases in which you may have high degrees of prediction, but you're incorporating protected variables or protected individuals in the case of the assessment that you cannot use because of law, that is a variable that you cannot incorporate.

So there are ways actually to perform the data science to be able to achieve, you know, a prediction. I'm not—

Senator SCHATZ. But if a police department like deals with a vendor, and they say, "We've got a predictive algorithm, and we can't show you how this predictive algorithm works," but in back of that, you know, inside of that black box, they've got census block stuff, they've got all kinds of stuff that you would be not allowed to use in policing, how do we even know?

Ms. Espinel, how worried should I be about this?

Ms. ESPINEL. So we talk about bias, but I would actually—I would love to address that as well because it's an issue I'm really passionate about and focused on, and I know we're running low on time here, but you also have sort of led into explainability and accountability, and that's really important, too. So I will try to briefly touch on both and I'm happy to continue this conversation.

Senator SCHATZ. Yes.

Ms. ESPINEL. So in terms of bias, I think there are really two parts of it. So part of it is how the AI systems are trained, how

they're built, in essence. And, obviously, as you point out, data can be inaccurate or it can be incomplete, or it can have its own biases that are going to skew the outcomes. And there are a number of things that can be done right now to try to help with that.

So part is making sure that data scientists that are building them have the tools and the training to try to counter that. Second, and you already raised this, I think this is another area, or you have another reason why diversity in tech is so important, and I think the more experience and background you have at the table as AI systems are being trained, the more helpful it's going to be to try to avoid that.

I think, third, as has been mentioned, there is a lot of research going on in this area, so continuing to support and invest in research that will help lessen the chances of bias in AI is very important.

And last I would say, you know, to the extent bias is discovered, obviously companies should be working immediately to try to address that.

So I think there are a number of things that can and should be happening now to try to counter that.

I think there's another part of this discussion, which we've heard less about, which I think is really important, which is how AI can be used, not trained and built, but how it can be used to try to counter bias and to try to broaden inclusion. And there are a number of really interesting examples here both in terms of hiring practices and in terms of broadening inclusion for people with diseases like—or people with conditions like autism or people that are visually impaired where AI can dramatically transform their ability to interact with society and in workplaces. And so I think having more discussion about how AI can and should be used to try to lessen bias and to try to broaden inclusion is very important.

I could talk more. I know we're running low on time.

Senator SCHATZ. I think my time is up. I'll let the other members ask their questions.

Senator WICKER. Senator Schatz's time is close to expiring.

[Laughter.]

Senator WICKER. But we'll take another round.

Senator Moran.

**STATEMENT OF HON. JERRY MORAN,
U.S. SENATOR FROM KANSAS**

Senator MORAN. I have four questions, and I'm glad to know that the standard has now increased beyond the 5 minutes.

[Laughter.]

Senator MORAN. Let me quickly try to ask these four questions. First about research and development in the Federal role. A number of us on this Committee are members of the Appropriations Committee. We would think of opportunities to be supportive of this endeavor by funding of NSF, NIH, DoD. What am I missing? Is there something out there that we ought to be paying attention to from an appropriations process that supports Federal research in this regard to AI? Don't pause very long here.

Dr. GIL. No, well, in addition to the agencies you listed, I think the DOE also has an important element of it in the intersection of

high-performance computing and artificial intelligence and how the computational platform to support historical approaches to be able to do things like modeling of chemical processes and sophisticated approaches to do that to combine it with the more statistical approaches that are being enabled now with artificial intelligence and machine learning. So I think being able to combine those two disciplines in the context of the DOE would be very helpful as well.

Senator MORAN. Thank you.

The significance—the difference between private research, business research, and government research, where do we see the focus? Let me say, where are the most dollars being spent? Is the private sector more engaged than the Federal Government?

Dr. GIL. I would say at this point it would be fair to say that in the private sector, certainly in the technology world, AI is the single most important technology in the world today. So—and the levels of investment that we're all making around that is commensurate with that statement.

Senator MORAN. Yes.

Dr. FELTEN. The private sector investment in AI research and development is much larger than the Federal investment currently. I would agree with the list of agencies that you listed, Senator, and that Dr. Gil listed. And I'd also commend to you the "National AI Research and Development Strategic Plan" that was published last October that was put together by the research agencies.

Senator MORAN. Thank you very much.

As we attempt to promote STEM education, in that broad phrase of "STEM," is that sufficient to describe the kind of intellectual and academic excellence that we need in order to develop AI? Is it something more than just promoting science, mathematics, engineering, and research, the traditional kinds of STEM things, as we pursue support of education?

Ms. Espinel.

Ms. ESPINEL. So that is very important, and I think trying to ensure that every child in every state, if they want to go into tech, they have the skills to do that, and that's a real viable, realistic career opportunity for them. I think that's critical.

I think there are other things that could be very helpful as well. So one of the things that we've been thinking about is trying to modernize vocational training. So for, you know, not just for—important for very young children learning as well, but as young adults are coming out of school, then thinking about where their career path could take them, I think there's a lot that could be done to try to improve those programs as well.

Senator MORAN. If anyone has suggestions in addition to the short time-frame that I have, that Senator Schatz didn't have, please let me know. We'd like to figure out how we focus our educational support in a way that adds this new dimension or additional dimension to what kids in grade school, middle school, are learning. They're the future.

Ms. ESPINEL. That's fantastic. Thank you.

Senator MORAN. As a Kansan, I need to ask a question about agriculture. This could be Dr. Gil, Ms. Espinel, Dr. Castro, or, well, really any of you, Dr. Bethel, where is the research taking—who are the leaders in research when it comes to agriculture? Is it the

universities or is it the private sector again who is focused on large data and what it can mean to increase productivity and efficiency and better on-the-farm income? Who should I be talking to that's fully engaged in this world?

Dr. GIL. Yes. I think there is wonderful work going on across a number of universities, and we can—I can give you some more details of some specific programs that are, you know, very well tailored to this. But certainly in the private sector there is a lot of activity that has had to do with focusing in the instrumentation aspect and the measurement of fields, particularly with satellite data, being able to combine also very unique datasets.

As we've aggregated datasets in terms of I alluded before in terms of like soil characteristics, you know, evapotranspiration models that we have, weather data, and be able to combine all these layers of data together to be able to have accurate forecasts and prediction, and to the degree that we have more autonomy also in the agricultural fields to be able to, you know, irrigate with more precision—right?—or be able to use fertilizers with more precision as well, the combination of all of those factors is what is increasing productivity and what it's enabled to do that.

Senator MORAN. Dr. Felten, my final question. You indicate in your testimony there is no clear technical path for achieving general AI, you have narrow and general. Is what you're telling me is that that's more science fiction, like more of a James Bond movie than where we are today?

Dr. FELTEN. That is the case today. We can't rule out the possibility that general AI may come along far in the future, but from a policymaking standpoint, narrow AI is what we have, and it, I think, should be setting the agenda. We should be alert for the possibility that sometime down the road general AI may come, but it's not close.

Senator MORAN. Thank you all very much.

Thank you, Mr. Chairman.

Ms. ESPINEL. Senator Moran, if I may, you mentioned advances in farming technology, and since you are from Kansas, I just wanted to let you know that we put out a study earlier this year looking at the impact of software across the United States, and Kansas is one of the states we saw the biggest jump in jobs. So over 30 percent growth in software jobs in Kansas, and you're up to nearly 40,000, and part of that is farming technology, but it is other types of software services as well. So Kansas is doing great.

Senator MORAN. I wouldn't want to forget the aviation world that we live in, too, in Kansas. Thank you very much.

Senator WICKER. Senator Peters.

Senator PETERS. Thank you, Mr. Chairman.

Senator WICKER. How's Michigan doing?

Senator PETERS. Yes, how's Michigan doing?

Ms. ESPINEL. Michigan is doing really, really well.

Senator WICKER. But not as good as Kansas, I'm sure.

[Laughter.]

Ms. ESPINEL. Well, Kansas did see a huge jump, but Michigan, \$13 billion in GDP from software into Michigan. And Michigan is definitely doing better in jobs overall in terms of numbers, maybe

not quite as big a jump year to year, but Michigan is a really—is a really strong state, and our—

Senator WICKER. And I guess New Mexico really isn't even in the game.

[Laughter.]

Ms. ESPINEL. We actually had an event in Detroit last week through our foundation talking about software and tech, and very focused on the educational system in Michigan and the great things that Michigan is doing to try to advance software and technology in the state, so thank you.

Senator WICKER. Senator Peters.

**STATEMENT OF HON. GARY PETERS,
U.S. SENATOR FROM MICHIGAN**

Senator PETERS. Thank you, Chairman Wicker, for bringing up that question so I didn't have to use my time for that. That was very well done.

[Laughter.]

Senator PETERS. And you're right, Michigan is moving very aggressively in this area, and it's primarily driven by self-driving cars and what's happening in that space, which is very exciting, something that I've been intimately involved in over the last few years and months. And we now have some significant legislation moving forward for that. In fact, it's been described to me by folks in the AI space that having self-driving vehicles may be the Moonshot for artificial intelligence, that when AI can pilot a car through a complex city environment with all sorts of things happening all around it, that means AI has developed to the point where it's going to be transformative in every single industry. It's going to be a key barometer of where we are going forward. So we are pleased that that's happening in Michigan.

In fact, we had General Mattis, who mentioned the four places for technology in the country, and Michigan was one of those four. So a little different vision than a lot of people here in Washington may have for my great state, so I appreciate the opportunity for that to come up.

The question I have—and I'm a believer in all of the wonderful things that you're talking about. I believe this is the most transformative technology in a long, long time, just as it is in the auto industry. It's probably as big as when the first car came off of the assembly line. We know what happened after that, in creating the American middle class, changing everything about our economy. We think the same thing will happen with AI.

And so there's incredible promise for it, but I think we also have to be very open to the potential downside to this. And I know some of you have addressed the employment issue, and I want to just talk about that because my experience has been the folks who are big proponents of the technology downplay the employment aspects. Folks who are scared probably overplay the employment aspects. And the truth is going to be somewhere in the middle.

And I think one thing it will have an impact on that employment growth and what we've been seeing in the economy recently is we have further concentration of industry and fewer and fewer companies that have larger shares. That has actually suppressed wage

growth, it has created a less dynamic environment when it comes to new business formation. I mean, we can go through the economic arguments associated with that.

And so, Mr. Castro, you mentioned that whoever comes up—whatever companies embrace AI will have a significant technological advantage when they do that. We see that in the auto industry. It's why the auto industry is racing to be first, or at least very early, knowing that there is probably going to be fewer car companies once AI is fully implemented as well.

So my question to you and other panelists, what sort of implications will AI have for the concentration of business in those companies or those industries, or I should say those individual companies within those industries, that have the resources to be able to utilize this? And will it be more difficult for small businesses?

Mr. CASTRO. Yes, so I think there are two effects. I mean, if you talk to a company like Amazon, it's a—you know, they're using AI more than anyone, and they're growing faster than anyone.

And so in some cases, we're going to see—especially when we're talking about global competitiveness, U.S. companies growing because of the technology, and that that growth will outpace the jobs offset. Of course, that won't happen everywhere, and in many cases, what we want to see is more productivity, which means fewer workers in a given space per output.

In those cases, and what we've seen historically in this phase, is that the new jobs are not necessarily AI jobs. It's not that, you know, everyone is going to now be, you know, building self-driving cars or designing them, it's that we see more people in other professions, more doctors, you know, classrooms with higher teacher-to-student ratios. The kinds of changes we often say we want to see and can't pay for right now, we can get in the future.

We did a really interesting study earlier this year that was looking at occupational change over the last 165 years, and it's actually at the lowest change, so the least disruption, that we've seen in this entire time period. And the reason for that is we get a lot of misperceptions when we see ATMs on the corner, and you think there are fewer of these jobs, when, in reality, you have more banks, you have more prosperity. And so the job losses are usually more visible than the job creation, which is why we have the skewed perception.

Senator PETERS. But I think we go beyond job losses and actually look at wage differentials and income inequality, and that's probably what I was alluding to, is when you have increased concentration, you have less dynamism in the economy, which I think is consistent with what you just said about the job churn has gone down, it's becoming less dynamic. Many economists believe that's a big reason why we have growing inequality in this country as well.

Certainly folks who are able to get these software jobs are going to do extremely well, and God bless them for doing that, but based on trends that we already see of stagnant wages, even though there are increases in productivity, doesn't necessarily translate into everyday wages for everyday folks, that that all could accelerate at a very quick pace that we should be thinking about, and I think it's important for us to be conscious of that impact, it's not just jobs, it's income inequality.

Now, Ms. Espinel, if you want to comment on that, please do.

Ms. ESPINEL. I was just going to say briefly that, one, I think you're right, it's something we need to be thinking about. You said, or maybe you were alluding to what Mr. Castro said, in terms of businesses using AI. I guess I would say, you know, we don't think that big businesses or concentrated businesses are the ones that should be using AI. I think our hope is that AI and the technology behind it will be democratized sufficiently so that small businesses will have the ability to use it as well and help them in whatever their business objectives are.

So I guess I would say, yes, I would agree. It would be a concern if we saw AI being used primarily by just some large companies, but I personally don't think that will be the future of its deployment, and I know, certainly speaking for our members, they would like for small businesses, for the public sector, for any type of organization that is trying to make decisions and trying themselves to make decisions in a way that is better informed to be able to have the use of the benefits of AI.

Senator PETERS. Well, I would agree that that's the goal, and we hope we could democratize it, but it hasn't necessarily played out that way. It does take significant amount of capital.

And as Mr. Castro mentioned, Amazon is using AI to the greatest extent of any other company, it's growing the fastest, and that's why we have brick-and-mortar retailers that are going out of business all over the country as well. It's great for Amazon, great for AI, great for productivity for Amazon, but it may not be so great for the mom-and-pop store that's there. And I would argue that mom-and-pop store thinks they can't have an AI system like Amazon, that's just simply not realistic for them.

And so we need to be thinking. I don't have the answers, but I think we need to be thinking about that or we're going to be facing some significant societal challenges in the future. But thank you for your comments.

Senator WICKER. Senator Udall.

**STATEMENT OF HON. TOM UDALL,
U.S. SENATOR FROM NEW MEXICO**

Senator UDALL. Thank you, Mr. Chairman.

And thank you to all the witnesses here today. I think this has been very excellent testimony. And obviously there are some very positive aspects to AI. But I wanted to ask about the bots and the software used to hurt consumers. The *New York Times* recently reported on so-called "Grinch bots," software used to predict the web links of the in-demand toys and other merchandise and to purchase the goods before the general public has access to these items. This practice has caused many of this holiday season's most popular items, such as *Fingerlings*, *Super Nintendo Classic Edition*, *Barbie Hello Dreamhouse*, to be sold out online within seconds. However, one can go to eBay and easily find these same products available for increased costs, sometimes dramatically increased.

Have either of your organizations worked with retailers to prevent software bots from taking advantage of Internet deals to jack up the prices of goods? Can you identify other ways to prevent

these software bots from using machine learning to become more and more sophisticated?

Ms. ESPINEL. I'm happy to take this as the parent of two young boys that are busy filling out their Christmas lists. So—and I actually think this is an example of an area where AI can be really helpful. So Dr. Gil and others have talked about the use of AI and cybersecurity because AI is really good at looking at large amounts of data, detecting patterns, and then predicting where threats might lie.

The Grinch bots I think is an area—AI is also great at fraud detection, and there are a lot of the characteristics of these Grinch bots that are similar. So part of it is, you know, if you see—if you see unusual patterns like huge, large, amounts of purchases that are happening very, very quickly, and these Grinch bots work so they can process transactions in just a few seconds, that is a pattern, and an unusual pattern.

And AI is really good at looking at patterns like this, and then alerting the retailers so they can shut down those purchases in much the same way that AI is being used now for credit card fraud detection and can help detect unusual patterns and then, you know, give your bank or you, as the credit card user, the ability to say, "I didn't approve those, so you need to block my card." That—this type of activity I think is actually a great example of how AI could be deployed to try to detect those unusual patterns and then tell the retailers to not process those transactions and shut it down.

Senator UDALL. Yes, we hope that that happens more because you can see the consumer damage, I think every day, and in a lot of the business coverage.

Dr. Bethel, in your testimony, you spoke of the need to access more and more data, including social media and mobile device location history to help refine the uses of artificial data. I'm concerned about the privacy implications that this kind of sharing could have. Could you speak to the ways to obtain relevant data while still protecting consumers' sensitive information? And do others on the panel share my concerns on privacy?

Please, Dr. Bethel.

Dr. BETHEL. I mean, privacy, anytime you're dealing with data, is going to be an issue of concern. And I think we have to be responsible in how we use that data and how we obtain the data. So there does probably need to be some consideration given to how that data is obtained, what that information is used for, and trying to ensure the privacy of the individuals that that information is coming from. It's critical to the success, I think, of AI, and being responsible with that.

So there may need to be some regulation in that to protect the general public from harm basically is where regulation I feel is needed. But we need to be aware of how that is happening and try to put in place measures to ensure the safety of that data.

Senator UDALL. Great. Thank you.

Please.

Mr. CASTRO. Thank you. On the question of privacy, there is actually some really interesting where we're seeing, you know, if you look at kind of trends, increased use of AI can actually increase pri-

vacy for individuals because it takes personal data out of the hands of humans, which is what generally people are concerned about. There are general privacy concerns when you actually start talking to people, is that, "I don't want you personally to see my medical record. I don't want this person over here to see my financial data or use it in a harmful way." But when they say, you know, "Are you OK with a computer seeing it?" suddenly everything is OK. And you see this in a lot of sensitive areas. This is one reason why people like to shop online for certain personal items, they don't want to see a clerk at checkout, but they're absolutely fine with Amazon having that information.

So in many cases, what we want to see is how we can have policies that actually shift it so companies can guarantee the data isn't ever touched by a human, it's only touched by a computer, and that is a privacy guarantee. But then we need to make sure we can actually hold companies accountable if they ever allow that data to escape in certain ways.

Senator UDALL. Yes.

Dr. FELTEN. AI raises the stakes in the privacy debate because it boosts both the uses of your data that you might like as well as the ones that you might not like. And so the importance of protecting data privacy continues and maybe is higher in a world with AI.

Dr. GIL. Yes. And I would just say from the perspective of IBM, we take the strong view that our clients' data is their own, and the insights that derive from the application of AI when we help them to do that are their own, too. So we take a very strong view of not using data for other purposes that are not the intended ones.

Senator UDALL. Great.

Ms. ESPINEL. I'll just say very briefly, I think it's important to distinguish between different types of data and different types of AI. Our BSA companies are at the forefront of protecting privacy, and much of the AI that we built is not built using personal data. So I think it's important to bear in mind.

Senator UDALL. Thank you very much.

I yield back, Mr. Chairman.

Senator WICKER. Senator Young.

STATEMENT OF HON. TODD YOUNG, U.S. SENATOR FROM INDIANA

Senator YOUNG. I thank our panelists for being here today. I want to build on some previous line of questioning about our labor markets and their preparation or lack thereof, as we start to move into an AI-driven economy. What sort of questions should policymakers be asking right now to ensure that we optimize the skills our workforce has to the extent possible to prepare for this, in many ways, exciting and promising new technology?

Dr. FELTEN. Well, I think—I think there's one set of questions with regard to education to make sure—especially in the K-12 system, that kids are getting some basic education in computer science. Things are more in hand, I think, at the university level; it's more a matter of resources there. The more difficult questions relate to adult works and displaced adult workers and what is available in terms of retraining or apprenticeship programs to help

them to make sure that those are available and to make sure that those are backed by good data on the effectiveness of programs is important.

Ms. ESPINEL. So I would say I think there are three things to think about. One is, you know, if you look at the situation today, the Department of Labor BLS is saying that by 2020 there's going to be 1.4 million jobs that need a computer science degree, and yet we only have about 400,000 graduates in the pipeline. So we have a gap in our labor pipeline today that definitely needs to be addressed if the United States is going to stay competitive in this area. So that's one thing we need to focus on.

Second is in terms of education. I think we need to be rethinking our educational system for young people to ensure that they have access to those skills and the opportunity to acquire them so that if they decide that they want to go into tech—and not every tech job requires a 4-year college degree, and I think that's also something we need to be better about explaining—that they have a realistic ability to do so. But that is going to require some rethink of the educational system that we have now. And I think we need to be modernizing the vocational training programs we have for young people that are coming out of school.

And then the third area, which Dr. Felten referred to, in terms of people that are in the workforce now, I think we need to not just be investing in reskilling and retraining programs, but, again, thinking differently about them. I think we need to do a better job of matching the skills that people have with the employment needs that are out there across the country, which is an area where I think there's a lot of work to be done, but a lot of potential.

Senator YOUNG. Yes. With respect to this issue of the labor markets and AI, what sort of assumptions are you hearing or reading about that you think either overstate some of the forces that will be unleashed as AI continues to develop and be adopted, perhaps strike you as a bit alarmist, or understate these forces?

Dr. GIL. I think—

Senator YOUNG. Yes, sir.

Dr. GIL. Oh, sorry. Dr. Felten was, you know, very helpful in providing this journey from the narrow form of artificial intelligence that we have today to a general form of artificial intelligence that ultimately could perform, you know, arbitrary, you know, tasks and domains. We're far away from that. So very often the discussion gets framed in terms of either the policy implications or the implications to the labor market associated with that future form of artificial general intelligence that frankly is decades away at best.

So in that sense, when the conversation is framed in that lens, it does come across as alarmist—right?—you know. When we talk about the more narrow form of artificial intelligence that exists today, it's more of a focus of, what domains can it have an impact? Where are proven examples that we can do better on that? And at best, you will perform some narrow specific tasks that can be complemented with human labor.

Senator YOUNG. Mr. Castro, do you have some thoughts?

Mr. CASTRO. Yes. I mean, so the number one study that everyone thinks about in terms of jobs is the study that came out of Oxford through research of Frey and Osborne that said 47 percent of U.S.

jobs would be eliminated by I think it was 2025 or 2030. That study was—first of all, it's not peer reviewed, and when you look at the data, they use a very kind of, you know, flawed methodology to come up with this estimate.

So they list in the back in their Appendix all the different professions that they say will be automated with AI, and it's things like fashion models and barbers, which they tried to walk a robot down the runway in Japan once, and they haven't done it since.

So, I mean, you know, kind of realistically, those numbers are very much inflated, and they're not actually tied to what we're seeing in the market today. So, you know, the main thing is I think when you see these studies, and there have been a number of studies that have used their data, it doesn't reflect reality.

Senator YOUNG. I'll just close here and indicate, you know, the reason we're focused on this I think is we see some incredible potential here. There are some studies that indicate AI has the potential to increase the rate of economic growth in the U.S. from 2.6 percent to 4.6 percent by 2035. I mean, that's just amazing. That would benefit all Americans. There are some serious policy things we need to wrestle with.

I've been partnering with Senator Cantwell, who has a lot of expertise and professional background in the area of technology, and we've developed legislation that would establish a Federal advisory committee to help us better understand some of these issues. So if you have thoughts moving forward about things that a Federal advisory committee should look at as we consider the policy implications and broader market implications of AI, I'd certainly welcome those, and I suspect Senator Cantwell would as well.

So thank you, Mr. Chairman.

Senator WICKER. What a nice segue to Senator Cantwell.

Senator YOUNG. Yes.

Senator WICKER Senator Cantwell is now recognized.

STATEMENT OF HON. MARIA CANTWELL, U.S. SENATOR FROM WASHINGTON

Senator CANTWELL. Thank you, Mr. Chairman. And I do look forward to working with Senator Young on this issue and the various aspects of investigation.

I wanted to bring up applications because one of the things that I think we should be thinking about is our role as an actual user. And one of the things I'm most interested in is AI's application for cybersecurity. One of our biggest threats obviously we face now is the threat of cybersecurity in all sorts of ways. I've seen some applications by MIT and I think University of Louisville, several entities that are both finding fault in code, basically doing a better job—why wait to find out some guy forgot to do the Apache patch, forgot to do the patch? You know? I mean, one employee at that company cost everybody a lot of money because he didn't put in a patch. AI could help us find errors in code or actually in some of these areas I think predict cyber attacks.

So I don't know who on the panel could speak to that, but to me, one of the applications that I hope that we will look at is the government's use of this as it relates to combating cybersecurity.

Dr. FELTEN. There's a huge opportunity there, and it's rapidly becoming a necessity, as the bad guys are adopting AI and automation in their cyber attacks. Government and other institutions need to be using more AI in defense in order to, as you said, Senator, find vulnerabilities before they're exploited in order to be able to react at machine speed when things start to go wrong, and in order to better understand what are the possible implications of the way that systems are set up so we don't get surprised in the way that institutions nowadays too often are by both that a breach occurs and by how bad the consequences are.

Dr. GIL. Yes, there are two dimensions are essential in this topic. One of it actually has to do with securing AI itself. The very models that we create to enable these predictions are actually vulnerable to attack, and that in itself, there are many steps that can be taken to be able to secure those models. In fact, you can even extract data from a model that has been created.

And the second I mention is the one you alluded to, which is AI itself has to be an integral component to protect against other AI-powered attacks. So in a way, we are going to have AI against AI in the realm of cybersecurity because, you know, some of the bad guys are already using these techniques to be able to attack our networks with the speed and accuracy and adaptability that without the presence of AI, and Ms. Espinel alluded to this already, it will be impossible to defend against.

Senator CANTWELL. Did you want to comment on that as well?

Ms. ESPINEL. Just I'll just say briefly that IBM is among the BSA members that are focused on this. We have a number of companies that are very focused on cybersecurity and are using AI to try to detect patterns and then predict threats before they happen. So I agree with you, that I think it's an area where AI is already being used, and there is potential for it to be doing even more.

Senator CANTWELL. Well, I like the notion just in the fact that because there are things called human error, that you can, you know, use this to look for faults in code and weak points and do—because that's what someone else is doing, right? And so the sooner that you can find that yourselves and detect that and create another security layer, the better.

Is there any other government application that you think we should be investigating on? Some people say on like health statistics and things of that nature, but I don't know if you—anybody else on the panel has—

Mr. CASTRO. We did look last week into government use of AI, and one of the biggest challenges is that we're just not measuring what we're doing within government, that there are lots of opportunities, but there's not a good place where if you're an agency, not even a CIO, but just a team manager basically, who wants to start using, you know, automated calendaring, you know, how do I actually go out and do that quickly? You know, can I quickly procure this? Is there an approved list of, you know, best practices from other agencies? Is there good information sharing? We're starting to do that through GSA, but we're not there yet.

So one of the things that this Committee could hopefully help do is to really push agencies to ask—bring the CIOs in here and ask them, "What are you doing around AI? Have you identified the top

three opportunities and are you pursuing them?" Just like, you know, many agencies were directed to find the high-value datasets, pick three, and get them out there as open data, we can do something similar around AI, pick the top three opportunity targets and pursue that over the next 12 months. Similarly—

Ms. ESPINEL. I would agree with that, but I think city planning is another area where AI could be really helpful. So like traffic congestion, you know, one of the things that city planners are trying to do is optimize traffic patterns and the changing of traffic lights to try to improve traffic flow, and that's something that AI is really good at. You know, there are a lot of variables, so traffic congestion can seem like a relatively simple question, but if you think of all the variables and traffic patterns, it's actually quite complicated in taking all that data and then giving city planners recommendations for how to optimize traffic flow is something AI can do really well.

Senator CANTWELL. Thank you.

Thank you, Mr. Chairman.

Senator WICKER. Dr. Bethel, you are doing practical applications on high-risk law enforcement situations, autonomous cargo transfer, and animal-assisted therapy. Do you receive Federal research dollars for that? And to what extent?

Dr. BETHEL. The law enforcement application we have submitted to NSF. We have currently one grant under review. The Therabot was funded under NSF funding. And some other projects that we're doing are also funded through the NSF or Department of Defense.

Senator WICKER. OK. Now, Senator Schatz mentioned a real concern with regard to law enforcement. What you all are doing, though, at Mississippi State, is scenarios where there is already a threat, and law enforcement needs information about how to respond, how to get inside the building. Is there a child there? Is there something explosive there? Has anyone raised the concerns that he raised, any of your law enforcement people or victims or defendants, raised concerns about data bias?

Dr. BETHEL. Law enforcement has not mentioned any data bias because when we are using these algorithms, we are looking at more like objects in the scene—

Senator WICKER. Something is already occurring.

Dr. BETHEL. Something is occurring. We're not trying to use it to target people, we're using it because the scene—some kind of event has occurred, and law enforcement is responding. And so we are trying to provide them as much information as we can prior to them going in so that hopefully they can make better, more safer, decisions when they enter into the environment. For instance, their protocol changes completely if there's a child inside the home that they're going into. They can't use a flash bang, they can't do things they would normally do.

So we can, by sending a robot in to be able to see what's happening prior, they can make decisions, so they're probably going to end up saving lives, both civilian and law enforcement lives, because—and they will have better performance because of it. So bias hasn't so far come into our discussions when we've been looking at law enforcement applications, but we are not using it, I think, in the manner in which Senator Schatz has indicated.

Senator WICKER. Now, with regard—first of all, on the law enforcement application, how long has Mississippi State been doing this?

Dr. BETHEL. Six years. I started training with them in 2011. We train monthly.

Senator WICKER. OK. AI, according to your testimony, is only as good as the data the system receives. How has your data improved over the 6 years, and could you elaborate on that?

Dr. BETHEL. So each training we do, we do video recordings, we do sensor recordings, we do different manners of data collection. The more examples of data that we can get and obtain, the better the system is at classifying information, doing the sensor fusion, and incorporating that information to make more informed decisions. So as time has gone on, we've been able to obtain more and more samples of data to be able to use that for the systems.

Senator WICKER. OK. Now, most of the panel, if not all the panel, seems to think we lead the world, it's just a question of whether we're going to continue leading the world. Senator Schatz said he doesn't think with regard to national policy we'll take a European approach. Rather than ask him what he means by that, let me ask the panel, are there mistakes that our international competitors, whether in Europe or Asia or Africa or wherever, are making that we need to avoid? Are there overreaches in terms of regulation that we need to avoid?

Everybody is trying to get ahead. The testimony is that China is making an important effort, the U.K., Japan, our neighbor to the north, but in terms of things we can avoid, mistakes that other countries have made, does anybody have a suggestion that we need to be mindful of as we work on Senator Young and Senator Cantwell's and Senator Schatz's legislation?

Yes, Mr. Castro.

Mr. CASTRO. Yes, so, I mean, the two big things I think some countries are thinking about is the right to explanation that's across the board or right to opt out of automated decisions. So the problem with those two policies is that it minimizes the ability or limits the ability of companies to deploy AI and many commercial applications.

So, for example, if you want to be a company, you know, just like before you had companies suddenly become online lenders, if you want to be a company that's an AI lender, streamline your entire business process through AI, you basically can't do that because anyone who got rejected for a loan application, for example, could ask for a human appeal, and you would have to have human support to do that.

And the problem with that kind of framework is it basically limits those types of business models, and it's especially problematic when you talk about global competition because it doesn't limit necessarily a foreign competitor from offering a similar service. Now, in certain industries, of course, you have to be licensed in the United States, so there are limits there, but overall, that would limit these kind of opportunities.

Ms. ESPINEL. Can I just say briefly? I think most other governments are still considering what their AI policy environment is going to look like. So many governments are having this discussion

right now. I do think that there are governments that are—seem to be skewing toward considering a more regulatory approach, a more heavy-handed regulatory approach, and that raises concerns.

And I think one issue I'd raise in particular I think for us, I think any regulatory approach that would compel the disclosure of an algorithm, that would compel companies to hand over source code or algorithms, to a government agency is one that raises a lot of concerns with us. We don't think it's actually going to be an effective way to address policy outcomes, and it raises a lot of competitive concerns to be handing over algorithms or source codes to governments.

Senator WICKER. Dr. Gil, on cyber attacks from hostile forces internationally and our defense against that, do I understand you to—well, no, let me rephrase it. Is it conceivable that artificial intelligence will be empowered by a hostile force to make a decision to go forward with a cyber attack without a human being at that end making the decision to pull the trigger?

Dr. GIL. Humans will definitely design AI-powered attacks to attack, you know, infrastructure or, you know, an industrial setting or another nation-state—

Senator WICKER. But whether that attack is made today or Valentine's Day next year, an artificial intelligence—

Dr. GIL. System.

Senator WICKER.—system might make that decision.

Dr. GIL. That's correct. That is possible—

Senator WICKER. How close is that to reality?

Dr. GIL.—because, I mean, think about it in the past for—you could design it with an explicit program model to be able to carry out such attacks, and we would have to, you know, create defense mechanisms against those kinds of attacks. But because the program was well stipulated, you could imagine somebody defending against the attack to be able to interpret what those rules of attack were.

Now, the moment that you're employing a more machine-learning-based approach where the type of attack could morph depending on the environment it's detecting, now being able to detect what is the form of the attack that is taking place requires another pattern detection mechanism. So that's why I was referring to—and Ms. Espinel was talking before—about you need AI to defend against AI-powered attacks because it's the only way to make it really adaptive to an adaptive attack.

Senator WICKER. OK. Senator Markey, do you mind if we—if we let the two witnesses respond here? And then I'll be generous with your time.

Dr. Felten?

Dr. FELTEN. Thank you. What you're talking about in terms of automated cyber attack is something we see already with computer viruses. A virus is software which autonomously spreads itself from place to place, and then does whatever it's programmed to do to cause harm at the places that it infects. So to bring AI into this would just be to have a more sophisticated, more adaptive form of virus. It's not fundamentally a new thing, it is—it's a path that we are already on, or a path that the bad guys are already on.

Senator WICKER. And quickly, Ms. Espinel.

Ms. ESPINEL. I still think it's more akin to making a recommendation than making a decision. Of course, you, the person, can decide that whatever recommendation the system gives you, you're going to have it automatically act on, but you still made that decision up front, and I don't think we should be—as a general matter, I don't think we should be abrogating decisionmaking authority, but I think that what you're talking about really is AI making a recommendation, and then the person who designed it deciding whether or not they're going to accept that recommendation automatically or not.

Senator WICKER. We're really talking about what's coming at us. Senator Markey.

**STATEMENT OF HON. EDWARD MARKEY,
U.S. SENATOR FROM MASSACHUSETTS**

Senator MARKEY. Thank you, Mr. Chairman, very much. I thank you all for being here.

The digital revolution's newest innovations—augmented reality, autonomous vehicles, drones—all of these industries use artificial intelligence, which rely heavily on a free and open Internet. Regrettably, these disrupting technologies may be dealt a major blow in their infancy, and that's because in just 2 days, the FCC will vote on a proposal to eliminate net neutrality.

Without enforceable net neutrality rules in place, broadband providers, like Comcast, Verizon, AT&T, and Spectrum, could block or slow down the content of innovators and businesses using AI all in an effort to leverage their gatekeeper role to favor their own content and generate additional profits. And what will replace these robust net neutrality protections? Nothing, absolutely nothing is going to replace those rules.

Dr. Felten, how would eliminating net neutrality protections impact the deployment and development of innovative technologies that use AI?

Dr. FELTEN. A lot of AI technologies operate within an institution or within a company's data center, and so those technologies would not be much affected by changes in network policy. To the extent that changes in network policy affect the ability of companies to deliver their products to consumers, that would obviously be a policy concern, but in my mind, it is important but also somewhat separate from the issue of development of AI.

Senator MARKEY. So you don't think that the additional cost to the developer, to the innovator, won't have an inhibiting impact upon their ability to go to the capital markets, raise the dough, in order to produce their innovation knowing that there's no guarantee that they can reach the 320 million Americans without paid prioritization or without a threat of throttling or blocking?

Dr. FELTEN. I would say that all pro-innovation policies and all policies that are designed to help small companies enter and compete provide value and are pro-innovation and valuable. As I said, I don't think that net neutrality plays a special role with respect to AI.

Senator MARKEY. But does it—

Dr. FELTEN. There are other areas of innovation.

Senator MARKEY. Will it play a role in your opinion?

Dr. FELTEN. Yes. I think those decisions do play a role in almost any area of innovation.

Senator MARKEY. OK. And on the question of child privacy, the Child Online Privacy Protection Act of 1998 is still the communications constitution for safeguarding children online, and that's a law I was able to get passed back in 1998.

As emerging technologies like AI are deployed, it's important that they honor core American values, including privacy. Dr. Felten, could AI technologies pose a threat to children's privacy? And is there a threat that AI technologies could produce inappropriate content for children?

Dr. FELTEN. This is an issue to pay attention to. AI does raise the stakes on policy—I'm sorry—AI raises the stakes on privacy discussions generally, and that's true with respect to children and others. And, of course, parents are very concerned about what their kids see and what happens, and that's one of the reasons why COPA, for example, requires parental consent before certain uses of data are allowed.

Senator MARKEY. So could relying on AI in children's toys negatively impact kids' ability to develop empathy if we substitute real people with computers that cannot fully understand emotion as humans do, Dr. Felten?

Dr. FELTEN. Well, I think kids are more interested in playing with other kids or using toys as a vehicle for playing with other kids. I'm less worried about kids bonding with something that is not human-like or not companionable. I think kids will reject those on their own.

Senator MARKEY. Yes. Earlier this year I wrote to Mattel with serious concerns about their plan to bring the first all-in-one voice-controlled smart baby monitor to the market. Mattel had planned for the device, *Aristotle*, to use artificial intelligence to help teach children and respond to their needs. After an outcry of questions, Mattel canceled that product.

Dr. Felten, what does that experience expose as to potential negatives of using artificial intelligence with children's devices?

Dr. FELTEN. Well, I think stories like this illustrate that it's important to understand the implications of the technologies that are being deployed. As I would imagine, that parents would be happy to, say, be notified if there is indication that their child is in distress, and AI may help to do that more effectively. But these issues of unintended consequences and safety are paramount, and that's one of the important aspects of clearing the road for responsible deployment of AI, is to make sure these issues are taken care of.

Senator MARKEY. Thank you.

Senator WICKER. Senator Cruz.

STATEMENT OF HON. TED CRUZ, U.S. SENATOR FROM TEXAS

Senator CRUZ. Thank you, Mr. Chairman.

Thank you to each of the witnesses for coming here this morning to testify.

A little over a year ago, the Subcommittee on Science and Space, which I chair, held the first congressional hearing on artificial intelligence. And then, as now, we heard testimony about the ex-

traordinary transformative process that we are engaged in right now and how AI in time can be expected to touch virtually every area of human endeavor, and indeed that this transformation may be of comparable import to the transformation we engaged in, in a prior era in the Industrial Age.

Anytime we're seeing dramatic transformations in our economy and our workforce and how we interact with each other, that poses the risk of dislocations, but it also poses policy and government and regulatory challenges for how to interact with the new terrain. In your judgment, what are the biggest barriers right now to developing and expanding AI and its positive impacts on our society and our economy?

Ms. ESPINEL. I'll head off another—so I think one of the biggest barriers is a lack of understanding, a lack of understanding about what AI is, what the actual technology is, and then what it does and what the—you know, what both the intended and unintended consequences are. And so I think, you know, this hearing, the legislation that Senator Schatz is working on, I think trying to increase our collective understanding is critical, it's fundamental.

I think there are a number of specific policy issues that would be helpful in terms of eliminating barriers. So, you know, one of those is AI is all about data, and so good data policy in various ways I think is very important. I think investing in research, which we've talked about already, investing both in government research and incentivizing private sector research, is very important. And then I think thinking about jobs and workforce development, both the jobs today, but what will happen tomorrow? And rethinking our educational system and our training and reskilling programs are vitally important. So those are the three specific areas, but I think a greater understanding needs to be part of all of those discussions.

Mr. CASTRO. So I agree that skills are very important. Especially much of this technology is new, and we need people that can basically rapidly be credentialed in how to deploy it. But when we look at the biggest opportunities for AI, it's really in some of the regulated industries. And so I think that's where we head up on challenges because there are two challenges: one, regulators aren't necessarily prepared with how to deal with this; and, two, they don't necessarily have the skill set or capabilities internally within the regulatory system to handle it. And so I think that's something we need to be very focused on, is asking questions—

Senator CRUZ. What agencies in the industries in particular?

Mr. CASTRO. So financial regulation, for example. Education is another example, especially when we're talking about using AI in primary education, there are a lot of questions about privacy that get raised. And the financial system it's the questions of, you know, do we have regulation basically to understand the technology? And we'll do things more than just ask, "Can I see the algorithm?" but be able to say, "Can I look at outcomes? Can I actually measure outcomes?" and then ask questions about, "OK, is this fair? Is this different than what we have now? Is this moving in the right direction?" and also have some of that regulatory flexibility. So we need kind of, you know, fewer cops and more sandboxes.

Senator CRUZ. So one area that has generated fears and concern is general AI, and scientists and innovators ranging from Stephen

Hawking to Bill Gates to Elon Musk have raised concerns. Stephen Hawking stated, quote, "Once humans develop artificial intelligence, it would take off on its own and redesign itself at an ever-increasing rate. Humans, who are limited by slow biological evolution, couldn't compete and would be superseded." Elon Musk has referred to it as, quote, "Summoning the demon."

How concerned should we be about the prospect of general AI? Or to ask the question differently, in a nod to *Terminator*, when does Skynet go online?

[Laughter.]

Dr. FELTEN. Hopefully never.

Senator CRUZ. That's the right answer.

Dr. FELTEN. I think there's a lot of debate within the technical community about how likely these sorts of scenarios might be. I think virtually everyone agrees that it would be far in the future. And generally the people who are most involved in AI research and development tend to be the most skeptical about the Skynet or existential risk type of scenarios.

In any case, the sorts of risks and concerns that exist now about AI are really baby versions of the ones that we would face with a more sophisticated general AI. And so the tactics we—the policies that make sense now to deal with the issues we face now are the same ones we would use to warm up for a general AI future if it comes. And so from a policy choice standpoint, the possibility of distant general AI seems less important to me.

Dr. GIL. Yes, I would completely agree. I think if you ask practitioners in the field when they would envision the possibility of that, I think everybody would say 20-plus years out. And whenever scientists says 20-plus years out is our code word for saying we just don't know. Right? We're nowhere near close to be able to do that.

So I do think that while it's an area of very important study, and there are many universities who are now creating a rate of monitoring of what is the progress of AI and the implications that it would have if we eventually reach general artificial intelligence, I do think it would be a mistake to guide our policy decisions at present based on the sort of like long-term hypothetical that we don't even have, even as practitioners, even a credible path to get there.

Senator CRUZ. Thank you.

Senator WICKER. Senator Cortez Masto.

**STATEMENT OF HON. CATHERINE CORTEZ MASTO,
U.S. SENATOR FROM NEVADA**

Senator CORTEZ MASTO. Thank you, Mr. Chairman.

And thank you to the panel members. I'm very excited about the conversation today as well.

Obviously, we are standing at the edge of a technological revolution that we must ensure will, and there has been discussion, take our labor force with it. So this is a timely conversation.

Workforce development has been a focus of mine since I've entered the Senate, and so has innovation. I've proudly been leading a bipartisan legislation on drone expansion, the use of smart technology in transportation, and in trying to spur the next generation of women to be equal at the forefront of STEM, specifically com-

puter programming, through the introduction of the Code Like a Girl Act. I've worked on this because I've seen the future of these developments through my state in Nevada.

And just last week I was visited by the leadership at the Truckee Meadows Community College in Reno, who, in partnership with Panasonic, has developed a curriculum to provide individuals to get the specific training that local jobs are hiring for. These are conversations that constantly are going on in my state.

So let me start here, Mr. Felten and Ms. Espinel. The skills gap has been discussed extensively today, and something that obviously is on all of our minds and how we address it. Obviously, we have China investing dramatically in the area of AI, so it begs the question of whether there are investments we can be making at the Federal level to help close any potential skills gaps. I'm curious your thoughts on that.

Ms. ESPINEL. So in terms of the skills gap specifically, we also talked a little bit about Federal support for research funding. But I would say a few, and then others may add to them.

I think so one is trying to improve access to computer science education at all states and at very early stages of education. So I think that's one area where the Federal Government could be helpful.

I think the second is in rethinking how our vocational training programs work. There are vocational training programs in place, but they could be streamlined and they could be better adapted to the world that we live in today. So I think that's an area where there's a lot that could done at the Federal level.

And then the third I would say is that I think there is now, and there will maybe be only an increasing sort of information gap between the skills that people have and employers knowing about those, and that seems like an area where there's a real deficiency now, so therefore real opportunity to try to create programs that not only create—either create pathways directly into employment or do a better job of matching up skills that people have with the jobs that employers have to offer.

Senator CORTEZ MASTO. Thank you.

Dr. FELTEN. There are opportunities to widen the pipeline at all stages starting with K-12, making sure that basic computer science education is available to every child at that level, and more advanced education at the high school level for those who are ready for it.

There are opportunities to improve—increase the number of teachers and trainee researchers at the university level, and that's education and research funding specifically in areas of AI and computer science at the university level. And then vocational and adult training and apprenticeship programs also are very important to get people on other career paths, give them an on-ramp into this area.

Senator CORTEZ MASTO. Go ahead.

Ms. ESPINEL. Can I just say one more thing?

Senator CORTEZ MASTO. Yes.

Ms. ESPINEL. There's a lot that BSA companies are doing in this area, IBM among them. And so I think another area is working

with the Federal Government both to scale up programs that are effective now and be more collaborative in this area I think is—

Dr. GIL. Just to touch on just one example of that. On the initiative, on a program, we started a number of years called P-TECH, which is a 9 through 14 educational program that combines both education, hands-on program on vocational, you know, schools. And when they graduate through this—right?—we are talking about the creations of new collar professions that have enough skills to be able to practice and benefit from the advances that are happening in AI without having to go through a full college education, and now that's touching, you know, more than 10,000 students.

Senator CORTEZ MASTO. Thank you. No, this is something that's happening in Nevada now. It seems like it's such common sense, but we just don't do it, which is that collaboration and partnership between whether it's a government, private sector, and then our education system or our skills vocational trade, to really say, "What are the jobs of the future? What are they going to look like?" work with those employers to develop the curriculum that you're going to need for that skilled workforce so that we can really start educating them now. And that starts I think at a very young age and working with our education system, but also the vocation trades that are out there as well. Not every child is going to go on to get a higher degree, but there are some that are absolutely and rightfully are going to go get that vocation or that trade and that skill that's necessary.

So I appreciate the comments today. I am running out of time, so I will submit the rest of my questions to this incredible panel for the record. And I appreciate you being here. Thank you.

Senator WICKER. Thank you very much.

Senator Blumenthal.

**STATEMENT OF HON. RICHARD BLUMENTHAL,
U.S. SENATOR FROM CONNECTICUT**

Senator BLUMENTHAL. Thanks, Mr. Chairman. I apologize that I missed some of your testimony so far, but I know that this panel has been very important and enlightening, and I thank you for being here.

As you well know, the success of autonomous vehicles is closely linked with the success of AI. An autonomous vehicle can only perform as well as the input it receives. A lot of us know the statistic that 94 percent of the 37,000 deaths on the road each year are attributable to human error—97 percent. The hope is that autonomous vehicles will eliminate that human error.

What we really lack is information on the extent, to what extent that is caused by human error will be possibly replaced by computer error. And we know computers are not infallible.

I wonder, Dr. Bethel, whether you could talk about some of the tasks in which humans still perform better than computers or an AI system in the context of autonomous vehicles, if you're able.

Dr. BETHEL. In the context of autonomous vehicles, humans are able to adjust to very rapidly changing, unpredictable environments, things happening in the environment. Our sensor systems and our onboard processing in autonomous vehicles is just not to a point where it can make those kind of adjustments that rapidly

currently with current technology to be able to adjust to that. So in those cases where you have an erratic driver who is doing unpredictable behaviors, it's really hard sometimes for the system to be able to detect that and react accordingly. In those cases, a human currently is making better decisions on autonomy—over an autonomous vehicle in that kind of environment.

Senator BLUMENTHAL. So if I can just extrapolate from your answer, a computer trying to deal with a drunk driver, either ahead or next to that computer, would have trouble because the drunk driver, by definition, is acting in not only unpredictable, but irrational and sometimes actually self-destructive ways.

Dr. BETHEL. Right. So it would be much more difficult for a computer to predict that kind of behavior than it would be for a human to slow down and react. I mean, it would react, but it probably is not going to be as effective as a human driver is at this point in the stage of where AI is.

Senator BLUMENTHAL. Are there ways to program a computer or create software that deals with those unpredictable situations?

Dr. BETHEL. To some extent, but I think there's a long way to go on that.

Senator BLUMENTHAL. In your testimony, you say, and I'm quoting, "A current limitation to the advancement of artificial intelligence is the quality and cost effectiveness of sensing capabilities to provide high-quality information or data to the system to make those digital decisions. We have a long way to go in the advancement of artificial intelligence—," "We have come a long way in the advancement of artificial intelligence," excuse me, "however, we still have a long way to go."

Sensing, perceiving, interpreting, surroundings are essential to driving a vehicle. Can you describe some of the limitations in the current computer vision and sensing technologies?

Dr. BETHEL. I'll probably need to get back to you on some of that because it's not exactly my area of expertise related to computer vision, but there are limitations in the sensing capabilities we currently have. Every sensor has its own limitations, so there's no perfect sensor out there.

There are also differences in processing power, so trying to be able to handle large amounts of data coming in from these sensors, and processing that in a timely manner can be a challenge. So that's another area. And computer vision has just mega amounts of data coming in that has to be processed, and so another limitation is the actual processing power to handle that, especially in a timely real-time manner onboard a system.

Senator BLUMENTHAL. At what point do you think it would be appropriate to completely remove the human from an autonomous vehicle?

Dr. BETHEL. Thank you for the question. I—depending on the application, I think there are applications currently where autonomy could be—fully autonomous systems are capable. I think it's not realistic anytime in the near future, and especially in autonomous cars, to be able to say that a fully autonomous car is going to be possible all the time.

Senator BLUMENTHAL. Thank you.
Thank you very much to the panel.

Thank you, Mr. Chairman.

Senator WICKER. Senator Schatz.

Senator SCHATZ. Dr. Felten, I just wanted to follow up on the question around autonomous weapons. It seems to me that this is an area that is different than a lot of these other ethical, societal, micro, macro, economic questions; this is about how we engage in warfighting. And so to the degree and extent that some of these algorithms are hackable, to the degree and extent that we have a body of history around how we're supposed to engage our military in an ethical manner, that this is an area where the Federal Government has to make policy.

And I'm wondering if you can give us any insight into, in the absence of policymaking, who's making these decisions? Is it defense contractors? Is it individual procurement officers? How is all of this getting decided?

Dr. FELTEN. Well, there does need to be a policy which deals with the very serious issues that you mentioned, both from the standpoint of what our military should be willing to do and what safeguards are needed on the systems that they're using consistent with the need for them to be as effective as possible in battle, and also how we deal with adversaries, who might not be as scrupulous in following the international humanitarian law.

And this is a national policy issue that ought to be a matter of policy discussion at the highest levels. And if it's done in a decentralized way, if each contractor, each contracting officer, does one thing, if the State Department goes to international arms control discussions and does their own thing, we get uncoordinated policy and we get a result that doesn't serve the American people well.

Senator SCHATZ. Thank you.

Senator WICKER. Well, I would like to thank everyone who has participated. And I think Senator Cortez Masto said this has been an incredible panel, and I would have to agree.

Ms. Espinel, you brought information about software and the way a number of states have benefited. Did you have information for all 50 states or just for the states represented by this panel?

Ms. ESPINEL. We have it for all 50 states. Our foundation software data put out a study in September for each of the 50 states.

Senator WICKER. OK. Well, then if you would, please enter that into the record.

Ms. ESPINEL. I'd be delighted to.

[The information referred to follows:]





Contents

- 1 The Findings: At a Glance
- 2 Letter From BSA | The Software Alliance
- 3 Software Is Making a Profound Impact on Our Economy
- 5 What Is Software?
- 7 Software-Driven Data Tools Are Extending Breakthrough Benefits Both Far and Wide
- 11 Software's Economic Impact Is Now Rippling Throughout the Economy
- 12 Charting a Path to an Even Brighter Economic Future
- 13 State Economic Impact
- 15 Letter From Third-Party Economists
- 16 Methodology

This report was written by BSA | The Software Alliance, incorporating the analysis done by The Economist Intelligence Unit (EIU). The EIU compiled these data and economic impact assessments using publicly available government data, maintaining full editorial control of the process and using industry standard approaches. Any views or opinions expressed in this document are not necessarily those of The Economist Intelligence Unit.



The Findings: At a Glance

UNITED STATES^a

Software is so much more than your desktop at work. Software is apps. Software is data. Software is cloud computing. It creates breakthroughs and drives growth in nearly every industry. Software empowers countless people and American businesses, and improves our lives each day in ways big and small. Along with all this progress comes the dramatic, positive impact software is making on our national economy each year. *The Economic Impact of Software*, a first-of-its-kind study from BSA | The Software Alliance conducted in 2016 by The Economist Intelligence Unit (EIU), captures the breadth of the software industry in the US and the sweeping economic impact it is making at state and national levels.

Total Value-Added GDP:

\$1.07 trillion

(includes indirect and induced impacts)^b

Direct Value-Added GDP:

\$475.3 billion

JOBS EMPLOYMENT

Direct:

2.5 million jobs

Total:

9.8 million jobs

(includes indirect and induced impacts)

From software developers and web designers to futurists, project coordinators, administrative assistants, and accountants, software creates jobs for a wide variety of professionals in today's workplaces. These numbers capture jobs created directly by the software industry, as well as jobs the software industry supports through indirect and induced impacts.



WAGES

Average Annual Wage for Software Developers:

\$108,760^c

A software developer's wage is more than twice the average annual wage for all US occupations, which was \$48,320 in 2015.^d



RESEARCH & DEVELOPMENT

\$52 billion^e

R&D Investment by Software Companies
17.2% of All Domestic Business R&D in the US^f

From developing new data analytics to driving breakthrough technologies like cognitive computing, the software industry's commitment to R&D continues to spur innovation at unprecedented rates.

^a All data is from 2014 unless otherwise indicated.

^b For definitions of "indirect" and "induced," see Methodology section on page 16.

^c US Department of Labor, Bureau of Labor Statistics, Occupational Employment Statistics. Data from May 2015.

^d Ibid.

^e National Science Foundation/National Center for Science and Engineering Statistics and US Census Bureau, Business R&D and Innovation Survey, 2012 Industry breakdown. Where data is not available for 2012, the most recent year is used.

^f National Science Foundation/National Center for Science and Engineering Statistics.

LETTER FROM BSA | THE SOFTWARE ALLIANCE

Almost every day, software breakthroughs are unlocking new opportunities to improve our lives in ways both big and small. Following years of substantial investments in software innovation, unprecedented technological advances now are transforming nearly every aspect of our lives, growing our economy, and greatly enhancing our lives. But for too long, software's significant economic impact has gone unseen or undermeasured.

To gain a better understanding of the software industry's role as an engine for economic growth, BSA commissioned The Economist Intelligence Unit to conduct an analysis, the first of its kind, on the economic contributions of software. We set out to help answer two fundamental questions: How broad of a role does the software industry directly play in growing the economy? What kinds of multiplier effects does software innovation play in the sweeping impacts being seen at both state and national levels?

The study's findings are profound. It finds that today in the US alone, the software sector is driving more than a trillion dollars a year in economic impact. Additionally, the software sector drives high-paying jobs, with software developers making more than twice the average annual wage for all US occupations. As software innovation continues to thrive, it is creating a ripple or multiplier effect that stimulates other parts of the IT sector and the economy. This dynamic benefits all 50 states in meaningful ways.

The software industry is poised to build upon today's critical contributions by driving even more widespread growth in the years ahead. The study shows that through investments in R&D, the software industry is investing in America's innovative future — laying the groundwork for a new wave of software-driven technological advances that promise to make businesses more productive, jobs more plentiful, and opportunities more pervasive.

This analysis can help thought leaders and policymakers better gauge the magnitude of the impacts that a thriving software industry can deliver. At the state level, the study's economic rankings can help state leaders better understand the relative success of their technology initiatives, and better focus on the smart policy choices that ignite broader software-driven economic gains.

Now, more clearly than ever, the software revolution isn't just the vital spark that can improve our lives and transform the way we work. It's an economic engine for creating the new jobs, industries, and solutions that are essential for bringing about a brighter future. If we want to rise and meet the challenges of both today and tomorrow, we need to seize upon what may be software's greatest untapped potential — its ability to fundamentally expand tomorrow's economic promise.

Sincerely,

Victoria A. Espinel
President and CEO, BSA | The Software Alliance



Software Is Making a Profound Impact on Our Economy

Software innovation's positive economic impact can be seen throughout our national economy — contributing to our GDP, creating jobs, lifting wages, and securing our future. To measure these impacts, we launched a unique analysis to explore whether these impacts can be seen across the country in all 50 states, whether software's growing economic potential carries through to other sectors of the economy, and whether these benefits extend to local economies, entrepreneurs, workers, consumers, and governments alike.

As part of BSA's *Economic Impact of Software* study, researchers from The Economist Intelligence Unit (EIU) sought to quantify software's economic impact by examining:

- ➲ The software industry's breadth as a direct driver of economic benefits; and
- ➲ The economic impact software indirectly plays in the sweeping gains being seen at both the state and national levels.

KEY FINDINGS¹

Taken together, the study's results indicate the software industry's impact is now rippling throughout the economy, benefiting all 50 states, boosting efficiency, creating jobs, and expanding our economic potential across numerous sectors.

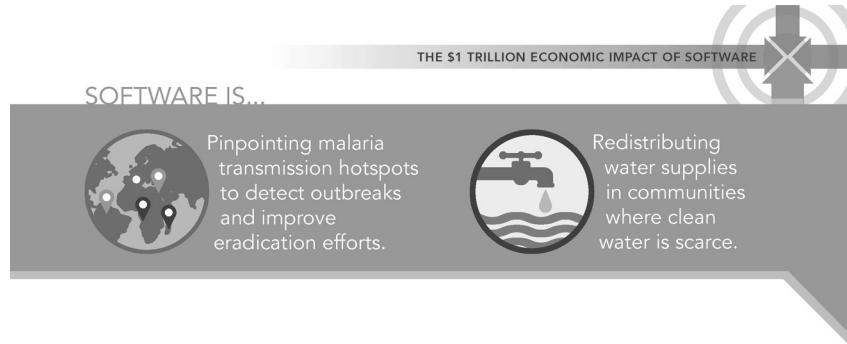
- ➲ **Software innovation adds more than a trillion dollars a year to the US economy:** The software industry directly added \$475.3 billion to the US economy in 2014. When factoring in its full effect (both indirect and induced), software was responsible for a total \$1.07 trillion of all US value-added GDP in 2014.²
- ➲ **Software supports nearly 10 million jobs nationwide:** From app developers to accountants, and from web designers to project managers, software creates jobs for a wide range of professionals across the US economy. While the software industry itself directly employed 2.5 million people in the US in 2014, research shows the software industry supports a total of 9.8 million jobs when including the indirect and induced jobs that software supports.
- ➲ **Software lifts wages through jobs that can pay more than twice the national average:** Not only does software support millions of good jobs, but these are great-paying jobs, too. For example, software developers earned \$108,760 on average in 2014 — more than twice what the average worker earns — \$48,320 a year — for all US occupations.³



Among the many results from the study, one thing is abundantly clear: software is a powerful catalyst for economic change with the potential to make businesses more productive and the economy more prosperous.

- ➲ **Software is driving economic gains in all 50 states:** Every state's economy and workforce are benefiting from new jobs that fit our modern digital economy and opportunities driven by software advances. The software sector in large states, including California, Texas, and New York, drives about a third of software's trillion-dollar economic impact on the GDP, while states like North Carolina and Massachusetts are benefiting from large investments into game-changing R&D in their states. Yet every state, no matter how big or small, rural or urban, is seeing tangible economic impacts from software innovation.
- ➲ **Software industry investments in the future will help expand America's long-term economic and innovative potential:** Given that well over half the economic growth in this country since the end of World War II has been directly attributable to technological innovation,⁴ by investing more than \$52 billion into cutting edge R&D in 2012,⁵ the software industry is helping to "pay it forward" by supporting breakthrough advances that can spur further innovation and an even more prosperous future. While the software industry makes up 2.7 percent of our economy, it contributes outsized sums to our future (17.2 percent of all domestic R&D).⁶

Among the many results from the study, one thing is abundantly clear: software is a powerful catalyst for economic change with the potential to make businesses more efficient and the economy more prosperous. Harnessing this potential in pragmatic ways can be the key to lifting our standards of living and accelerating American prosperity.



What Is Software?

Software isn't just the vital code that brings your desktop computer to life, or that puts the "smart" in your smartphone. Software is the foundational language at the very heart of our digital revolution — transforming data into ideas, and ideas into action. It brings technology to life in countless ways, and empowers us to do everyday things in ways that only recently seemed unimaginable. Software is:

- ➊ Helping people hear for the very first time.⁷
- ➋ Building smarter homes that learn from us to save energy.
- ➌ Pinpointing malaria transmission hotspots to detect outbreaks and improve eradication efforts.
- ➍ Redistributing water supplies in communities where clean water is scarce.
- ➎ Routing travelers and packages more efficiently so they reach their destinations sooner.
- ➏ Connecting families and friends worldwide in new and novel ways.
- ➐ Helping farmers monitor crops, pests, and soil conditions.
- ➑ Empowering doctors with new visualization tools that recognize cancer tumors earlier.

Software advances are improving our lives in thousands of ways every day. We use software almost ubiquitously throughout our daily lives for everything from helping us with simple tasks and conveniences to fundamentally expanding our human potential. For example:

 <p>The apps we use every day are software. When you use apps to connect with friends, check the weather, find the fastest route to work, share the latest viral video, or bank online — you are using software. These software apps have now become vital to our daily lives.</p>	 <p>Data is driven by software. Data's transformative impact is being driven by software innovations that enable everything from smarter cities and better weather predictions, to life-saving health breakthroughs and improved crop yields. Already, 90 percent of business leaders cite data innovation as one of the key resources and a key differentiator for businesses on par with basic resources like land, labor, and capital.⁵</p>
 <p>Cloud computing is enabled by software. Cloud computing — enabled by software — is transforming the way companies work by revolutionizing how computing power and storage is bought, sold, and delivered. It enables companies both big and small to affordably access data, share computing power, and collaborate in new ways — at any time, from any device, from anywhere around the globe.</p>	 <p>Industry solutions are powered by software. Software helps optimize, analyze, and visualize everything from aerodynamics systems to our global financial infrastructure. Throughout our airways, railways, and roadways, software is helping make transportation more efficient to save time, money, fuel, and lives. In classrooms, software is transforming the way teachers teach and students learn. In manufacturing, software is improving the way products are designed, built, and distributed.</p>

These are just a few examples of the pervasive ways we benefit from software. One particularly exciting advance is the way these benefits are being magnified and amplified throughout broad sectors of the economy. New software advances are sparking unprecedented advances that extend benefits more deeply into almost every sector of the economy — to make businesses work better, introduce new jobs, and make opportunities more profound.

The New York City Fire Department is combining 7,500 data points across 17 city-agency data streams with the help of clever software that uses predictive analytics and artificial intelligence to predict which of New York City's 1 million buildings are at high risk for fires.

Software-Driven Data Tools Are Extending Breakthrough Benefits Both Far and Wide

The economic impacts identified in this study are more than just numbers. Software's positive impact can be seen all around us, but perhaps none more significantly than the way software-driven data innovation is leading to benefits throughout broad sectors of the economy. Over the past decade, software innovation has unleashed unprecedented advances, grown our economies, improved our security, and increased our standards of living.

Data analytics is the software tool used to synthesize and make sense of data. It helps us use data to make more informed decisions across a range of disciplines, or to discover unexpected insights from within seemingly unrelated data. These software tools enable startups, governments, small businesses and large enterprises to sift through mountains of data to find nuggets of "information gold." This gives companies a competitive advantage, researchers a breakthrough moment, and governments the ability to better serve their citizens more effectively and efficiently.

Software is cultivating agricultural growth

To increase crop yield and boost production, farmers are often turning to software that uses data from seeds, satellites, sensors, and tractors to make better decisions about what to grow, how to grow it, and how to track food freshness from farm to fork. With the amount of farmland shrinking but the number of mouths to feed growing, farmers have quickly found they can use software to harvest data in new ways to increase yields and grow their bottom lines. For example:

- ➲ Software processing data coming from a "Lettuce Bot" mounted on a tractor can analyze real-time images of lettuce crops to improve yields by precisely feeding each plant with the nutrients it needs to maximize output. By harnessing data more effectively, farmers are able to target resources to reduce input costs, and reduce pesticide and chemical use, while improving yields by five or 10 bushels an acre.⁹
- ➲ In response to President Obama's Innovation Challenge, the US Department of Agriculture, using Microsoft Azure, used key climate data sets to explore solutions for increasing food resiliency and strengthening food systems.¹⁰



The city of Los Angeles, California, is sharing data with app providers to improve driving, reduce congestion, and promote safety, while the app providers share real-time, crowd-sourced data with the city to help city police, fire, transportation, sanitation, and street services better perform their jobs.

Software is enabling smarter community solutions

Software isn't just the tool that puts the "smart" in our smartphones. It's also helping create smarter solutions for communities to address pressing needs through smart city initiatives. Software is enabling the creation of smart cities where communities are tackling key challenges like improving the delivery of city services, reducing traffic congestion, fighting crime, cutting costs, and boosting local economies. These communities are using software that harnessed data from sensors and systems to improve the city safety, efficiency, and quality of life for their residents. In an era of constrained capital where cities still must compete to grow their economies and improve quality of life, these software-enabled systems are being deployed in a variety of ways to increase energy efficiency, improve city planning, boost transportation efficiency, transform water and wastewater management, and improve service delivery. For example:

- ➲ The New York City Fire Department is combining 7,500 data points across 17 city-agency data streams with the help of clever software that uses predictive analytics and artificial intelligence to predict which of New York City's 1 million buildings are at high risk for fires.¹¹
- ➲ Charlotte, North Carolina, is harnessing smart city software and sensors to achieve a 20 percent reduction in energy usage — saving millions of taxpayer dollars in the process.¹²

➲ The city of Los Angeles, California, is sharing data with app providers to improve driving, reduce congestion, and promote safety, while the app providers share real-time, crowd-sourced data with the city to help city police, fire, transportation, sanitation, and street services better perform their jobs.¹³

➲ The city of Chicago, Illinois, has deployed a city-wide network of 500 lamppost-mounted sensors to monitor air quality, and uses software to identify environmental issues like pest infestations that could be connected to the incidence of asthma.¹⁴

Software is helping people live longer, healthier lives through better diagnoses, lower health costs, and more personalized treatments

As populations age and health care costs rise, software is being used to improve health monitoring to better understand diseases, lower health costs, diagnose faster, revolutionize medical research, improve fitness, and help people live longer, healthier lives. For example, health care professionals are turning to clinical decision support system software to give doctors unprecedented insight into their patients' health, enable better decision making, and inform better treatment options for patients. Software systems that analyze data from disparate sources help make faster and more reliable diagnoses in a complex data environment — proving to be beneficial in more

IBM's Watson helps augment human decision making by crunching data on thousands of previous cancer cases through powerful data analytics software. Watson has helped doctors improve diagnostic accuracy rates to as high as 90 percent for lung cancer compared with a human diagnostic accuracy rate of around 50 percent.

than 70 percent of cases.¹⁵ According to McKinsey, if the health care sector were to use data more effectively to drive efficiency and quality, the sector could save more than \$300 billion every year.¹⁶ For example:

- ➲ At Memorial Sloan Kettering in New York and at MD Anderson Cancer Center in Texas, oncologists have trained IBM's Watson — a powerful cognitive system that thinks more like a human than a computer — to help doctors make better cancer treatment choices. IBM's Watson helps augment human decision making by crunching data on thousands of previous cancer cases through powerful data analytics software. Through natural language processing — the reading of billions of electronic health records and images — Watson has helped doctors improve diagnoses of various types of cancer, leading to diagnostic accuracy rates as high as 90 percent for lung cancer compared with a human diagnostic accuracy rate of around 50 percent.¹⁷
- ➲ In Africa, the 3D printing of prosthetics is drastically improving the quality of life for amputees, as 3D models can be created quickly and at a fraction of the cost of traditional prosthetics. Autodesk, a 3D entertainment and design software company, recently partnered with the University of Toronto and hospitals in Uganda to give people access to prosthetics and fit them for their daily needs.¹⁸

Software is driving transportation gains that can save time, money, fuel, and even lives

At a time of ever-increasing congestion on the road, leaders are turning to software innovations to reduce the hours we spend in congestion every year,¹⁹ connect cars in ways that save as many lives as the introduction of seat belts,²⁰ and enable truck fleets to find quicker routes, perform real-time engine diagnostics, and identify unsafe driving habits before they become a problem.²¹ For example:

- ➲ New cars today are so packed with software and sensors that they can generate up to 25 GB of data per hour²² and contain more than 10 million lines of software code just to process the data.²³ This real-time data processing is helping drivers avoid collisions, improve engine performance, save fuel, and save time.
- ➲ As more than 25 million packages are routed and delivered every day — about 300 every second²⁴ — delivery companies are turning to data-driven software to deliver the right product to the right person in the fastest, safest, and most affordable way. Package delivery companies like UPS are using vehicle sensor data combined with software that plans more efficient routes to drive decisions that save millions of gallons of fuel and reduce emissions by the equivalent of taking thousands of cars off the road for a year.²⁵



New cars today are so packed with software and sensors that they can generate up to 25 GB of data per hour and contain more than 10 million lines of software code just to process the data. This real-time data processing is helping drivers avoid collisions, improve engine performance, save fuel, and save time.

- ➲ New York City's "Vision Zero" system uses an intelligent software-based traffic control and surveillance platform to collect historical traffic incident data to measure changes in patterns that inform improvements in roads, highways, turning lanes, traffic circles, and other infrastructure. As a result, in 2015 New York City experienced the fewest traffic-related deaths in any year since 1910, despite a doubling of the city's population during the past century.²⁶

In the manufacturing sector, software is improving the way products are designed, built, and distributed. By one estimate, better use of data in manufacturing can yield up to a 50 percent decrease in product development time, cut assembly costs, and reduce working capital up to 7 percent.²⁷ In fact, in a Microsoft study, IDC estimates that manufacturing companies that take full advantage of their data are poised to achieve a \$371 billion data dividend over four years.²⁸ Software also is radically improving our ability to design and build things in smarter ways — leading to a revolution

that moves us beyond mere mass production to a world of mass customization. Software-enabled digital fabrication technologies like 3D printing can drive down the cost of complexity, enable rapid prototyping, reduce energy costs by 50 percent, and cut material costs by 90 percent.²⁹

In the financial sector, software is at the heart of our global financial system — zipping dollars around the globe, boosting efficiency, improving compliance, and cutting fraud. As our financial exchanges generate four to five terabytes of data a day, software is often the crucial tool that enables experts to make sense of it all. In one survey, 71 percent of banking and financial markets firms reported that the use of information and analytics is creating a competitive advantage for their organizations.³⁰ Similarly, for companies like Visa, software-enabled data analytics has helped the company identify \$2 billion in annual credit card fraud, and given it the chance to address those vulnerabilities before that money was lost.³¹



Software's Economic Impact Is Now Rippling Throughout the Economy

It is now becoming increasingly clear that software innovation is transforming the way we do almost everything: from the way we make movies to the way we make medicines; from how we forecast fashion to the way we farm food; from how we educate to the very way we innovate. Software is no longer just the tool that helps bring families closer even when far apart, or that enables us to work even when we are away from our desks. Today's software innovation is an engine that fuels broad sectors of our economy.

In addition to the specific software solutions often at the core of high-growth sectors of the economy,³⁷ software's pervasiveness also is driving economic change that is boosting the underlying productivity of many key sectors of our economy.

These gains stem from software's ability to make it easier and faster for workers to perform their jobs — in the process enabling companies to be more responsive to customer needs; unleashing greater human creativity and performance; bringing products and services to markets faster; lowering product and service costs; and expanding markets at global scale through the Internet. In fact, a study in the *Harvard Business Review* found that companies in the top third of their industry that use software for data-driven decision making are, on average, 5 percent more productive and 6 percent more profitable than their competitors.³⁸

So not only does the software sector help create more good, high-wage, high-skill jobs, it also benefits virtually every other industrial sector by making them more efficient. This potential is so meaningful that a recent report from GE estimates

that if companies used data to become just 1 percent more efficient, global GDP would grow by \$15 trillion by 2030.³⁹ It means the next big thing may be the billions of small connected devices made infinitely smarter through software as we connect nearly everything that can be connected to the Internet, to enable even more powerful solutions through what many call the Internet of Things.

Today's software revolution also is about helping us make smarter decisions. Rather than making big decisions based on gut or pure instinct, software innovations now are enabling us to better harness data in ways that help us make better-informed decisions, with greater precision, faster speed, and bigger impact. According to one survey of IT decision makers across a range of industries, 59 percent report that improving the quality of decision making is the primary goal driving investments in data technologies.⁴⁰

Yet despite all of the real-world benefits that software is driving today, we have seen only a fraction of its full potential. We now generate 2.5 quintillion bytes of data each day — so much that at least 90 percent of the world's data today has been created in just the last two years alone, and the amount of data produced is now doubling every two years.⁴¹ And sophisticated software code is helping us make greater sense of this growing data.



Charting a Path to an Even Brighter Economic Future

This study enables us to better understand the potential that software innovation can deliver as leaders grapple with how to chart a course for an even brighter economic future. The findings are clear: across the country, software innovation is helping drive more than a trillion dollars a year in economic impact, supporting nearly 10 million jobs, and delivering opportunity in all 50 states. And the best may still be ahead as the software industry continues to make large investments in R&D to help unlock future gains.

It is now becoming increasingly apparent that the software industry's greatest untapped potential lies not only in what it enables individuals to do today, but in its ability to fundamentally transform our economic opportunities for tomorrow.

With vast new opportunities on the horizon, how quickly we reap its benefits will be shaped by decisions being made by policymakers today. With pragmatic policies that foster a climate where software innovation can continue to thrive; that tap human talent and tenacity; harness innovation and investment; expand capacity and capabilities; and enable data to flow freely across borders, policymakers can help maximize the benefits that software innovation can deliver.



State Economic Impact

Software makes big economic contributions in all 50 states. Throughout each of the 50 states, software innovation can be seen delivering an enormous positive impact. It is directly contributing to state economies, supporting a wide range of software-related jobs, and driving outsized investments in state R&D to spur continued innovation and advance a more prosperous future.

Every state's economy and workforce are benefiting from new jobs that fit our modern digital economy and opportunities driven by software advances. The software sector in large states, including California, Texas, and New York, drives about of a third of software's trillion-dollar economic impact on GDP, while states like North Carolina and Massachusetts are benefiting from large investments into game-changing R&D in their states. Yet every state, no matter how big or small, rural or urban, is seeing tangible economic impacts from software innovation.

State	STATE ECONOMIC IMPACT				
	GDP	EMPLOYMENT		SOFTWARE R&D	
	Direct Contribution (\$million)	Direct (Jobs)	Total (Jobs)*	Investments (\$million)	As % of all R&D**
United States	\$475,301	2,536,908	9,820,443	\$52,006	17.21%
Alabama	\$3,289	25,241	52,108	\$142	11.06%
Alaska	\$248	1,325	2,839	\$4	10.26%
Arizona	\$5,190	40,834	87,244	\$235	4.74%
Arkansas	\$1,466	11,238	14,503	\$68	22.30%
California	\$90,526	408,143	949,916	\$20,522	25.12%
Colorado	\$11,633	73,371	128,931	\$694	16.91%
Connecticut	\$5,401	31,600	95,390	\$285	3.88%
Delaware	\$839	5,899	13,397	\$39	1.61%
District of Columbia	\$4,424	23,104	41,612	\$44	16.92%
Florida	\$15,812	107,522	197,220	\$976	18.73%
Georgia	\$12,596	85,773	148,748	\$858	22.14%
Hawaii	\$791	5,485	15,724	\$22	11.70%
Idaho	\$680	4,972	12,392	\$26	2.41%
Illinois	\$16,126	97,568	262,447	\$545	4.19%
Indiana	\$4,138	27,293	54,122	\$232	3.84%
Iowa	\$2,496	17,186	28,665	\$86	4.88%
Kansas	\$2,267	16,355	31,836	\$88	4.25%

*includes indirect and induced effects

**all R&D refers to all domestic business R&D in the US

State	STATE ECONOMIC IMPACT				
	GDP	EMPLOYMENT		SOFTWARE R&D	
		Direct Contribution (\$million)	Direct (Jobs)	Total (Jobs)*	Investments (\$million)
Kentucky	\$2,154	20,493	24,676	\$43	4.00%
Louisiana	\$1,234	9,606	18,042	\$26	7.16%
Maine	\$701	4,907	10,776	\$29	10.55%
Maryland	\$11,938	77,561	149,768	\$490	12.16%
Massachusetts	\$22,572	114,114	341,406	\$3,153	18.03%
Michigan	\$8,481	57,556	124,254	\$331	2.22%
Minnesota	\$8,182	48,191	73,316	\$567	9.12%
Mississippi	\$691	5,730	7,023	\$14	5.13%
Missouri	\$7,132	46,757	107,656	\$442	6.33%
Montana	\$476	4,279	7,067	\$11	10.48%
Nebraska	\$2,291	17,074	21,185	\$243	42.04%
Nevada	\$1,415	9,195	21,523	\$53	8.36%
New Hampshire	\$2,162	13,706	28,697	\$141	7.59%
New Jersey	\$16,420	86,256	232,836	\$639	4.04%
New Mexico	\$706	5,306	9,102	\$29	6.52%
New York	\$37,160	147,361	476,335	\$2,331	19.87%
North Carolina	\$9,325	70,134	127,278	\$1,433	22.91%
North Dakota	\$606	4,638	5,019	\$11	4.95%
Ohio	\$10,961	72,249	177,543	\$292	3.77%
Oklahoma	\$1,499	10,457	23,149	\$33	7.13%
Oregon	\$4,220	28,460	66,942	\$559	10.84%
Pennsylvania	\$13,301	81,603	229,603	\$992	10.65%
Rhode Island	\$1,300	9,658	19,418	\$21	4.67%
South Carolina	\$2,356	17,443	29,543	\$55	3.41%
South Dakota	\$286	2,246	2,819	\$10	8.93%
Tennessee	\$3,724	24,651	48,745	\$54	3.76%
Texas	\$29,977	199,999	339,723	\$2,214	14.60%
Utah	\$4,451	37,066	84,935	\$456	21.37%
Vermont	\$633	4,693	9,539	\$199	42.70%
Virginia	\$25,833	162,987	439,948	\$660	13.88%
Washington	\$27,564	115,133	297,699	\$9,061	62.52%
West Virginia	\$676	4,289	8,379	\$7	2.30%
Wisconsin	\$5,710	39,269	58,446	\$531	12.84%
Wyoming	\$168	945	2,159	\$3	10.00%

*includes indirect and induced effects

**all R&D refers to all domestic business R&D in the US



LETTER FROM THIRD-PARTY ECONOMISTS

The future competitiveness of the US will demand extraordinary competence in software capabilities — not just in the IT sector, but across the economy.

In a world underpinned by ever-more powerful, affordable, and public digital technology platforms, software is fast becoming the key source of economic value and competitive advantage in business. As a result, software is redefining nearly every industry.

Although no study of this nature can ever be perfect, it is vital to measure the industry as best we can and interpret findings carefully. This study from BSA | The Software Alliance, conducted by The Economist Intelligence Unit, is well designed, uses the best data available and draws reasonable, thought-provoking conclusions.

The findings of the study are important and merit attention. The software sector has a large economic impact in terms of value addition, employment, and wages. Its impact is extensive; every industry uses software. Beyond the direct contribution to GDP, software also is an enabler of productivity growth across the economy. Moreover, it offers a growing number of high-paying and professionally satisfying jobs across many sectors. And it continues to invest substantively in R&D, which increases the likelihood of its ongoing success.

The future competitiveness of the US will demand extraordinary competence in software capabilities — not just in the IT sector, but across the economy. The US already is the world leader in software. Maintaining this leadership position will require ongoing investment in R&D at the national and corporate levels, a commitment to funding related education including computer science, and a comprehensive understanding of the dynamics of the software industry and its economic impact. This study represents a welcome step in that direction.

Dr. Vijay Gurbaxani, Taco Bell Endowed Professor of Information Systems and Computer Science and Director, Center for Digital Transformation, University of California, Irvine Paul Merage School of Business

Dr. Vidyanand Choudhary, Associate Professor, Information Systems, University of California, Irvine Paul Merage School of Business

METHODOLOGY

In 2016, BSA | The Software Alliance commissioned The Economist Intelligence Unit (EIU) to assess the economic impact of the software industry. The EIU collected and analyzed the most recent data available from several recognized and reputable sources. These sources included the EIU itself, IMPLAN, the National Science Foundation, the US Bureau of Economic Analysis, the US Bureau of Labor Statistics, and the US Census Bureau.

To estimate the total contributions of the software industry to the US economy, the EIU analyzed the direct contributions and estimated indirect and induced impacts using various economic multipliers. The economic contribution analysis presented in this paper uses input-output models, which describe the full inter-industry transactions between producers and intermediate and final consumers, to compute multipliers. Multipliers allow for the estimation and isolation of the direct, indirect, and induced contributions of an industry to economic outcomes (e.g., value-added GDP, employment, and wages). Direct and indirect contributions are estimated using different multipliers:

1. **Direct contributions:** The levels of output or employment from the software industry directly.
2. **Indirect impacts:** The indirect impacts estimate the inter-industry economic activity resulting from the direct contributions (e.g., purchases of inputs). These indirect effects look backward at the linkages of the software industry in the economy, and the demands inputs from other sectors, like real estate and other professional services. This demand generates additional output (and jobs) from those sectors, that wouldn't exist if it weren't for that software industry demand. As a result, the indirect multipliers estimate this additional output from other industries that is attributable to the software industry.

3. **Induced impacts:** Induced impacts take the next step — identifying the additional economic activity supported by spending on goods and services by households whose income was affected by the direct contributions and indirect impacts. The software industry pays its employees but also supports incomes in other sectors, like real estate. These jobs come with additional wage payments, which increase total earnings to people working in these upstream sectors. These people then buy more goods and services, which generates additional demand (and output) across the broader economy. Induced multipliers estimate this additional output from increased general demand due to higher total wages paid to people in the software industry and people in industries that supply to the software industry.

The modern definition of the software industry used in the study reflects recent technological advancements in the software industry — from one that focused on tangible and packaged software products to one that includes software related services like the cloud based software as a service (SaaS), cloud storage and computing, mobile app development and hosting. As a result, the EIU analysis has defined the US software industry to include the following software sub-industries:

- NAICS 5112: Software Publishers
- NAICS 5415: Computer Systems Design and Related Services
- NAICS 518: Data Processing, Hosting and Related Services
- NAICS 519130: Internet Publishing and Broadcasting and Web Services

The EIU compiled these data and economic impact assessments using publicly available government data, maintaining full editorial control of the process and using industry standard approaches. Any views or opinions expressed in this document are not necessarily those of The Economist Intelligence Unit.



NOTES

- ¹ All data is from 2014 unless otherwise indicated.
- ² Value includes indirect and induced impacts. For definitions of "indirect" and "induced," see Methodology section on page 16.
- ³ US Department of Labor, Bureau of Labor Statistics, Occupational Employment Statistics. Data from May 2015.
- ⁴ Patrick Gallagher, "Innovation as a Key Driver of Economic Growth & Competitiveness" (Remarks at FedScoop's US Innovation Summit, Washington, DC, June 20, 2012), available at <http://www.nist.gov/director/speeches/innovation-summit.cfm>.
- ⁵ National Science Foundation/National Center for Science and Engineering Statistics and US Census Bureau, Business R&D and Innovation Survey 2012 industry breakdown. Where data is not available for 2012, the most recent year is used.
- ⁶ National Science Foundation/National Center for Science and Engineering Statistics.
- ⁷ Significant hearing loss meant Sarah Chumeran had always worn hearing aids to help her navigate her days. But then Sarah received an Esteem hearing implant. This tiny technology made it possible for her to hear her daughter's voices for the first time. It also gave her the self-confidence to be more independent. David Chen, "Microsoft Technology Helped Sarah Chumeran Hear for the First Time," The Official Microsoft Blog, February 1, 2014, available at <http://blogs.microsoft.com/forense/2014/02/01/microsoft-technology-helped-sarah-chumeran-hear-for-the-first-time/#sm.001wmyzd910kpc4pb5>.
- ⁸ "Capgemini Report Shows Rising Impact of Big Data on Decision-Making," Press Release, June 12, 2012, available at <https://www.capgemini.com/news/capgemini-report-shows-rising-impact-of-big-data-on-decision-making>.
- ⁹ Katherine Noyes, "Groping Up on Every Farm: Big Data Technology," Fortune.com, May 30, 2014, available at <https://fortune.com/2014/05/30/groping-up-on-every-farm-big-data-technology/>.
- ¹⁰ "Announcing the USDA-Microsoft Innovation Challenge Awardees," Microsoft Research Blog, January 27, 2016, available at https://blogs.msdn.microsoft.com/msr_en/2016/01/27/announcing-the-usda-microsoft-innovation-challenge-awardees/.
- ¹¹ "Report to the President Technology and the Future of Cities," Executive Office of the President, President's Council on Advisors on Science and Technology, February 2016, available at https://www.whitehouse.gov/sites/default/files/microsites/ostp/PCAST/pcast_cities_report_final_3_2016.pdf.
- ¹² "Helping the City of Charlotte Envision a More Sustainable Future," Smart Cities Council, available at <http://smartcitiescouncil.com/resources/helping-city-charlotte-envision-more-sustainable-future>.
- ¹³ "Report to the President Technology and the Future of Cities."
- ¹⁴ "Report to the President Technology and the Future of Cities."
- ¹⁵ Today's experienced clinician needs close to two million pieces of information to practice medicine and doctors subscribe to an average of seven journals, representing more than 2,500 new articles each year, making it almost impossible to keep abreast with the latest information about diagnosis, prognosis, therapy and related health issues. "Clinical Decision Support Systems: The Time Has Come," Frost & Sullivan, available at <http://www.frost.com/press-releases/129788>.
- ¹⁶ James Mankoski et al., "Big Data: The Next Frontier for Innovation, Competition, and Productivity," McKinsey Global Institute, May 2011, available at http://www.mckinsey.com/insights/business_technology/big_data_the_next_frontier_for_innovation.
- ¹⁷ "IBM's Watson Gets Its First Piece of Business in Healthcare," Forbes.com, February 2, 2013, available at <http://www.forbes.com/sites/bruceleupin/2013/02/08/ibms-watson-gets-its-first-piece-of-business-in-healthcare/#18a2d2f944b1>.
- ¹⁸ "3D Printing for Developing Countries: The Untapped Potential," Autodesk Education Community, October 1, 2015, available at <http://sustainabilityworkshop.autodesk.com/blog/3d-printing-developing-countries-untapped-potential>.
- ¹⁹ Traffic congestion already costs drivers more than \$100 billion annually in wasted fuel and lost time, costing drivers more than \$1,700 and 100 hours in lost time a year. "Economic and Environmental Impact of Traffic Congestion in Europe and the U.S.," INRIX, available at <http://www.inrix.com/economic-environment-cost-congestion/>.
- ²⁰ Traffic accidents kill an estimated 1.24 million people a year globally, 93 percent of accidents are based on human error. New connected vehicle sensors and automated decision making (even without becoming fully autonomous) could have as big of a safety effect as seatbelts — reducing injuries and fatalities by as much as 50 percent. Claire Cain Miller, "If Robots Drove, How Much Safer Would Roads Be?" New York Times, June 10, 2014, available at <http://www.nytimes.com/2014/06/10/upshot/if-robots-drove-how-much-safer-would-roads-be.html?ref=technology&r=0>.
- ²¹ Connected trucks have already enabled companies to save millions of gallons of gas and reduce emissions by the equivalent of taking thousands of cars off the road for a year. Alex Mayasy, "Why UPS Trucks Don't Turn Left," Priceonomics, April 4, 2014, available at <http://priceonomics.com/why-ups-trucks-don-t-turn-left/>.
- ²² Ford's "Sync 3" car infotainment model generates up to 25 GB of data per hour. Nicole Hemsoth, "How Ford Is Putting Hadoop Pedal to the Metal," Datamami, March 16, 2013, available at http://www.datamami.com/datamami/2013/03/16/how_ford_is_putting_hadoop_pedal_to_the_metal.html.
- ²³ The Chevy Volt contains over 10 million lines of software code, and software developer is one of the fastest growing technical professions in Southeast Michigan, a region long known for its manufacturing prowess. Jaclyn Trop, "Detroit, Embracing New Auto Technologies, Seeks App Builders," New York Times, June 30, 2013, available at <http://www.nytimes.com/2013/07/01/technology/detroit-embracing-new-auto-technologies-seeks-app-builders.html?r=0>.
- ²⁴ UPS daily volume is 16.9 million packages and documents per day. "UPS Fact Sheet," available at <https://www.pressroom.ups.com/pressroom/ContentDetailsViewer.page?ConceptType=FactSheets&id=1426321563187-193>. FedEx volume is 10.5 million per day. "About FedEx Corporate Fact Sheet," available at http://about.van.fedex.com/fedex_corporation.
- ²⁵ Alex Mayasy, "Why UPS Trucks Don't Turn Left."
- ²⁶ Allison Shapiro, "Can We Use Data to Stop Deadly Car Crashes? Pacific Standard, January 26, 2016, available at <https://pmag.com/can-we-use-data-to-stop-deadly-car-crashes-c186a45a00s#:cd5p1d1jcbl>.
- ²⁷ James Manyika et al., "Big Data: The Next Frontier for Innovation, Competition, and Productivity."
- ²⁸ Suzanne Choney, "Data Smart: Strategies for Customers Are Yielding Early but Impressive Returns," The Official Microsoft Blog, May 22, 2014, available at <http://blogs.microsoft.com/forense/2014/05/22/data-smart-strategies-for-customers-are-yielding-early-but-impressive-returns/>.
- ²⁹ "Additive Manufacturing: Pursuing the Promise," US Department of Energy, August 2012, available at https://www1.eere.energy.gov/manufacturing/pdfs/additive_manufacturing.pdf.
- ³⁰ Steve LaValle et al., "Analytics: The New Path to Value," IBM Institute for Business Value, 2010, available at http://www-01.ibm.com/mmon/ibm/carbon/reports/whitepaper_type_id=PMSE&category_id=GB_TL_USD18&file_id=GB0303371USE&att_id=GB032745EN_PDF.
- ³¹ Steve Rosenbaum, "Visa Says Big Data Identifies Billions of Dollars in Fraud," Wall Street Journal, March 11, 2013, available at <http://blogs.wsj.com/cio/2013/03/11/visa-says-big-data-identifies-billions-of-dollars-in-fraud/>.
- ³² Marc Andreessen, "Why Software Is Eating the World," Wall Street Journal, August 20, 2011, available at <http://www.wsj.com/articles/SB1000142405311903480945761225091529460>.
- ³³ Andrew McAfee and Erik Brynjolfsson, "Big Data: The Management Revolution," Harvard Business Review, October 2012, available at <https://hbr.org/2012/01/big-data-the-management-revolution/ar>.
- ³⁴ A 1 percent productivity increase may seem small, but as Jeff Immelt, CEO of GE, puts it, "tell an oil guy you can use software to save him one percent on something, and that guy will be your friend for life." Peter C. Evans and Marco Annunziata, "Industrial Internet: Pushing the Boundaries of Minds and Machines," November 26, 2012, available at <http://files.gerperspectives.com/wp-content/uploads/2012/11/ge-industrial-internet-vision-paper.pdf>.
- ³⁵ Lynn Holmlund, "CEOs Call for Big Data and IT Continues to Lead Investment Decisions," Press Release, January 7, 2014, available at <http://www.idgenterprise.com/news/press-release/ceos-call-for-big-data-and-it-continues-to-lead-investment-decisions>.
- ³⁶ "Bringing Big Data to the Enterprise," IBM, <http://www-01.ibm.com/software/data/bigdata/what-is-big-data.html>.

BSA | The Software Alliance (www.bsa.org) is the leading advocate for the global software industry before governments and in the international marketplace. Its members are among the world's most innovative companies, creating software solutions that spark the economy and improve modern life.

With headquarters in Washington, DC, and operations in more than 60 countries, BSA pioneers compliance programs that promote legal software use and advocates for public policies that foster technology innovation and drive growth in the digital economy.

The Software Alliance

BSA

www.bsa.org

BSA Worldwide Headquarters
20 F Street, NW
Suite 800
Washington, DC 20001
T: +1.202.872.5500
F: +1.202.872.5501

BSA Asia-Pacific
300 Beach Road
#25-08 The Concourse
Singapore 199555
T: +65.6292.2072
F: +65.6292.6369

BSA Europe, Middle East & Africa
2 Queen Anne's Gate Buildings
Dartmouth Street
London, SW1H 9BP
United Kingdom
T: +44.207.340.6080
F: +44.207.340.6090

Senator WICKER. I would like to enter into the record a letter from the Information Technology Industry Council and a letter from the Electronic Privacy Information Center without objection. And that's so ordered.

[The information referred to follows:]

INFORMATION TECHNOLOGY INDUSTRY COUNCIL
Washington, DC, December 11, 2017

Hon. ROGER WICKER, Chairman,
Hon. BRIAN SCHATZ, Ranking Member,
Subcommittee on Communications, Technology, Innovation, and the Internet,
U.S. Senate Committee on Commerce, Science, and Transportation,
Washington, DC.

Dear Chairman Wicker and Ranking Member Schatz:

In advance of your hearing on “Digital Decision-Making: The Building Blocks of Machine Learning and Artificial Intelligence,” I am writing to thank you for interest in and attention to the exciting innovation that is Artificial Intelligence (AI). ITI represents more than 60 of the world’s leading information and communications technology (ICT) companies. Our companies are the most dynamic and innovative companies from all corners of the ICT sector, including the development and deployment of AI. I submit this letter on behalf of ITI and its members, and respectfully request that you enter it into the hearing record.

Artificial intelligence (AI) technology is an integral part of our daily lives, work, and existence. It’s already made an important mark on much of our society and economy, and the exciting part is that we’re just seeing the beginning of its benefits.

Go to any hospital and medical research center and you will see how doctors and medical providers use AI to save lives. For example, the company Berg uses AI to analyze large amounts of oncological data to create a model of how pancreatic cancer functions, enabling us to develop chemotherapy to which cancer cells are more responsive.

Educators across the country use AI to enhance the potential for future generations to grow and learn. Thanks to IBM’s Teacher Advisor, a new tool based on its Watson cognitive computing platform, third-grade math teachers can develop personalized lesson plans. Teacher Advisor analyzes education standards, sets targets for skills development, and uses student data to help teachers tailor instructional material for students with varying skill levels.

AI also makes day-to-day of life easier—for everyone. Many of our everyday tasks, like making shopping lists and ordering groceries, are streamlined through devices like Alexa. And, through AI technology, researchers at the International Islamic University of Malaysia have developed the Automatic Sign Language Translator (ASLT) that uses machine learning to interpret sign language and convert it into text, easing communications for many.

We know AI will revolutionize the way we do business and our overall economy. It’s projected AI will add between \$7.1 trillion and \$13.17 trillion to the global economy by 2025.¹

There’s no question AI will continue to transform our lives, society, and economy for the better. We understand, however, that there are many questions about this technology and that with transformative innovation, there are going to be points of tension. The tech industry is committed to working with all stakeholders to identify and resolve challenges.

In conjunction with the global leaders in AI innovation, ITI recently published *AI Policy Principles*. These principles are designed to be guidelines for the responsible development and deployment of AI as we develop partnerships with governments, academia, and the public.

Our Policy Principles are a conversation catalyst, encouraging all stakeholders, public and private, to collaborate to create smart policies that allow this emerging technology to flourish while addressing the complex issues that arise out of its growth and deployment. Given the reach of AI, we think this kind of partnership and engagement is critical to advance the benefits and responsible growth of AI while also endeavoring to answer the public’s questions about the use of this nascent technology.

¹<https://www.itic.org/resources/AI-Policy-Principles-FullReport2.pdf>

We look forward to working with this Committee, other members of Congress, academia, industry partners, and the public to advance AI responsibly. Thank you, again, for holding the timely and important hearing.

DEAN GARFIELD,
President and CEO,
Information Technology Industry Council (ITI)

ELECTRONIC PRIVACY INFORMATION CENTER
Washington, DC, December 12, 2017

Hon. JOHN THUNE, Chairman,
Hon. BILL NELSON, Ranking Member,
U.S. Senate Committee on Commerce, Science, and Transportation,
Washington, DC.

Dear Chairman Thune and Ranking Member Nelson:

We write to you regarding the “Digital Decision-Making: The Building Blocks of Machine Learning and Artificial Intelligence” hearing.¹ EPIC is a public interest research center established in 1994 to focus public attention on emerging privacy and civil liberties issues.² EPIC has promoted “Algorithmic Transparency” for many years.³

Democratic governance is built on principles of procedural fairness and transparency. And accountability is key to decision making. We must know the basis of decisions, whether right or wrong. But as decisions are automated, and organizations increasingly delegate decisionmaking to techniques they do not fully understand, processes become more opaque and less accountable. It is therefore imperative that algorithmic process be open, provable, and accountable. Arguments that algorithmic transparency is impossible or “too complex” are not reassuring.

It is becoming increasingly clear that Congress must regulate AI to ensure accountability and transparency:

- Algorithms are often used to make adverse decisions about people. Algorithms deny people educational opportunities, employment, housing, insurance, and credit.⁴ Many of these decisions are entirely opaque, leaving individuals to wonder whether the decisions were accurate, fair, or even about them.
- Secret algorithms are deployed in the criminal justice system to assess forensic evidence, determine sentences, to even decide guilt or innocence.⁵ Several states use proprietary commercial systems, not subject to open government laws, to determine guilt or innocence. The Model Penal Code recommends the implementation of recidivism-based actuarial instruments in sentencing guidelines.⁶ But these systems, which defendants have no way to challenge are racially biased, unaccountable, and unreliable for forecasting violent crime.⁷
- Algorithms are used for social control. China’s Communist Party is deploying a “social credit” system that assigns to each person government-determined favorability rating. “Infractions such as fare cheating, jaywalking, and violating family-planning rules” would affect a person’s rating.⁸ Low ratings are also assigned to those who frequent disfavored websites or socialize with others who have low ratings. Citizens with low ratings will have trouble getting loans or

¹Digital Decision-Making: The Building Blocks of Machine Learning and Artificial Intelligence, 115th Cong. (2017), S. Comm. on Commerce, Science, and Transportation, <https://www.commerce.senate.gov/public/index.cfm/hearings?ID=7097E2B0-4A6B-4D92-85C3-D48E1008C8FD> (Dec. 12, 2017).

²EPIC, About EPIC, <https://epic.org/epic/about.html>.

³EPIC, Algorithmic Transparency, <https://epic.org/algorithmic-transparency/>.

⁴Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 Wash. L. Rev. 1 (2014).

⁵EPIC v. DOJ (Criminal Justice Algorithms), EPIC, <https://epic.org/foia/doj/criminal-justice-algorithms/>; Algorithms in the Criminal Justice System, EPIC, <https://epic.org/algorithmic-transparency/crim-justice/>.

⁶Model Penal Code: Sentencing § 6B.09 (Am. Law. Inst., Tentative Draft No. 2, 2011).

⁷Julia Angwin et al., *Machine Bias*, ProPublica (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.

⁸Josh Chin & Gillian Wong, *China’s New Tool for Social Control: A Credit Rating for Everything*, Wall Street J., Nov. 28, 2016, <http://www.wsj.com/articles/chinas-new-tool-for-social-control-a-credit-rating-for-everything-1480351590>

government services. Citizens with high rating, assigned by the government, receive preferential treatment across a wide range of programs and activities.

- In the United States, U.S. Customs and Border Protection has used secret analytic tools to assign “risk assessments” to U.S. travelers.⁹ These risk assessments, assigned by the U.S. Government to U.S. citizens, raise fundamental questions about government accountability, due process, and fairness. They may also be taking us closer to the Chinese system of social control through AI.

In a recent consumer complaint to the Federal Trade Commission, EPIC challenged the secret scoring of young athletes.¹⁰ As EPIC’s complaint regarding the Universal Tennis Rating system makes clear, the “UTR score defines the status of young athletes in all tennis related activity; impacts opportunities for scholarship, education and employment; and may in the future provide the basis for ‘social scoring’ and government rating of citizens.”¹¹ As we explained to the FTC, “EPIC seeks to ensure that all rating systems concerning individuals are open, transparent and accountable.”¹²

In re Universal Tennis, EPIC urged the FTC to (1) Initiate an investigation of the collection, use, and disclosure of children’s personal information by Universal Tennis; (2) Halt Universal Tennis’s scoring of children without parental consent; (3) Require that Universal Tennis make public the algorithm and other techniques that produce the UTR; (4) Require that Universal Tennis establish formal procedures for rectification of inaccurate, incomplete, and outdated scoring procedures; and (5) Provide such other relief as the Commission finds necessary and appropriate.¹³

“Algorithmic Transparency” must be a fundamental principle for consumer protection. The phrase has both literal and figurative dimensions. In the literal sense, it is often necessary to determine the precise factors that contribute to a decision. If, for example, a government agency or private company considers a factor such as race, gender, or religion to produce an adverse decision, then the decision-making process should be subject to scrutiny and the relevant factors identified.

On October 12, 2016, The White House announced two reports on the impact of Artificial Intelligence on the U.S. economy and related policy concerns. *Preparing for the Future of Artificial Intelligence* concluded that “practitioners must ensure that AI-enabled systems are governable; that they are open, transparent, and understandable; that they can work effectively with people; and that their operation will remain consistent with human values and aspirations.”¹⁴

Some have argued that algorithmic transparency is simply impossible, given the complexity and fluidity of modern processes. But if that is true, there must be some way to recapture the purpose of transparency without simply relying on testing inputs and outputs. We have seen recently that it is almost trivial to design programs that evade testing.¹⁵ And central to the science and innovation is the provability of results.

Europeans have long had a right to access “the logic of the processing” concerning their personal information.¹⁶ That principle is reflected in the U.S. in the publication of the FICO score, which for many years remained a black box for consumers, establishing credit worthiness without providing any information about the basis of score.¹⁷

The continued deployment of AI-based systems raises profound issues for democratic countries. As Professor Frank Pasquale has said:

⁹EPIC v. CBP (Analytical Framework for Intelligence), EPIC, <https://epic.org/foia/dhs/cbp/afii/>.

¹⁰EPIC, EPIC Asks FTC to Stop System for Secret Scoring of Young Athletes (May 17, 2017), <https://epic.org/2017/05/epic-asks-ftc-to-stop-system-f.html>; See also Shanya Possess, Privacy Group Challenges Secret Tennis Scoring System, Law360, May 17, 2017, <https://www.law360.com/articles/925379>; Lexology, EPIC Takes a Swing at Youth Tennis Ratings, June 1, 2017, <https://www.lexology.com/library/detail.aspx?g=604e3321-dfc8-4f46-9afc-abd47c5a5179>

¹¹EPIC Complaint to Federal Trade Commission, *In re Universal Tennis* at 1 (May 17, 2017).

¹²*Id.*

¹³*Id.* at 13.

¹⁴Preparing for the Future of Artificial Intelligence, (Oct 2016), Executive Office of the President, National Science and Technology Council, Comm. on Technology, https://obama-whitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.

¹⁵Jack Ewing, *In '06 Slide Show, a Lesson in How VW Could Cheat*, N.Y. Times, Apr. 27, 2016, at A1.

¹⁶Directive 95/46/EC—The Data Protection Directive, art 15 (1), 1995, <http://www.dataprotection.ie/docs/EU-Directive-95-46-EC—Chapter-2/93.htm>.

¹⁷Hadley Malcom, *Banks Compete on Free Credit Score Offers*, USA Today, Jan. 25, 2015, <http://www.usatoday.com/story/money/2015/01/25/banks-free-credit-scores/22011803/>.

Black box services are often wondrous to behold, but our black box society has become dangerously unstable, unfair, and unproductive. Neither New York quants nor California engineers can deliver a sound economy or a secure society. Those are the tasks of a citizenry, which can perform its job only as well as it understands the stakes.¹⁸

We ask that this Statement from EPIC be entered in the hearing record. We look forward to working with you on these issues of vital importance to the American public.

Sincerely,

MARC ROTENBERG,
President,
EPIC.

CAITRIONA FITZGERALD,
Policy Director,
EPIC.

CHRISTINE BANNAN,
Policy Fellow.
EPIC.

Senator WICKER. The hearing record will remain open for 2 weeks. During this time, Senators are asked to submit any questions for the record. Upon receipt, the witnesses are requested to submit their written answers to the Committee as soon as possible.

Thank you very much. The hearing is now adjourned.
[Whereupon, at 12 p.m., the hearing was adjourned.]

¹⁸Frank Pasquale, *The Black Box Society: The Secret Algorithms that Control Money and Information* 218 (Harvard University Press 2015).

A P P E N D I X

RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. AMY KLOBUCHAR TO DR. CINDY L. BETHEL

Question. Political ads on the Internet are more popular now than ever. In 2016, more than \$1.4 billion was spent on digital advertisements and experts project that number will continue to increase. In October, I introduced the Honest Ads Act with Senators Warner and McCain, to help prevent foreign interference in future elections and improve the transparency of online political advertisements. We know that 90 percent of the ads that Russia purchased were issue ads meant to mislead and divide Americans. Increasing transparency and accountability online will benefit consumers and help safeguard future elections.

Dr. Bethel, could making more data about political advertisements publicly available help improve the performance of algorithms designed to prevent foreign interference?

Answer. More data that is high quality and has varied examples of content may be helpful in improving AI algorithms in general. The data though needs to have features that are detectable to learn from as part of the machine learning process. It is not clear from the political advertisements that have been promoted, as in the previous election, it was not evident what features could be used to learn from to detect involvement by foreign parties of these political advertisements. More data does not always equate to better results. There needs to be sufficient variations in key features that can be detected to be able to develop and test algorithms that will be effective and will have beneficial and meaningful results.

RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. TOM UDALL TO DR. CINDY L. BETHEL

Question. Can you speak of some of the ways that government funded artificial intelligence development is now being used in the private sector?

Answer. AI research has been funded for many years through the National Science Foundation, National Institutes of Health, the Department of Defense, USDA, among other agencies. It is possible that there has been government-funded research into artificial intelligence that has moved directly into the private sector for inclusion in product development. Generally, concepts, algorithms, and advancements developed as part of artificial intelligence research, has been integrated into product developments in the private sector but it is typically not a direct transfer from research into the private sector.

Currently, the research I am performing with Therabot™, the robotic therapeutic support companion, funded by the National Science Foundation, developed algorithms associated with artificial intelligence and machine learning that are being used in this application. This is a project that is planned for commercialization and making it available to the public. Mississippi State University has received government funding for projects that include artificial intelligence and have leveraged and transitioned some of the that knowledge into industry-based projects that benefit the private sector.

The private sector has been active in funding their own advancements of artificial intelligence and machine learning. Many of the developments and advancements in applications of AI have been developed from private sector research and development groups. There are also cases, where a researcher receives government-funded grants for the development of AI and machine learning, and then may later transition into a private sector position and takes that knowledge with him or her to advance product development that benefits the private sector.

There are numerous research developments that have been government funded under SBIR/STTR mechanisms that are joint funding for industry and academic researchers. NIH and USDA have been very active in transitioning the technology developments using their funds into the private sector and commercially available

products. NSF and DoD also has programs such as Innovation Corps (I-Corps) that transition developments into commercialized and publicly available products. These have been successful programs available to researchers who have been funded under government grants and contracts.

I am not sure of all of the government-funded research to date in artificial intelligence, so I am not sure exactly which projects have benefited or ended up being applied in the private sector. That would be a research project in itself that may be a worthwhile effort.

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. MAGGIE HASSAN TO
DR. CINDY L. BETHEL

Question 1. Artificial intelligence, or AI, holds tremendous promise for individuals who experience disabilities. For example, Google and Microsoft have technologies to process language and speech and translate it into a text format to assist individuals who are deaf and hard of hearing. Other technologies will go even further to improve the lives of people with disabilities and I would like to learn more from the panel about what we can expect. What other specific technologies are you aware of in the AI space that will help people who experience disabilities?

Answer. There are numerous technologies that use artificial intelligence that are being developed to assist people with different types of disabilities to improve their quality of life. There is the development and use of brain-machine interfaces for operating wheelchairs and other devices through detecting signals sent with intention from the brain. There are exoskeletons that are being developed and prosthetics that learn and detect the signals and impulses from the nervous system to be able to customize how these devices work and enhance the capabilities of the end users. There are in-home assistive robots that are being developed to assist disabled and elderly people in their homes with reminders to take medications, to fetch different items, and to remotely monitor users so they can remain in their homes longer. The Therabot™ robot is being developed as a home therapy support tool to provide comfort and support for people who have experienced post-traumatic stress or other types of mental health disorders that can be used to detect and alert clinicians or others when problems are detected. There are continual enhancements to technologies for the hearing and visually impaired users. These are just some of the many examples that use artificial intelligence to improve quality of life for people with different types of disabilities.

Question 2. How will manufacturers and developers work to perfect this technology so that it can truly be a reliable tool for these individuals?

Answer. These technologies will need to go through extensive testing and user studies/clinical trials to ensure the safety of the developments before they are sold to the public or used by the public. Edge cases need to be tested for events that may not occur frequently but have the possibility of happening. Once these types of technologies are developed and tested, then standards need to be established to ensure ongoing quality of the products for safe use as would be the case of any product used in medical applications or for consumer use.

Question 3. What more can Congress do to assist with these efforts?

Answer. While research and development is occurring it is important to not establish highly restrictive legislative policies or it will stifle the creativity and development by researchers. Once something is established and has been tested, then it may be necessary to legislate standards of practice for the protection and safety of the public using these items. This would be later in the process. Providing funding to support and an environment supportive of development of these items would allow the U.S. to stay on the top of research developments that use artificial intelligence and machine learning. Legislation should be limited to restrictions and standards for consumer and user safety.

Question 4. As we see machine learning and AI increasingly embedded in products and services that we rely on, there are numerous cases of these algorithms falling short of consumer expectations. For example, Google and Facebook both promoted fraudulent news stories in the immediate wake of the Las Vegas Shooting because of their algorithms.¹

YouTube Kids is a service designed for children, and marketed as containing videos that are suitable for very young children. In November, YouTube Kids promoted

¹ NYT: After Las Vegas Shooting, Fake News Regains Its Megaphone, Kevin Rose, 10/02/2017
<https://www.nytimes.com/2017/10/02/business/las-vegas-shooting-fake-news.html>

inappropriate content due to algorithms.² While the use of machine learning and AI holds limitless positive potential, at the current point, it faces challenges where we should not risk getting it wrong.

Should there be any formal or informal guidelines in place for what tasks are suitable to be done by algorithms, and which are still too important or sensitive to turn over; and what more can be done to ensure better and more accurate algorithms are used as you work to better develop this technology?

Answer. In cases, where the algorithms are related to safety or life critical decisions then it may be necessary to have a human in the loop for sanity checks to ensure the best possible decision is made. When it comes to children, this would be case when the system needs to be thoroughly tested with a human involved to ensure the system is working well and there needs to be testing and validation that occurs to ensure those “edge” cases or rarer situations are also tested for no matter how unlikely it is for it to occur. Validation and testing should be performed extensively with adults prior to using the system with children and then tested with children with adult supervision. Mistakes can happen, but everything possible needs to be done to prevent issues that could potentially cause harm. In the case of decisions to weaponize autonomous systems there should be a human in the loop when it comes to decisions that impact human lives. There does need to be established standards and benchmarks to assist developers in testing to ensure the safety of a product before it is put in the hands of the public.

Question 5. Machine learning and AI hold great promise for assisting us in preventing cybersecurity attacks. According to an IBM survey of Federal IT managers, 90 percent believe that artificial intelligence could help the Federal Government defend against real-world cyber-attacks. 87 percent think AI will improve the efficiency of their cybersecurity workforce.³

While this is promising, the Federal Government currently faces a shortage of qualified cybersecurity employees, and to make matters worse, the pipeline of students studying these topics is not sufficient to meet our needs. A recent GAO report found that Federal agencies have trouble identifying skills gaps, recruiting and retaining qualified staff, and lose out on candidate due to Federal hiring processes.

The George Washington University Center for Cyber & Homeland Security recently released a report titled “Trends in Technology and Digital Security” which stated:

“Traditional security operations centers are mostly staffed with tier one analysts staring at screens, looking for unusual events or detections of malicious activity. This activity is similar to physical security personnel monitoring video cameras for intruders. It is tedious for humans, but it is a problem really well-suited to machine learning.”⁴

What effect will machine learning and AI will have on cybersecurity; and how do you think the Federal Government can best leverage the benefits offered by machine learning and AI to address our cybersecurity workforce shortage?

Answer. The use of artificial intelligence and machine learning can definitely help with the tedious task of identifying potential threats to security. The initial evaluation could be performed using computer systems and they are adept at detecting anomalies and maybe even ones that a human may not readily detect. In cases where it is a threat to life or safety of humans involved, then it may need to be verified by a cybersecurity trained specialist. Overall, the use of good algorithms and machine learning systems could help fill the gap that has occurred with the lack of trained cybersecurity personnel. If the initial detection work can be performed by the computer systems, then it would require less personnel to verify those findings. Overall, the use of well-developed AI and machine learning systems could be leveraged to address some of the workforce shortage issues associated with cybersecurity professionals. It is also important to better recruit for these programs and to show the benefits of being involved in this type of career. There are limitations to the government hiring practices and pay scales, but these can be overcome. It may require changes though to these practices to entice students entering and choosing their career fields to consider careers in these areas.

²NYT: *On YouTube Kids, Startling Videos Slip Past Filters*, Sapna Maheshwari, 11/04/2017 <https://www.nytimes.com/2017/11/04/business/media/youtube-kids-paw-patrol.html>

³INFORMATION MANAGEMENT: *AI seen as key tool in government's cybersecurity defense*, Bob Violino, 11/30/2017 <https://www.information-management.com/news/artificial-intelligence-seen-as-key-tool-in-governments-cybersecurity-defense>

⁴<https://cchs.gwu.edu/sites/cchs.gwu.edu/files/downloads/Fall%202017%20DT%20symposium20compendium.pdf>

**RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. TOM UDALL TO
DANIEL CASTRO**

Question. In your testimony, you discussed how regulators will not need to intervene because the private sector will address artificial intelligence's problems—such as bias and discrimination. However, there have been studies that show implicit bias even when artificial intelligence is deployed. For example, a study¹ about using AI to evaluate resumes found that candidates with names associated with being European American were 50 percent more likely to be offered an interview than candidates with names associated with being African-American. What role should the Federal Government play where there is implicit bias and discrimination—particularly when companies are required to be "Equal Opportunity Employers"?

Answer. This is an important question. To clarify, regulators will need to continue to intervene to address specific policy goals, such as ensuring non-discrimination in hiring practices. However, policymakers do not necessarily need to create new laws and regulations only for AI to achieve those goals. Existing laws that make these practices illegal still apply, regardless of whether or not a company uses AI to discriminate against a protected class. For example, a company cannot circumvent its obligations in Title VII of the Civil Rights Act to discriminate against a particular race in its hiring practices simply by using an algorithm to review job applicants.

There are additional steps policymakers can take to reduce bias. One way to assess bias, whether it be in analog processes or digital algorithms, is to have businesses conduct disparate impact analyses. For example, if a company is using an AI system to screen job applicants, and it has concerns about potential racial bias, it should test this system to assess its accuracy. If government agencies are early adopters of such AI-driven services, they can help identify potential areas of concern. However, disparate impact analysis is only possible if organizations have data available to them. Moreover, regulators can also identify practices that are known to have a disparate impact, such as using certain criteria for making a credit or housing decision, and discouraging businesses from using those methods.²

In addition, there are likely areas where additional policy is needed to protect workers. For example, Congress should consider passing laws such as the Employment Non-Discrimination Act (ENDA) to ensure that data about sexual orientation and gender identity cannot be used to unfairly harm workers.³ These types of laws address specific concerns of vulnerable populations, but they do not apply only to AI.

Finally, policymakers should recognize that using AI can often reduce discrimination by limiting the potential for both implicit and explicit human bias. For example, a company that uses AI to screen applicants has the potential to reduce implicit bias of human managers in hiring practices. And while AI systems may not be perfect at the outset, as people identify problems, they will be able to more quickly resolve these issues. The same is not true for strictly human processes, where eliminating bias, is much more difficult.

**RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. GARY PETERS TO
DANIEL CASTRO**

Question. Many have predicted that AI will have a profound effect on the labor market. Most predict that low-wage, routine-based jobs will be under the most pressure for replacement by AI. Meanwhile, recent advancements in technology has led to job creation that will mostly require highly-skilled, highly-educated workers. What evidence have you seen regarding businesses incorporating this labor shift into their business plans?

Answer. Some of these predictions are not based on sound analysis. Bureau of Labor Statistics (BLS) projections show that the fastest growing jobs are not in high-skilled occupations. For example, the industry that BLS projects will have the most job growth between 2016–2026 is the "food services and drinking places" in-

¹ <http://www.sciencemag.org/news/2017/04/even-artificial-intelligence-can-acquire-biases-against-race-and-gender>

² Travis Korte and Daniel Castro, "Disparate Impact Analysis is Key to Ensuring Fairness in the Age of the Algorithm," Center for Data Innovation (2015), <http://datainnovation.org/2015/01/disparate-impact-analysis-is-key-to-ensuring-fairness-in-the-age-of-the-algorithm/>.

³ For more on this topic as it relates to data, see: Joshua New and Daniel Castro, "Accelerating Data Innovation: A Legislative Agenda for Congress," Center for Data Innovation (2015), <http://datainnovation.org/2015/05/accelerating-data-innovation-a-legislative-agenda-for-congress/>.

dstry.⁴ These are not high-wage jobs. Increased use of AI can yield higher rates of automation and hopefully fewer of these low-wages jobs. The way to achieve higher-wage jobs is by increasing productivity. In particular, increasing productivity in low-skill jobs will grow wages in these occupations.

Some companies have taken steps to address disruption in the workforce. For example, Google and Facebook have made substantial commitments to funding for job retraining programs.⁵ However, overall, U.S. companies are investing less in training now than they were 15 years ago. A knowledge tax credit, where corporates receive a credit for qualified expenditures on worker training, would help address this problem.⁶

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. MAGGIE HASSAN TO
DANIEL CASTRO

Question 1. Artificial intelligence, or AI, holds tremendous promise for individuals who experience disabilities. For example, Google and Microsoft have technologies to process language and speech and translate it into a text format to assist individuals who are deaf and hard of hearing.

Other technologies will go even further to improve the lives of people with disabilities and I would like to learn more from the panel about what we can expect.

What other specific technologies are you aware of in the AI space that will help people who experience disabilities?

Answer. AI will have widespread benefits for people with disabilities. As noted in the question, one of the major areas of impact AI will have is by allowing more people to interface with computer systems using their voice, instead of a keyboard. In particular, the combination of AI with the Internet of Things, will give people with many types of disabilities a better quality of life as they will now be able to control more of the world around them. These functions will allow more people with disabilities to participate in the workforce, go to school, and be more active in their communities. In addition, AI can be used to create smart agents that automate specific tasks, such as scheduling meetings, setting a thermostat, or re-ordering groceries. While these types of actions are conveniences for some people, for people with significant disabilities, they can be empowering and allow individuals significantly more autonomy and independence.

Question 2. How will manufacturers and developers work to perfect this technology so that it can truly be a reliable tool for these individuals?

Answer. One way to improve is by having industry work more closely with different populations of people with disabilities throughout the design and testing of new products. Working closely with different groups helps developers better anticipate user needs and pursue universal design.

Question 3. What more can Congress do to assist with these efforts?

Answer. One significant challenge is that the need to design for accessibility for people with disabilities is still underappreciated among technologists. One way to change this is to address this problem at the colleges and universities training the next generation of computer scientists and engineers. For example, Congress could establish NSF-funded Centers of Excellence for Accessible Design to prioritize this skillset and develop more curriculum. In addition, Congress should explore ways to encourage and support more people with disabilities to pursue careers in technology-related fields so they can be involved from the outset in the design and creation of more technologies. Finally, Congress should work to increase access to technology for people with disabilities, including by ensuring that programs designed to close the digital divide, such as PC or Internet access, are updated for newer technologies.

Question 4. As we see machine learning and AI increasingly embedded in products and services that we rely on, there are numerous cases of these algorithms falling short of consumer expectations. For example, Google and Facebook both promoted fraudulent news stories in the immediate wake of the Las Vegas Shooting because

⁴ See “Projections of industry employment, 2016–2026,” Bureau of Labor Statistics, <https://www.bls.gov/careeroutlook/2017/article/projections-industry.htm>.

⁵ “Google pledges \$1 billion to prepare workers for automation,” Engadget, October 13, 2017, <https://www.engadget.com/2017/10/13/grow-with-google/>.

⁶ See Rob Atkinson, “How a knowledge tax credit could stop decline in corporate training,” *The Hill*, <http://thehill.com/blogs/pundits-blog/finance/235018-how-a-knowledge-tax-credit-could-stop-decline-in-corporate>.

of their algorithms.⁷ YouTube Kids is a service designed for children, and marketed as containing videos that are suitable for very young children. In November, YouTube Kids promoted inappropriate content due to algorithms.⁸ While the use of machine learning and AI holds limitless positive potential, at the current point, it faces challenges where we should not risk getting it wrong. Should there be any formal or informal guidelines in place for what tasks are suitable to be done by algorithms, and which are still too important or sensitive to turn over; and what more can be done to ensure better and more accurate algorithms are used as you work to better develop this technology?

Answer. AI, much like humans, is fallible. There should always be some oversight of AI, just as there should always be some oversight of humans. It is not a problem if AI systems make mistakes, unless these mistakes go undetected. So the key objective, whether a decision is being made by a computer or a human, is whether there is sufficient oversight appropriate to the level of risk for the individuals involved. This will likely be context dependent. This is one reason why it is inappropriate to talk about industry-wide regulation of AI, and much more appropriate to talk about industry-specific regulations of AI. For example, the Department of Transportation may have specific requirements for the types of oversight it wants for autonomous vehicles that looks very different than the type of oversight the Securities and Exchange Commission needs for AI-driven stock trading.

Question 5. Machine learning and AI hold great promise for assisting us in preventing cybersecurity attacks. According to an IBM survey of Federal IT managers, 90 percent believe that artificial intelligence could help the Federal Government defend against real-world cyber-attacks. 87 percent think AI will improve the efficiency of their cybersecurity workforce.⁹ While this is promising, the Federal Government currently faces a shortage of qualified cybersecurity employees, and to make matters worse, the pipeline of students studying these topics is not sufficient to meet our needs. A recent GAO report found that Federal agencies have trouble identifying skills gaps, recruiting and retaining qualified staff, and lose out on candidate due to Federal hiring processes. The George Washington University Center for Cyber & Homeland Security recently released a report titled “Trends in Technology and Digital Security” which stated:

“Traditional security operations centers are mostly staffed with tier one analysts staring at screens, looking for unusual events or detections of malicious activity. This activity is similar to physical security personnel monitoring video cameras for intruders. It is tedious for humans, but it is a problem really well-suited to machine learning.”¹⁰

What effect will machine learning and AI will have on cybersecurity; and how do you think the Federal Government can best leverage the benefits offered by machine learning and AI to address our cybersecurity workforce shortage?

Answer. AI is very good at specific tasks, such as pattern recognition and anomaly detection. This means that it will be useful for identifying attacks in real time, and it will be an especially important line of defense against zero-day attacks (*i.e.*, attacks that use a previously undisclosed vulnerability). AI might also help developers eliminate certain types of vulnerabilities which may be identifiable at the outset, much like a spell-checker or grammar-checker can review documents. However, AI will not be a panacea, as many cybersecurity risks are the result of poor implementation and a lack of adherence to best practices.

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. GARY PETERS TO VICTORIA ESPINEL

Question 1. I am concerned by recent reports in *Nature*, *The Economist*, and *Wall Street Journal* about large tech firms monopolizing the talent in AI and machine learning. This concentration of talent can lead to several negative outcomes including long-term wage stagnation and income inequality.

⁷ NYT: After Las Vegas Shooting, Fake News Regains Its Megaphone, Kevin Rose, 10/02/2017 <https://www.nytimes.com/2017/10/02/business/las-vegas-shooting-fake-news.html>

⁸ NYT: On YouTube Kids, Startling Videos Slip Past Filters, Sapna Maheshwari, 11/04/2017 <https://www.nytimes.com/2017/11/04/business/media/youtube-kids-paw-patrol.html>

⁹ INFORMATION MANAGEMENT: AI seen as key tool in government's cybersecurity defense, Bob Violino, 11/30/2017 <https://www.information-management.com/news/artificial-intelligence-seen-as-key-tool-in-governments-cybersecurity-defense>

¹⁰ <https://cchs.gwu.edu/sites/cchs.gwu.edu/files/downloads/Fall%202017%20DT%20symposium%20compendium.pdf>

In your opinion, what steps or incentives might mitigate this concentration, encourage AI-experts to work at small and medium enterprises, or launch their own start-up with the goal of growing a business (rather than having a goal of being bought out by one of the tech giants)? Similarly, what incentives might encourage AI experts to become educators and trainers to help develop the next generation of AI experts?

How can the Federal Government compete with the tech giants to attract experts needed to develop and implement AI systems for defense and civil applications?

Answer. Artificial intelligence (AI) is a burgeoning field, with market dynamics that are quickly evolving. While the competition for AI expertise is certainly fierce, it is important to remember that the economic benefits of AI will be spread throughout the economy. By helping people make better data-driven decisions, AI is stimulating growth in every industry sector. It is helping to optimize manufacturing, improve supply chains, secure networks, and enhance products and services.

The history of the technology industry suggests that innovation will continue to emerge from enterprises of all sizes. Indeed, BSA's membership is a testament to just how fiercely competitive the enterprise technology landscape is. Datastax, DocuSign, Salesforce, Splunk and Workday are just a few of the young companies that have disrupted the industry over the past 10 years and contributed to a wave of innovation that has made the U.S. software industry the envy of the world. Moreover, despite intense competition for AI expertise, small and medium-sized firms continue to play an incredibly important role in driving AI innovation. In fact, a recent study found that there are currently more than 2,000 AI startups that have raised almost \$30 billion in funding.¹

Because AI will be a huge driver of the global economy in the years ahead, it is vital that we examine the important issues that you have raised to ensure that the United States remains the global hub for AI innovation. There are three specific ways in which the government can increase the talent pool and attract that talent to the government for defense and civil applications. First, the government should increase its commitment to STEM education so that the United States prepares more men and women for AI-related careers. As part of this effort, government agencies can explore partnerships with academic institutions to provide internships to students studying in this field, which could pique their interests in pursuing careers in public service, and develop opportunities for academic researchers to share their technical expertise with government agencies. Second, the government should increase its funding for AI research, providing targeted investments into the "high risk, high reward" areas of basic research that are typically underfunded by the private sector.² Third, the government should be ambitious in its goals. The greater the vision for how AI will improve government services and capabilities, the better it will do in attracting talent.

Question 2. Many have predicted that AI will have a profound effect on the labor market. Most predict that low-wage, routine-based jobs will be under the most pressure for replacement by AI. Meanwhile, recent advancements in technology has led to job creation that will mostly require highly-skilled, highly-educated workers. What evidence have you seen regarding businesses incorporating this labor shift into their business plans?

Answer. The benefits of AI will be widespread, likely enhancing operations in every industry. As a result, AI also will likely create shifts in the labor market across the economy. The precise impact of AI on employment is uncertain. However, it is clear that AI will create new opportunities within existing jobs and new roles that require skills that the current workforce does not yet have. As a result, many BSA companies have launched initiatives to train employees, youth, and military veterans to help meet the demands of the future labor market. BSA would like to work with Congress to ensure we have the right programs and resources in place

¹See Vala Afshar, *AI Is Transformational Technology and Major Sector Disruptor*, Huffington Post (Dec. 5, 2017), https://www.huffingtonpost.com/entry/ai-is-transformational-technology-and-major-sector_us_5a259dbfe4b05072e8b56b6e.

²See Jason Furman, *Is this Time Different? The Opportunities and Challenges of Artificial Intelligence, AI Now: The Social and Economic Implications of Artificial Intelligence in the Near Term* (July 7, 2016), available at <https://goo.gl/pzFDYw> ("In 2015, American businesses devoted almost 1.8 percent of GDP to research and development, the highest share on record. But government investments in R&D have fallen steadily as a share of the economy since the 1960s. While business investment is critical, it is not sufficient. Basic research discoveries often have great social value because of their broad applicability, but there tends to be underinvestment in basic research by private firms because it is difficult for a private firm to appropriate the gains from such research. In fact, while the private sector accounts for roughly two-thirds of all spending on R&D, it is important to keep in mind that it largely invests in applied research while the Federal Government provides 60 percent of the funding for basic research.").

for the jobs of the future. We would be happy to come in and discuss with you the initiatives of the software industry that address this important issue.

**RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. MAGGIE HASSAN TO
VICTORIA ESPINEL**

Question 1. Artificial intelligence, or AI, holds tremendous promise for individuals who experience disabilities. For example, Google and Microsoft have technologies to process language and speech and translate it into a text format to assist individuals who are deaf and hard of hearing. Other technologies will go even further to improve the lives of people with disabilities and I would like to learn more from the panel about what we can expect. What other specific technologies are you aware of in the AI space that will help people who experience disabilities?

Answer. There are numerous ways in which AI is being used to improve the lives of people who experience disabilities. Below, I highlight a few examples.

- Visual impairment—Microsoft recently released an intelligent camera app that uses a smartphone's built-in camera functionality to describe to low-vision individuals the objects that are around them. See Microsoft, *Seeing AI*, <https://www.microsoft.com/en-us/seeing-ai/>. The app opens up new possibilities for the visually impaired to navigate the world with more independence.
- Autism—IBM researchers are using AI to develop tools that will help people with cognitive and intellectual disabilities, such as autism, by breaking down complex sentences and phrases to help them better understand normal speech and communicate more effectively. See <https://www-03.ibm.com/able/content-clarifier.html>.
- Accessible public transportation—As part of a public-private partnership, an innovative project is underway that aims to help disabled people and those with special needs access public transportation by providing real-time information through an Internet of Things system that helps them find the right track, platform, train, and place to board, and alerts them when to disembark. See David Louie, *Artificial Intelligence Research Is Helping the Disabled Use Public Transportation*, (July 12, 2017), <http://abc7news.com/technology/ai-being-used-to-help-disabled-using-public-transportation/2210112/>.
- Mobility Impairments—Microsoft's Windows 10 operating system introduced Eye Control, a built-in eye tracking feature that enables people with motor neurone disease and other mobility impairments to navigate their computers. See Tas Bindi, *Microsoft Using AI to Empower Living With Disabilities*, Zdnet (Nov. 15, 2017), <http://www.zdnet.com/article/microsoft-using-ai-to-empower-people-living-with-disabilities/>.
- Alzheimer's—Researchers in Italy and Canada have developed machine-learning algorithms to help identify patients that are at risk of developing Alzheimer's. In early tests, the technology has identified changes in the brain that lead to Alzheimer's almost a decade before clinical symptoms would appear. See Daisy Yuhas, *Doctors Have Trouble Diagnosing Alzheimer's. AI Doesn't*, NBC News (Oct. 30, 2017), <https://www.nbcnews.com/mach/science/doctors-have-trouble-diagnosing-alzheimer-s-ai-doesn-t-ncna815561>.

As researchers continue to apply AI to new settings, the myriad ways in which AI is used to enhance the lives of people with disabilities will only increase.

Question 2. How will manufacturers and developers work to perfect this technology so that it can truly be a reliable tool for these individuals?

Answer. BSA members that design and offer AI products and services have strong incentives to ensure that the technology is reliable, as they understand that building trust and confidence in AI systems is integral to successful and widespread deployment of AI services.

There are a number of strategies that companies already employ to accomplish this objective. For example, one key priority is ensuring access to vast, robust, and representative data sets. Because AI technologies process and learn from data inputs, ensuring sufficient quantity and quality of data used to train AI systems is very important to enhancing the reliability of these services.

In addition, another key step companies take to enhance reliability is testing their AI systems to ensure that they operate as intended, and making appropriate adjustments where they identify errors.

Companies also recognize the need to protect AI systems from cyberattacks and are investing heavily in the development of advanced security tools.

As companies continue to seek to expand the capabilities of AI technologies, investment in research and development will continue to be important to unleash the full potential of innovation, strengthen cybersecurity, and enhance the overall reliability of AI systems.

Question 3. What more can Congress do assist with these efforts?

Answer. Congress can play a very important role in facilitating the deployment of AI services that help people with disabilities by ensuring, more broadly, that the U.S. maintains a flexible policy framework that spurs innovation in AI.

Specifically, as I highlighted in my testimony, I think Congress can assist with these efforts in three key ways. First, Congress should pass the OPEN Government Data Act, which recognizes that government-generated data is a national resource that can serve as a powerful engine for creating new jobs and a catalyst for economic growth, and that it is incredibly valuable in fostering innovation in AI and other data-driven services.

Second, Congress should support efforts to promote digital trade and facilitate data flows. In a global economy, real-time access to data around the world has become increasingly critical for AI and other digital services to function. As a result, Congress should support modernizing trade initiatives, such as NAFTA, that seek to facilitate digital trade and limit inappropriate restrictions on cross-border data transfers.

Third, Congress should promote U.S. investment in AI research, education, and workforce development to ensure that the U.S. remains globally competitive. Strategic investment in education and workforce development can help ensure that the next generation and our current workforce are prepared for the jobs of the future. In addition, promoting public sector and incentivizing private sector research will be essential to unlocking additional capabilities that AI can provide.

Question 4. As we see machine learning and AI increasingly embedded in products and services that we rely on, there are numerous cases of these algorithms falling short of consumer expectations. For example, Google and Facebook both promoted fraudulent news stories in the immediate wake of the Las Vegas Shooting because of their algorithms.³ YouTube Kids is a service designed for children, and marketed as containing videos that are suitable for very young children. In November, YouTube Kids promoted inappropriate content due to algorithms.⁴ While the use of machine learning and AI holds limitless positive potential, at the current point, it faces challenges where we should not risk getting it wrong. Should there be any formal or informal guidelines in place for what tasks are suitable to be done by algorithms, and which are still too important or sensitive to turn over; and what more can be done to ensure better and more accurate algorithms are used as you work to better develop this technology?

Answer. Because AI is ultimately a technology that is intended to help people and organizations make better uses of data, I would be reluctant to prescribe any bright line rules about when its use may or may not be appropriate. However, it is important for companies that develop AI systems, and their customers, to consider the unique risks and potential unintended consequences that can arise when AI is deployed in particular settings. While AI is an invaluable tool for making sense of large quantities of data, there are settings where the intuition of subject matter experts will remain important. For instance, while AI systems certainly have an important role to play in helping to diagnose patients, they are a resource for a medical professional to consider in making a diagnosis or prescribing treatment, they should not be replacing a doctor's judgment.

Question 5. Machine learning and AI hold great promise for assisting us in preventing cybersecurity attacks. According to an IBM survey of Federal IT managers, 90 percent believe that artificial intelligence could help the Federal Government defend against real-world cyber-attacks. 87 percent think AI will improve the efficiency of their cybersecurity workforce.⁵

While this is promising, the Federal Government currently faces a shortage of qualified cybersecurity employees, and to make matters worse, the pipeline of students studying these topics is not sufficient to meet our needs. A recent GAO report found that Federal agencies have trouble identifying skills gaps, recruiting and retaining qualified staff, and lose out on candidate due to Federal hiring processes.

³ NYT: After Las Vegas Shooting, Fake News Regains Its Megaphone, Kevin Rose, 10/02/2017
<https://www.nytimes.com/2017/10/02/business/las-vegas-shooting-fake-news.html>

⁴ NYT: On YouTube Kids, Startling Videos Slip Past Filters, Sapna Maheshwari, 11/04/2017
<https://www.nytimes.com/2017/11/04/business/media/youtube-kids-pau-patrol.html>

⁵ INFORMATION MANAGEMENT: AI seen as key tool in government's cybersecurity defense, Bob Violino, 11/30/2017 <https://www.information-management.com/news/artificial-intelligence-seen-as-key-tool-in-governments-cybersecurity-defense>

The George Washington University Center for Cyber & Homeland Security recently released a report titled “Trends in Technology and Digital Security” which stated:

“Traditional security operations centers are mostly staffed with tier one analysts staring at screens, looking for unusual events or detections of malicious activity. This activity is similar to physical security personnel monitoring video cameras for intruders. It is tedious for humans, but it is a problem really well-suited to machine learning.”⁶

What effect will machine learning and AI will have on cybersecurity; and how do you think the Federal Government can best leverage the benefits offered by machine learning and AI to address our cybersecurity workforce shortage?

Answer. AI tools are revolutionizing network security, helping analysts parse through hundreds of thousands of security incidents per day to weed out false positives and identify threats that warrant further attention by network administrators. By automating responses to routine incidents and enabling security professionals to focus on truly significant threats, AI-enabled cyber tools are helping enterprises stay ahead of their malicious adversaries. For instance, AI has helped an enterprise security operations center “reduce the time to remediate spearphishing attacks from three hours per incident to less than two minutes per incident.”⁷ Importantly, AI is also helping to train the next generation of security analysts, teaching them to more quickly identify threats that need to be escalated through the chain of command.⁸ Greater deployment of AI is therefore a critical factor for addressing the cyber workforce shortage, which experts now estimate will climb 1.8 million positions by 2022.

However, AI alone will not solve the cyber workforce shortage. It is therefore incumbent on governments and industry to work collaboratively to grow the pipeline of cyber talent. To that end, BSA recently launched a new cybersecurity agenda⁹ that highlights four pathways for developing a 21st century cybersecurity workforce:

- *Increase access to computer science education:* Expand cybersecurity education for K–12 as well as in undergraduate computer science programs, increase scholarships, and incentivize minority students.
- *Promote alternative paths to cybersecurity careers:* Launch careers through apprenticeship programs, community colleges, cybersecurity “boot camps,” and government or military service.
- *Modernize training for mid-career professionals:* Reform Trade Adjustment Assistance, and update other mid-career re-training programs, to provide American workers with high-demand cybersecurity and IT skills as digitalization transforms the global economy.
- *Improve the exchange of cybersecurity professionals between the government and private sector:* Enable private sector experts to join the government for periodic or short-term assignments.

**RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. AMY KLOBUCHAR TO
DR. DARIO GIL, PH.D.**

Question. While I was at the hearing there was significant discussion about the future security applications for machine learning and artificial intelligence. As the Ranking Member on the Rules Committee, I am working with Senators Lankford, Harris and Graham on a bill to upgrade our election equipment to protect against cyber-attacks. The Department of Homeland Security recently confirmed that hackers targeted 21 states’ election systems in the run-up to the 2016 election. As we prepare for 2018 and beyond, we must ensure that our election systems are secure, both from a hardware and a software perspective because election security is national security. Dr. Gil, can artificial intelligence and machine learning be used to identify and prevent cyber-attacks?

Answer. The power of AI, like most machine learning techniques, lies in identifying broader trends, building models of what normal and expected behavior is and

⁶<https://cchs.gwu.edu/sites/cchs.gwu.edu/files/downloads/Fall%202017%20DT%20symposium%20compendium.pdf>

⁷See Robert Lemos, *AI Is Changing SecOps: What Security Analysts Need to Know*, TechBeacon (Dec. 19, 2017), <https://techbeacon.com/ai-changing-secosps-what-security-analysts-need-know>.

⁸*Id.*

⁹BSA/The Software Alliance, A Cybersecurity Agenda for the Connected Age, available at www.bsa.org/~media/Files/Policy/BSA_2017CybersecurityAgenda.pdf.

flagging anomalies. There has been tremendous value shown by deploying AI in the field of security, and we use it in a plethora of use cases: in, applications to flag systems and networks, to monitor for anomalies and raise alerts when these behaviors change. Such anomalous behavior may indicate an attack (or a benign error). AI has also been leveraged for generalizations to make it easier to identify new instances of known attacks. It can be used to learn about malware and exploits of vulnerabilities and can use that to detect new infections and intrusions better than rule-based systems. AI and AI-based techniques can also help in hardening security protections to make it more difficult for attackers to successfully exploit a system. For example, automating tests (fuzzing) to probe for vulnerabilities using reinforcement learning can be more efficient than an exhaustive scan. It should be borne in mind that while there is no panacea for security, AI is a very powerful tool that can be employed to increase security when combined with a standard suite of best practices.

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. TOM UDALL TO
DR. DARIO GIL, PH.D.

Question 1. As you are aware, New Mexico is home to two national laboratories—Sandia and Los Alamos. Can you speak of any partnership you have with the national laboratories?

Answer. We have had, and continue to have, a number of partnerships with Los Alamos and Sandia National Labs. At this time, we cannot comment on individual projects, but we value the joint work we have with the labs and all of our external partners.

Question 2. Can you speak of some of the ways that government funded artificial intelligence development is now being used in the private sector?

Answer. Government-funded AI development is currently being used by IBM in the following ways:

- U.S. Army Research Labs funded the development of technologies for rule-based systems and policy management. The technology developed is being used now by IBM to improve and support the ability of employees in a company to configure computer software to efficiently execute high-volume, highly transactional process functions, boosting capabilities and saving time and money.
- U.S. Army Research Labs funded the development of AI enabled algorithms for analyzing and understanding properties of large numbers of moving objects. The technology developed has been used to create commercial cloud-based services such as The Weather Company LongShip service to predict the influence of weather impacts on traffic. It also has been incorporated into commercial software products such as DB/2.

Looking more broadly, there are a variety of ways in which government funded research can be used by the private sector. These include:

- Government funded technology has been used to create dual-purpose technologies—those which serve the needs of the private sector as well as the needs of the government agency which sponsored the work. One expression of this is when the private sector produces a COTS (Custom Off the Shelf) product which allows the government to meet their specific requirement at a lower overall cost.
- Government funded research (primarily in basic sciences) has a research horizon which is typically longer than the private sector. As a result, government funded research has been used to create technology at the earlier stage. There are many instances where government support has been used for technologies at the TRL (Technology Readiness Level) of 1–4, and the successful ones from this level of exploration have led to commercialized products at TRL level 5 and above.
- Government funded technology has been used to produce open source software, which private sector companies develop further and use to create new offerings. In many cases, the open source software is used by academics and the public at large for knowledge creation.
- Government funded alliances have driven collaboration between private sector researchers, academics and government researchers. This cross-fertilization of researchers working towards a common goal has been beneficial to all parties—including employment opportunities for students, infusion of new ideas in industry activities, improvements in government, and formation of lasting collaborations.

RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. GARY PETERS TO
DR. DARIO GIL, PH.D.

Question. A major challenge AI and machine learning developers need to address is the ability to ensure prolonged safety, security, and fairness of the systems. This is especially true of systems designed to work in complex environments that may be difficult to replicate in training and testing, or systems that are designed for significant learning after deployment. Dr. Gil, you testified that IBM is looking to build trust in AI by following a set of principles to guide your development and use of AI systems. Would you please provide more detail about how these principles are being implemented? How will these principles prevent a system designed to learn after deployment from developing unacceptable behavior over time?

Answer. The currently available AI products, such as factory robots, personal digital assistants, and healthcare decision support systems, are designed to perform one narrow task, such as assemble a product, provide a weather forecast or make a purchase order, or help a radiologist interpret an X-ray. When these technologies learn after deployment, they do so in the context of that narrow task, and do not have the ability to learn other tasks on their own, even similar ones. The kind of AI systems that can acquire new skills and perform new cognitive and reasoning tasks autonomously are actively being researched. This effort does not include only the development of underlying learning and reasoning capabilities; the AI research community is actively pursuing the capabilities that would ensure the safety, security and fairness of these systems.

To start with, the principles of safe design are applied to a wide variety of engineered systems, such as trains, safety breaks, industrial plants, flight autopilot systems, and robotic laser surgery. Some of these principles apply directly to the design of AI systems, some will be adapted, and new ones will have to be defined. For example, it is possible to constrain the space of outcomes or actions a robot can perform, to ensure that it does not accidentally come into contact with human workers and cause injury. Similarly, robots in complex environments that encounter completely new situations could be designed to require human intervention. Another direction is to embed principles of safe and ethical behavior in the AI reasoning mechanisms, so that they can distinguish between right and wrong actions.

With respect to the fairness of the AI systems, we are currently pursuing a range of efforts aimed at developing and embedding in our services and offerings techniques for bias detection, certification, and mitigation. For example, we have developed algorithms that can de-bias training data so that any AI system that learns from such data does not discriminate against protected groups (e.g., those defined by race or gender). We also are working on using blockchain to ensure the integrity of an AI system by making sure that it is secure, auditable and used as intended. We also are developing capabilities to enhance the explainability and interpretability of AI systems, so that unacceptable behaviors can be easily discovered and removed.

IBM has established the following principles for the artificial intelligence/cognitive era:

Purpose: The purpose of AI and cognitive systems developed and applied by the IBM company is to augment human intelligence. Our technology, products, services and policies will be designed to enhance and extend human capability, expertise and potential. Our position is based not only on principle but also on science. Cognitive systems will not realistically attain consciousness or independent agency. Rather, they will increasingly be embedded in the processes, systems, products and services by which business and society function—all of which will and should remain within human control.

Transparency: For cognitive systems to fulfill their world-changing potential, it is vital that people have confidence in their recommendations, judgments and uses. Therefore, the IBM company will make clear:

- When and for what purposes AI is being applied in the cognitive solutions we develop and deploy.
- The expertise that informs the insights of cognitive solutions, as well as the methods used to train those systems and solutions.
- The principle that clients own their own business models and intellectual property and that they can use AI and cognitive systems to enhance the advantages they have built. We will work with our clients to protect their data and insights, and will encourage our clients, partners and industry colleagues to adopt similar practices.

Skills: The economic and societal benefits of this new era will not be realized if the human side of the equation is not supported. This is uniquely important with cognitive technology, which augments human intelligence and expertise and works collaboratively with humans. Therefore, the IBM company will work to help students, workers and citizens acquire the skills and knowledge to engage safely, securely and effectively in a relationship with cognitive systems, and to perform the new kinds of work and jobs that will emerge in a cognitive economy.

Data: Since AI is heavily based on data, IBM has developed a framework of best practices for data stewardship¹ that ensures great care and responsibility in data ownership, storage, security, and privacy. IBM abides by these practices and, as a result, serves as a data steward providing transparent and secure services. For example, we write client agreements with full transparency and will not use client data unless they agree to such use. We will limit that use to the specific purposes clearly described in the agreement. IBM does not put ‘backdoors’ in its products for any government agency, nor do we provide source code or encryption keys to any government agency for the purpose of accessing client data.

We are working on a range of efforts aimed at developing and embedding in our services and offerings techniques for bias detection, certification, and mitigation. For example, we are working on improving the accuracy of directly interpretable decision-support algorithms, such as decision trees and rule sets, as well as enhancing the interpretability of deep learning neural net models.

Moreover, as we develop innovative AI systems, we are guided by the principles of safety engineering. Some of these principles could be directly applied to the design of AI systems, some will be adapted, and new ones will have to be defined. For example, robots in complex environments that encounter completely new situations could be designed to require human intervention. Another direction is to embed principles of safe and ethical behavior in their reasoning mechanisms, so they can distinguish between right and wrong actions.

Finally, we are working to develop AI systems that act according to human values that are relevant for the scenarios and communities in which such systems will be deployed. This means constraining the learning, reasoning and optimization machinery inside AI systems to behavioral constraints. These constraints will ensure that the actions of the AI system comply with values and guidelines that humans define as appropriate for the specific use case and application. Such behavioral constraints should be learned *offline* (*i.e.*, by training the system with data or via simulation), modified only by humans, and given a higher priority compared to *online* policies (outcomes that an AI system learns post-deployment, which are based on reinforcement learning or other machine learning approaches aimed at reward maximization and optimization).

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. MAGGIE HASSAN TO
DR. DARIO GIL, PH.D.

Question 1. Artificial intelligence, or AI, holds tremendous promise for individuals who experience disabilities. For example, Google and Microsoft have technologies to process language and speech and translate it into a text format to assist individuals who are deaf and hard of hearing. Other technologies will go even further to improve the lives of people with disabilities and I would like to learn more from the panel about what we can expect. What other specific technologies are you aware of in the AI space that will help people who experience disabilities?

Answer. AI technologies are enabling many exciting new assistive functions by enhancing machines’ ability to see, hear, interpret complex signals, and operate in the real world through action and dialog. Essential building blocks for these new capabilities are machine vision, speech to text, text to speech, natural language understanding and generation, emotion recognition, and machine learning to interpret sensor data.

For example, with *AI vision*, it is now becoming possible to describe an image, a local environment, or a video to a person with visual impairment. Further, these technologies will soon support wearable assistants that can recognize people, objects, landmarks and obstacles in the environment, and guide a person safely to an unfamiliar destination.

AI speech to text capabilities, coupled with natural language understanding, enable a quadriplegic individual to control their environment through speech commands, providing a new level of autonomy. *Machine learning* techniques can trans-

¹IBM: *Data Responsibility@IBM*, [Khttps://www.ibm.com/blogs/policy/wp-content/uploads/2017/10/IBM_DataResponsibility-A4_WEB.pdf](https://www.ibm.com/blogs/policy/wp-content/uploads/2017/10/IBM_DataResponsibility-A4_WEB.pdf)

late brain and nerve signals into commands for prosthetic limbs and convey a sense of touch.

AI natural language understanding and generation enable communication of knowledge in the form most easily understood by an individual, whether that means generating a description of a graph for a blind person, reading text aloud for a person with dyslexia, simplifying a complex document for a person with an intellectual disability, or, one day, translating between spoken languages and sign languages used by people who are deaf and hard of hearing.

AI embedded in autonomous vehicles, intelligent wheelchairs and interactive assistance robots will provide physical independence and assistance for many. For example, IBM, the CTA (Consumer Technology Association) Foundation, and Local Motors are exploring applications of Watson technologies to developing the world's most accessible self-driving vehicle, able to adapt its communication and personalize the overall experience to suit each passenger's unique needs.

Machine learning on sensor data from instrumented environments can support an older adult in living independently and safely at home by learning their normal patterns of behavior and providing assistance and alerts.

Just as importantly, AI technologies will benefit people with disabilities by analyzing websites and applications, finding and fixing accessibility problems in a more automated way than was previously possible.

Question 2. How will manufacturers and developers work to perfect this technology so that it can truly be a reliable tool for these individuals?

Answer. Perfecting AI technologies will take significant experimentation and depends on the availability of data. For these core technologies to be reliable, applications for people with disabilities should be explored as early as possible and used to drive requirements. This includes consulting with, and testing by, people with disabilities in realistic environments.

At IBM, we are exploring the potential of AI technologies to support people with disabilities, and older adults through several initiatives and collaborations. IBM researchers are exploring how Watson's language-processing software could help people with cognitive disabilities by simplifying text, how older adults' patterns of activity can be learned, how a blind person can navigate and find objects effectively using machine vision, and how AI can enable our accessibility test tools to move from pointing out problems to actively suggesting solutions.

Secondly, it is essential that people with disabilities are represented adequately in training data, to prevent new forms of discrimination from emerging. For example, an authentication system based on voice recognition should be able to recognize people with dysarthric speech. Manufacturers and developers applying AI technologies should incorporate mechanisms to recognize and gracefully handle exceptions, falling back on human judgment for cases that are outside their training.

Question 3. What more can Congress do to assist with these efforts?

Answer. For AI to deliver the promised economic and societal benefits to a broader range of people, including people with disabilities, both policy support and public investment from the U.S. Government are critical.

Given the great diversity of human abilities, it is a challenge for manufacturers and developers to ensure diversity in training data. For example, speech recognition training data should ideally include people who stutter. Government investment in initiatives to make diverse data broadly available would accelerate our ability to make AI technology more inclusive, and to apply AI techniques to new accessibility problems. Government support for controlled studies with people with disabilities will also accelerate the inclusion of people with disabilities.

Access to other forms of data is also critical. An indoor-outdoor navigation system for blind people relies on public outdoor maps, but indoor maps are privately owned. A centralized mechanism to share such maps for accessibility purposes would remove a practical barrier to the widespread use of such systems. AI vision techniques depend on the use of images or video to describe people and objects to people with visual impairment. Government leadership is needed to address privacy concerns of individuals and copyright concerns of organizations over the use of images of their faces or products for accessibility purposes. There is a copyright exception for converting books into braille, and a similar solution could be effective here.

Secondly, policy support can help to counter the danger of new forms of discrimination. For example, reinstating the Department of Justice rulemaking on accessibility guidelines for public websites would emphasize the importance of accessibility, and spur efforts by industry to include people with disabilities in development.

Most of the examples in question 1 describe ways that AI technologies can assist people with sensory or physical impairments. There is a need to foster standards

and policies to address the needs of people with cognitive disabilities, which will encourage application of AI technologies to these challenges.

Question 4. As we see machine learning and AI increasingly embedded in products and services that we rely on, there are numerous cases of these algorithms falling short of consumer expectations. For example, Google and Facebook both promoted fraudulent news stories in the immediate wake of the Las Vegas Shooting because of their algorithms.² YouTube Kids is a service designed for children, and marketed as containing videos that are suitable for very young children. In November, YouTube Kids promoted inappropriate content due to algorithms.³ While the use of machine learning and AI holds limitless positive potential, at the current point, it faces challenges where we should not risk getting it wrong. Should there be any formal or informal guidelines in place for what tasks are suitable to be done by algorithms, and which are still too important or sensitive to turn over; and what more can be done to ensure better and more accurate algorithms are used as you work to better develop this technology?

Answer. AI is already more capable than humans in narrow domains, some of which involve delicate decision making. Humanity is not threatened by them, but many people could be affected by their decisions. Examples are autonomous online trading agents, media and news services, and soon autonomous cars. Even though AI algorithms are usually evaluated based on their accuracy, that is, their ability to produce correct results, this is only one component of a bigger picture. We need to be able to assess the impact of their decisions in the narrow domains where they will function.

To understand the suitability of an AI system with respect to performing a specific task, one must consider not only their accuracy, but also the context, the possible errors, and the consequences on the impacted communities. Furthermore, the assessment of risk should be carried out with respect to both the risk of “doing it” and the risk of “not doing it”, as in many fields we already know the consequences of wrong decisions made by humans. For example, melanoma detection from skin images is a task that AI algorithms can perform at high levels of accuracy. Even though there is still a possibility of error, it is beneficial to deploy such systems in healthcare decision support, in a way that would augment human decision-making process. On the other hand, let us consider automated trading systems. A bad decision in these systems may be (and has been) a financial disaster for many people. That will also be the case for self-driving cars. Some of their decisions will be critical and possibly affect lives. Because sectors like finance and transportation can carry large risks, protections have always been in place through existing regulations. These existing protections are properly designed to provide consumer protection even with the advent of new technologies like AI.

Finally, we believe that in many applications, rather than considering only fully autonomous AI solutions, the most effective approach is to build AI systems that support humans and work with them in performing a task. For example, in a breast cancer detection study, it has been shown that doctors and AI working together achieve a higher degree of accuracy than just doctors or AI separately.

Question 5. Machine learning and AI hold great promise for assisting us in preventing cybersecurity attacks. According to an IBM survey of Federal IT managers, 90 percent believe that artificial intelligence could help the Federal Government defend against real-world cyber-attacks. 87 percent think AI will improve the efficiency of their cybersecurity workforce.⁴

While this is promising, the Federal Government currently faces a shortage of qualified cybersecurity employees, and to make matters worse, the pipeline of students studying these topics is not sufficient to meet our needs. A recent GAO report found that Federal agencies have trouble identifying skills gaps, recruiting and retaining qualified staff, and lose out on candidate due to Federal hiring processes.

The George Washington University Center for Cyber & Homeland Security recently released a report titled “Trends in Technology and Digital Security” which stated:

“Traditional security operations centers are mostly staffed with tier one analysts staring at screens, looking for unusual events or detections of malicious activity.”

²NYT: After Las Vegas Shooting, Fake News Regains Its Megaphone, Kevin Rose, 10/02/2017 <https://www.nytimes.com/2017/10/02/business/las-vegas-shooting-fake-news.html>

³NYT: On YouTube Kids, Startling Videos Slip Past Filters, Sapna Maheshwari, 11/04/2017 <https://www.nytimes.com/2017/11/04/business/media/youtube-kids-paw-patrol.html>

⁴INFORMATION MANAGEMENT: AI seen as key tool in government’s cybersecurity defense, Bob Violino, 11/30/2017 <https://www.information-management.com/news/artificial-intelligence-seen-as-key-tool-in-governments-cybersecurity-defense>

This activity is similar to physical security personnel monitoring video cameras for intruders. It is tedious for humans, but it is a problem really well-suited to machine learning.”⁵

What effect will machine learning and AI have on cybersecurity; and how do you think the Federal Government can best leverage the benefits offered by machine learning and AI to address our cybersecurity workforce shortage?

Answer. AI and machine learning will be a disruptive force in the field of cybersecurity by providing the potential for aiding both in the defense and protection of critical infrastructure, leveling the playing field between large nation states and smaller niche players.

From a defensive standpoint, AI has shown promise in automating defenses, such as probing systems for weaknesses, including software vulnerabilities and configuration errors. Penetration testing and bug finding tools have benefited tremendously from AI techniques in improving their efficiency to more quickly evaluate systems for weaknesses and increase coverage of the evaluated space. Security monitoring tools have also benefited greatly from AI and will continue to do so as AI systems improve. Automation can be leveraged to process suspicious alerts and events that warrant investigation, performing many of the rote tasks typically performed by low-level analysts. These automated tools will provide an analyst a more complete picture of the events unfolding, highlight meaningful information and context, triage, and allow the analyst to provide a higher-level response. This can allow security analysts to investigate far more alerts than are currently possible, and hopefully make fewer errors in how those alerts are processed. Security operations can be conducted at machine-scale as opposed to human-scale.

The Federal Government can learn from the experience in research, industry and academia in leveraging AI to develop and deploy the next generation of AI-powered defenses that will be necessary to protect the Nation’s critical infrastructure. This requires significant leadership and outreach on behalf of the Government to industry and academia on the following fronts:

- Declare AI Leadership in Cyber Security as a national research and development priority.
- Evolve and develop the Nation’s cybersecurity strategy to address the AI-powered threats to critical infrastructure with AI-powered defenses.
- Initiate U.S. Government programs, through various policy and funding agencies (e.g., OSTP, DARPA, IARPA, NSF, NIST etc.) to fund and sponsor leading edge research in areas of intersection between AI and security
- Set policies and standards for procurement of next generation security controls by the U.S. Government.

RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. AMY KLOBUCHAR TO
DR. EDWARD W. FELTEN, PH.D.

Question. Political ads on the Internet are more popular now than ever. In 2016, more than \$1.4 billion was spent on digital advertisements and experts project that number will continue to increase. In October, I introduced the Honest Ads Act with Senators Warner and McCain, to help prevent foreign interference in future elections and improve the transparency of online political advertisements. We know that 90 percent of the ads that Russia purchased were issue ads meant to mislead and divide Americans. Increasing transparency and accountability online will benefit consumers and help safeguard future elections. Dr. Felten, can machine learning be used to help identify issue ads and stop misinformation from spreading online?

Answer. Yes, machine learning can be useful in several ways. Machine learning can help to classify the nature or topic of ads, to distinguish issue ads from others and to characterize the issue being addressed by an ad. Machine learning can be useful in determining the source of an ad, including in identifying when a single source is trying to disguise itself as a set of separate, independent sources. More broadly, machine learning can be helpful in identifying misinformation and disinformation campaigns, and in targeting countermeasures to maximize the impact on a harmful campaign while minimizing collateral damage.

Three caveats are in order, however. First, more research will be necessary to take full advantage of these opportunities. That research is best done using realistic datasets derived from platforms’ experience with past disinformation campaigns.

⁵<https://cchs.gwu.edu/sites/cchs.gwu.edu/files/downloads/Fall%202017%20DT%20symposium%20compendium.pdf>

Second, machine learning methods will necessarily be less than perfectly accurate. Not only will they fail to spot some disinformation campaigns, they will also sometimes misclassify content or a user as malicious when they are in fact benign. Appropriate use of machine learning in this setting will require both a careful technical evaluation of the likelihood of errors, and a policy approach that recognizes the harm that might be done by errors. Finally, machine learning systems for detecting anomalies depend on a variety of data sources and signals, and the success of machine learning depends on the characteristics of those sources and signals. Real-world data is sometimes erroneous and often incomplete, in ways that could frustrate the use of machine learning for this application or render it less accurate. Where data signals derive from the votes or clicks of users, the resulting system may be subject to gaming or manipulation, so such signals should be used with caution, especially in systems that aim to limit disinformation.

**RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. TOM UDALL TO
DR. EDWARD W. FELTEN, PH.D.**

Question. In your testimony, you discussed how adoption of artificial intelligence can inadvertently lead to biased decisions. What specific steps should the Federal Government and other users take to improve the data and ensure that datasets minimize societal bias—especially with regard to vulnerable populations?

Answer. The results of a machine learning system can only be as accurate as the dataset on which the system was trained. If a community is underrepresented in the dataset, relative to its representation in the population, then that community is likely to be poorly served by the system, as the system will not put enough weight on the characteristics of the underrepresented community.

Practitioners should take care to ensure that datasets are representative of the population, to the extent possible. Where this is not possible, deficiencies in the dataset should be noted carefully, and steps should be taken to mitigate the deficiencies. For example, it is sometimes possible to correct for a group's underrepresentation in a data analysis or machine learning procedure by putting greater weight on data points that represent that group. Additional statistical methods exist that can counteract the effect of non-representative datasets.

Another common source of error or bias in machine learning occurs when a system is tasked with learning from examples of past decisions made by people. If those past decisions were biased, the machine is likely to learn to replicate that bias. Whenever a system is trained based on past human decisions, care should be taken to consider the social and historical context of those past decisions and to look for indications of bias in the system's output, and anti-bias techniques should be used in designing or training the system if possible.

In addition to technical measures in the design and use of AI systems, the possibilities of bias—whether that AI will introduce bias or that AI will open opportunities to measure and counteract human bias—should be taken into account in making policy decisions. Consulting with technical experts, and including technical expertise in the policymaking conversation, are important steps toward good policy in this area.

**RESPONSE TO WRITTEN QUESTION SUBMITTED BY HON. GARY PETERS TO
DR. EDWARD W. FELTEN, PH.D.**

Question. I am concerned by recent reports in Nature, The Economist, and Wall Street Journal about large tech firms monopolizing the talent in AI and machine learning. This concentration of talent can lead to several negative outcomes including long-term wage stagnation and income inequality.

In your opinion, what steps or incentives might mitigate this concentration, encourage AI-experts to work at small and medium enterprises, or launch their own start-up with the goal of growing a business (rather than having a goal of being bought out by one of the tech giants)? Similarly, what incentives might encourage AI experts to become educators and trainers to help develop the next generation of AI experts?

How can the Federal Government compete with the tech giants to attract experts needed to develop and implement AI systems for defense and civil applications?

Answer. One approach to making smaller companies attractive to top AI talent is to adopt pro-competition policies generally. Because AI often relies on large datasets, and those datasets are more likely to be held by large companies, there may be a natural tendency toward concentration in AI-focused industry sectors. Public policy can help to ensure that smaller companies can be viable competitors.

The Federal Government can also provide some large, high quality datasets that may be useful to individuals and companies of all sizes.

At present, the demand for highly skilled AI experts exceeds the supply, leading to a scarcity of those experts in all but the best-funded environments. In the long run, steps to increase the education and training of AI professionals are the most important means to strengthen our national talent base and broaden the availability of expertise.

The talent pipeline can be widened at every stage. At the K–12 level, access to a good computer science course should be available to every student. A bipartisan coalition of states and nonprofit actors is working toward this goal. At the university and graduate level, access to education is limited by the number of trained faculty available to teach advanced AI and machine learning courses.

It is difficult to underestimate the importance of supporting a large and robust public research community. This ensures that access to the latest knowledge and techniques in AI is available to the public and not limited to a few companies' researchers. It widens the talent pipeline because AI research funding enables faculty hiring in AI, which increases the national capacity to train AI leaders. Federally-funded research projects serve as the main training ground for the next generation of research leaders.

RESPONSE TO WRITTEN QUESTIONS SUBMITTED BY HON. MAGGIE HASSAN TO
DR. EDWARD W. FELTEN, PH.D.

Question 1. Artificial intelligence, or AI, holds tremendous promise for individuals who experience disabilities. For example, Google and Microsoft have technologies to process language and speech and translate it into a text format to assist individuals who are deaf and hard of hearing. Other technologies will go even further to improve the lives of people with disabilities and I would like to learn more from the panel about what we can expect. What other specific technologies are you aware of in the AI space that will help people who experience disabilities?

Answer. There are many examples, of which I will highlight three here.

First, self-driving vehicles will improve mobility and lower the cost of transportation for people who are unable to drive. These vehicles will have major safety benefits in the long run, and they are already starting to benefit people with disabilities. Maintaining policies to encourage safety-conscious testing and deployment of self-driving vehicles will benefit all Americans, and especially those with disabilities.

Second, computer vision and image interpretation systems have the potential to help those with visual disabilities process information about their surroundings. These systems are demonstrating an increasing capacity to identify specific objects and people in complex scenes, and to model and predict what might happen next, such as warning of potential dangers.

Third, AI can help to identify barriers to accessibility. For example, Project Sidewalk at the University of Maryland combines crowdsourced data collection with AI techniques to build a database of curb, ramp, and sidewalk locations, and analyze it to identify accessibility problems. This can help city planners and property owners recognize accessibility failures.

Question 2. How will manufacturers and developers work to perfect this technology so that it can truly be a reliable tool for these individuals?

Answer. As with any new, complex technology, careful testing is needed to understand the implications of using a system. Such testing must be done in a realistic environment and must involve the community of potential users.

The best design practices are user-centered, meaning that the potential user community for a product is involved throughout the design process, from initial concept exploration through final testing. This is especially important if the designer might experience the world differently than the user community.

Question 3. What more can Congress do to assist with these efforts?

Answer. Three significant things that Congress can do are (1) provide funding for research on applications of AI for use by people with disabilities; (2) work with agencies to ensure they are giving proper attention to these issues and the interests of people with disabilities; and (3) highlight the need for work in this area and highlight the successes of those already working in the area.

Question 4. [citations omitted] As we see machine learning and AI increasingly embedded in products and services that we rely on, there are numerous cases of these algorithms falling short of consumer expectations. For example, Google and Facebook both promoted fraudulent news stories in the immediate wake of the Las

Vegas Shooting because of their algorithms. YouTube Kids is a service designed for children, and marketed as containing videos that are suitable for very young children. In November, YouTube Kids promoted inappropriate content due to algorithms. While the use of machine learning and AI holds limitless positive potential, at the current point, it faces challenges where we should not risk getting it wrong. Should there be any formal or informal guidelines in place for what tasks are suitable to be done by algorithms, and which are still too important or sensitive to turn over; and what more can be done to ensure better and more accurate algorithms are used as you work to better develop this technology?

Answer. In considering a switch from human to AI based decision making, we should not demand perfection of the AI system. The alternative to AI is often to rely on human judgment, which is also prone to bias and mistakes. Instead of demanding perfection of the AI system, an organization needs to understand the potential consequences of adopting the AI system well enough to conclude with justified confidence that switching to an automated system is an improvement and on balance the effects are benign.

There are also considerations of scale. AI can operate at larger scale (that is, on larger number amounts of data or more decisions per second) than any human organization can hope to achieve. As a result, in some cases the choice in designing a function is not between using AI and using a human, but rather between using AI and not providing the function at all. The tremendous value provided by many of today's information, communication, and publishing tools relies at least in part on the use of AI.

That said, the potential risks of AI-based systems must be considered and addressed. It is too early to establish formal guidelines, because not enough is known about how best to address these problems. Informal guidelines are needed, and the industry and other stakeholders should be encouraged to develop them collaboratively. Multi-stakeholder groups such as the Partnership on AI may be useful venues for these discussions. The guidelines, best practices, and technical tools to address these problems will evolve with time.

Switching a process based on human decision-making to one based on AI can have unpredictable consequences, so experimentation is needed in a safe environment to adequately understand the implications of such a change before it is made. Organizations can be more transparent by publishing information about what kinds of testing and analysis were done in preparation for the introduction of AI into an existing process, thereby enabling stakeholders to understand why a change was made and query the organization if concerns remain.

In many cases, an organization will have valid reasons, such as trade secrets or user privacy, to refrain from publishing the full details of how a system works. This does not preclude the organization from publishing information about how it tested the system and evaluated the pros and cons of adopting it.

Question 5. [citations omitted] Machine learning and AI hold great promise for assisting us in preventing cybersecurity attacks. According to an IBM survey of Federal IT managers, 90 percent believe that artificial intelligence could help the Federal Government defend against real-world cyber-attacks. 87 percent think AI will improve the efficiency of their cybersecurity workforce. While this is promising, the Federal Government currently faces a shortage of qualified cybersecurity employees, and to make matters worse, the pipeline of students studying these topics is not sufficient to meet our needs. A recent GAO report found that Federal agencies have trouble identifying skills gaps, recruiting and retaining qualified staff, and lose out on candidate due to Federal hiring processes. The George Washington University Center for Cyber & Homeland Security recently released a report titled "Trends in Technology and Digital Security" which stated:

"Traditional security operations centers are mostly staffed with tier one analysts staring at screens, looking for unusual events or detections of malicious activity. This activity is similar to physical security personnel monitoring video cameras for intruders. It is tedious for humans, but it is a problem really well-suited to machine learning."

What effect will machine learning and AI will have on cybersecurity; and how do you think the Federal Government can best leverage the benefits offered by machine learning and AI to address our cybersecurity workforce shortage?

Answer. As the GWU report suggests, cybersecurity tasks of a routine nature can be automated, thereby reducing the need for human operators do to lower-level work. This can free up workers to concentrate on higher-level tasks requiring more skill and judgment, and can help to mitigate the Federal Government's cybersecurity personnel shortage.

Notwithstanding these opportunities, the Federal Government will continue to face challenges in recruiting and retaining the best technology talent. Many proposals exist to address these challenges by improving hiring authorities, pay scales, and working conditions for Federal technology workers, and by instituting or expanding training and scholarship-for-service programs.

Question 6. Mr. Felten, as we heard you recently served as the deputy U.S. Chief Technology Officer at the White Office of Science and Technology Policy. One of the projects you worked on in that office was a major report on Artificial Intelligence. That report was one of the many important projects taken on by the Office of Science and Technology Policy in recent years. And it's extremely disappointing that President Trump has failed to nominate leaders for that office, now more than ten months into his presidency. That's the longest a president has gone without a science advisor since the Office of Science and Technology Policy was established in law in 1976. I've led two letters to President Trump urging him to nominate well-qualified experts to lead this office, but so far we have seen nothing from this administration. As a former leader at the Office of Science and Technology Policy, could you please explain why this office is important, and what kinds of qualities you look for in good nominees for this office?

Answer. OSTP's importance derives from the central role of science and technology in nearly every area of policy. In major policy areas such as national security and defense, transportation, education, and the economy, technology is critical to the most important challenges and opportunities. Making the best decisions in these areas requires input from and dialog with the technical community. Congress assigned that role within the White House to OSTP.

AI is just one of the areas of interest for OSTP, but it connects to many important policy questions. What should DoD's policy be on autonomous weapons systems, and what position should the United States take in international talks about such weapons? How will AI-driven automation affect the job market, and how can American schoolchildren and adults be educated and trained for the future workplace? What needs to be done to improve highway safety as automated vehicles become practical? How can American farmers, journalists, and businesses be freed to use drones, while strengthening our defenses against potential terrorist uses of the same technology? How will changes in information technology affect the mission of the Intelligence Community, and what kinds of people and capabilities will the IC need in the future? How will cybersecurity concerns affect all of these goals? Each of these questions can be better answered with the help of technical advisors who have deep domain knowledge, connections to the relevant technical communities, and a seat at the policy table.

A successful OSTP Director will be a trusted advisor to the President and the President's senior advisors, a liaison to departments and agencies on science and technology issues, and an ambassador to scientific and technical communities in the United States and around the world.

A candidate for OSTP Director should be a highly respected member of the scientific/technical community, with a reputation for technical knowledge and policy judgment. The candidate should be able to work successfully across disciplines, acquiring knowledge and providing advice across many subject areas with appropriate staff support. They should be able to work successfully within the unusual administrative and legal environment of the White House, and they should be able to recruit, motivate, and lead a team of highly-skilled domain experts and policy advisors.

Because the subject matter of science and technology is so extensive, and the United States is blessed with leading experts in so many specialties, no one person can hope to have the knowledge, experience, and connections needed to provide advice in all technical areas. A successful OSTP Director will recruit a team of topic-area advisors who can provide context and guidance in specific areas and can expand OSTP's "surface area" in coordinating with agencies, outside experts, and the public.



This page intentionally left blank.

