



Hearing on

The Need for Transparency in Artificial Intelligence

**Senate Committee on Commerce, Science, & Transportation
Subcommittee on Consumer Protection, Product Safety, and Data Security**

**September 12, 2023, at 2:30 p.m.
Russell Senate Office Building
Washington, DC**

**Testimony of Victoria Espinel
CEO
BSA | The Software Alliance**

**Testimony of Victoria Espinel
CEO of BSA | The Software Alliance
Hearing on The Need for Transparency in Artificial Intelligence**

**Before the Senate Committee on Commerce, Science, & Transportation
Subcommittee on Consumer Protection, Product Safety, and Data Security**

September 12, 2023

Good afternoon Chairman Hickenlooper, Ranking Member Blackburn, and members of the Subcommittee. My name is Victoria Espinel, and I am the CEO of BSA | The Software Alliance.¹

BSA is the leading advocate for the global enterprise software industry.² Our members are at the forefront of developing cutting-edge services — including AI — and their products are used by businesses across every sector of the economy. I commend the Subcommittee for convening today's hearing and thank you for the opportunity to testify. I also appreciate this Committee's longstanding focus on AI, including your efforts to establish the National Artificial Intelligence Initiative and your outreach to BSA to learn about how our companies are implementing the AI Risk Management Framework released earlier this year by the National Institute of Standards and Technology.

Nearly six years ago, I testified before this Committee at a hearing focused on the building blocks of machine learning and artificial intelligence.³ Chair Cantwell and Senator Young also introduced one of the first AI bills that year. Since then, the building blocks we discussed in 2017 have come together at a rapid pace. Traditional benchmarks for measuring how AI performs tasks like recognizing and classifying images or understanding text are becoming obsolete, as researchers launch new methods to measure progress.⁴

As I said then, AI is a foundational technology that drives products and services that people use every day. It also raises important policy issues, which are core to our work at BSA. We undertook a year-long project to work with member companies to develop the BSA Framework to Build Trust in AI,⁵ which was released in 2021 and is designed to help organizations mitigate the potential for unintended bias in AI systems. Built on a vast body of research, the BSA

¹ I am a member of the National AI Advisory Committee, but I am testifying in my capacity as CEO of BSA.

² BSA's members include: Adobe, Alteryx, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, Databricks, DocuSign, Dropbox, Elastic, Graphisoft, IBM, Informatica, Juniper Networks, Kyndryl, MathWorks, Microsoft, Okta, Oracle, Palo Alto Networks, Prokon, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Unity Technologies, Inc., Workday, Zendesk, and Zoom Video Communications, Inc.

³ Testimony of Victoria Espinel, Hearing on Digital Decision-Making: The Building Blocks of Machine Learning and Artificial Intelligence, Before the Senate Committee on Commerce, Science, & Transportation Subcommittee on Communications, Technology, Innovation, and the Internet (Dec. 12, 2017), *available at* <https://www.bsa.org/files/policy-filings/12122017BSAAITestimony.pdf>.

⁴ AI Index Steering Committee at the Institute for Human-Centered AI at Stanford University, The AI Index 2023 Annual Report (April 2023), page 114, *available at* https://aiindex.stanford.edu/wp-content/uploads/2023/04/HAI_AI-Index-Report_2023.pdf (AI Index 2023).

⁵ See *Confronting Bias: BSA's Framework to Build Trust in AI* (June 2021), *available at* <https://www.bsa.org/reports/confronting-bias-bsas-framework-to-build-trust-in-ai>. BSA has testified before the United States Congress and the European Parliament on the Framework.

Framework sets out a lifecycle-based approach for performing impact assessments to identify risks and highlights best practices for mitigating those risks.

Best practice documents, like BSA's Framework, have moved the policy debate forward, but I am here today to say that they are not enough, and legislation is needed. Thoughtful AI legislation will benefit the US economy by creating new guardrails that build trust in the use of AI technologies. It will protect consumers by ensuring AI developers and deployers take required steps to mitigate risks; and it will set the US as a leader in the global debate about the right way to regulate AI. Fortunately, Congress can leverage tools that already exist to create legislation that requires companies to identify and mitigate risks associated with high-risk AI systems.

I. Congress Should Act Now to Adopt Meaningful AI Legislation

Congress should not wait to enact legislation that creates new obligations for companies that develop and use AI in high-risk ways.

Legislation will not only protect consumers from real risks of harm, but will also create trust in AI technologies that will benefit the economy broadly. Consumers and businesses already rely on a wide range of AI-powered services, but they will only continue to adopt new AI technologies if they trust that those products and services are developed and deployed responsibly. Because companies of all sizes and in all industry sectors can benefit from AI, thoughtful AI legislation is important to promoting the United States economy. Countries that support the broad adoption of AI will see the greatest growth in jobs and economic prosperity.

To enact legislation, Congress should take advantage of the considerable work that governmental organizations, civil society advocates, and industry groups have put into identifying the risks of using AI in different contexts and the concrete steps organizations can take to mitigate those risks. Although these proposals have important differences, they collectively form a basis for action. For example, there are fundamental objectives on which everyone should agree: AI, in any form, should not be used to commit illegal acts. It should not be used to compromise privacy, facilitate cyberattacks, exacerbate discrimination, or create physical harm. AI that is developed and deployed responsibly, that improves our lives and makes us safer, should flourish. But Congress should go farther, to create new rules for companies that develop and deploy high-risk AI systems.

By enacting legislation, Congress will also ensure the United States is not just a leader in developing AI technology but is a leading voice in the global debate about regulating AI. The window to lead conversations about AI regulation is rapidly closing, as other governments are moving to shape the rules that will govern AI's future. By the end of this year, the European Union is expected to finalize its AI Act, which is on pace to be the most comprehensive AI law enacted. Japan is leading a G7 effort to establish common standards for AI governance. In November, the United Kingdom will host a global AI summit, focused on the safe and responsible deployment of AI. These governments are not alone. Australia, Brazil, Canada, China, Korea, Singapore, and Thailand are among the long list of countries that are examining the right policy approaches for addressing AI risks.⁶

⁶ See Australia: BSA Comments on Supporting Safe and Responsible AI Innovation in Australia (July 26, 2023), *available at* <https://www.bsa.org/files/policy-filings/07262023safeai.pdf>; Brazil: BSA Recommendations to the Jurists Committee on AI Regulation in Brazil (May 30, 2022), *available at* <https://www.bsa.org/files/policy-filings/en05302022airegbrazil.pdf>; Korea: BSA Comments on Bill on Fostering AI Industry and Securing Trust (Feb. 13, 2023), *available at* <https://www.bsa.org/files/policy-filings/en02132023kraitrust.pdf>; Singapore: BSA Comments on

The United States has recognized the enormous benefits of working with other countries on AI governance issues, and is already participating in a range of global efforts. We support the United States' active involvement in these fora, including the US-EU Trade and Technology Council (TTC), the Global Partnership on AI (GPAI), the Organisation for Economic Co-Operation and Development (OECD), and the G7 Hiroshima AI project. But the US voice in those efforts will be stronger if the United States adopts national AI legislation that creates clear guardrails for how companies develop and deploy high-risk AI systems.

My message to Congress is: Do not wait. You can adopt AI legislation now that creates meaningful rules to reduce risks and promote innovation. We urge you to do so.

II. The US Economy Will Benefit from Legislation that Builds Trust in AI

Businesses of all sizes across all industries are looking for ways to adopt AI tools to grow. Indeed, consumers and businesses today already rely on a wide range of services powered by enterprise AI — and the AI our companies create is all around us. While these AI systems may not get the most attention, they are increasingly integrated in the economy and everyday life. For example:

- **Consumers.** As consumers, AI-powered services remind us when we forget to attach a document to an email, or prompt us to open a particular document so that we can pick up where we left off. We expect to auto-complete forms, chat with virtual assistants that are available even when customer service representatives are not, and to use AI-powered apps to identify and reach personal financial goals, or to predict how a sports player will perform in an upcoming game.
- **Businesses.** Businesses across every industry sector are integrating AI into their products and services, including retail stores that use AI to predict when demand for a product will surge so that they can keep shelves stocked, warehouses that rely on AI-powered logistics planning to minimize supply chain shortages, and manufacturers that use AI to detect safety concerns on factory floors. Across industry sectors, businesses of all sizes also use AI to identify and manage common documents, detect fraudulent transactions, and guard against cybersecurity threats.⁷

The economic benefits of AI are not limited to one industry sector or one business model. Instead, the promise that AI may one day impact every industry is quickly turning into a commercial reality and driving digital transformation across sectors. Airlines now use AI systems to more efficiently clean planes between flights; farmers use AI to analyze large amounts of weather information to maximize their harvest; manufacturers use AI to test new prototypes, and construction companies build AI-generated “digital twins” of real-life cities to understand the impacts of a proposed design.

Public Consultation on Proposed Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems (Aug. 30, 2023), *available at* <https://www.bsa.org/files/policy-filings/08302023bsaaiguideines.pdf>; Thailand: BSA Comments on Draft Bill on Promotion and Support of National AI Innovation (Aug. 18, 2023), *available at* <https://www.bsa.org/files/policy-filings/08182023bsanatlai.pdf>.

⁷ For more examples of everyday uses of AI, see BSA, *Everyday AI for Consumers*, *available at* <https://www.bsa.org/files/policy-filings/08012023aicustomers.pdf>, and BSA, *Everyday AI for Businesses*, *available at* <https://www.bsa.org/files/policy-filings/08012023aibusiness.pdf>.

Enterprise software companies are at the leading edge of this transformation, creating the products and services relied on by other companies.⁸ I have included an extensive list of examples from BSA's members in an annex to this testimony, but want to highlight a handful of ways that companies in all industries are using enterprise-powered AI:

- In healthcare, a large pharmacy chain uses an AI-powered platform to forecast demand and redistribute medications across thousands of store locations and to deliver near real-time insights and recommendations for pharmacists to provide more personalized advice to patients. This helps managers understand the supply chain, store labor and productivity, patient vaccine scheduling, and prescription pickup processes.
- In manufacturing, a car maker used AI-based generative technology to redesign a seat bracket, which secures seat belt fasteners to seats and seats to floors, that is 40% lighter and 20% stronger than the previous iteration. Changes like these can help reduce the amount of material needed to build a car and make vehicles more fuel efficient.
- In agriculture, the research division of an enterprise software provider partnered with a climate risk company to develop AI models capable of providing more accurate long-range weather predictions. Traditional weather forecasting methods can provide accurate predictions for a seven-day window. By leveraging AI, the researchers are developing new forecasting models to provide accurate predictions of weather trends two- to six-weeks out from a given date. By providing reliable extended forecasts, these tools will help water managers predict snowpack and water availability for irrigation, hydropower, and other critical agricultural and environmental uses.

The importance of AI to the US economy is inescapable. Innovations in electricity and personal computers created investment booms of as much as 2% of US GDP as those technologies were adopted into the broader economy.⁹ Economic forecasts estimate that AI could have an even bigger impact on GDP. By 2025, investment in AI is expected to approach \$100 billion in the United States and \$200 billion globally.¹⁰ Generative AI alone could add up to \$4.4 trillion of value to the global economy every year.¹¹

Why does this matter for policymakers?

To realize these economic benefits, consumers and businesses must trust that AI is developed and deployed responsibly. While the adoption of AI can unquestionably be a force for good, it can also create real risks if not developed and deployed responsibly. Setting thoughtful rules for AI is therefore central to the vitality of our economy. Industries of all kinds

⁸ BSA's policy and educational resources on AI are available at <https://ai.bsa.org/>. For additional information about AI's adoption across industry sectors see BSA, Artificial Intelligence in Every Sector, *available at* <https://www.bsa.org/files/policy-filings/06132022bsaaieverysector.pdf>.

⁹ Goldman Sachs, AI Investment Forecast to Approach \$200 Billion Globally By 2025 (Aug. 1, 2023), *available at* <https://www.goldmansachs.com/intelligence/pages/ai-investment-forecast-to-approach-200-billion-globally-by-2025.html>.

¹⁰ *Id.*

¹¹ McKinsey & Company, The Economic Potential of Generative AI (June 2023), *available at* https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier?gclid=Cj0KCQjwI8anBhCFARIsAKbbpyTeLnz0c6i4X2UTnmWdO1KGQnE1mUR8ErJSrM0eMnWDxgfaZukt_L0aAqP7EALw_wcB#.

and businesses of all sizes are looking for ways to use AI to grow, but they will only adopt AI-powered products and services that they trust. Countries that best facilitate responsible and broad-based AI adoption will see the greatest economic and job growth in the coming years.

III. Congress Should Enact Legislation That Builds Trust in AI

Now is the time for Congress to adopt thoughtful legislation that addresses known risks of AI systems.

There will continue to be a range of significant and evolving policy debates around AI. But there's no need to wait to pass legislation that creates meaningful guardrails against the AI risks that exist today. You can — and should — build on existing regulatory efforts by setting rules across the economy to address concerns about the potential for bias when AI systems are used in high-risk ways.

A. Legislation Should Create New Obligations for High-Risk Uses of AI

AI legislation can build on work that's already been done.

To anticipate and address potential harms, developers of AI used in high-risk situations should understand the practical implications of its use, including through the use of impact assessments. Those who train AI systems for high-risk uses, or are deploying AI systems in high-risk contexts, should also be required to understand how the tools they create or deploy might result in unintended outcomes and act to mitigate those risks.

We urge Congress to adopt legislation that creates meaningful guardrails for high-risk uses of AI. That legislation should require companies to:

- (1) establish risk management programs to identify and mitigate risks across AI systems;
- (2) conduct annual impact assessments for high-risk uses of AI, and
- (3) publicly certify that they have met these requirements.

Congress can build on tools that already exist today — and require companies to adopt those tools to identify and mitigate risks associated with high-risk AI systems. My testimony focuses on two of those tools: risk management programs and impact assessments.

B. Risk Management Programs

The goal of risk management is establishing repeatable processes for identifying and mitigating potential risks that can arise throughout the lifecycle of an AI system. Risk management is particularly important in contexts like AI, privacy, and cybersecurity, where the combination of quickly evolving technologies and highly dynamic threat landscapes can render traditional approaches ineffective. Rather than evaluating a product or service against a static set of requirements that can rapidly become outdated, risk management programs integrate compliance responsibilities into the development process to help identify and mitigate risks throughout a product or service's lifecycle.

Risk management programs have two key elements: (1) a governance framework to support the organizations' risk management functions, and (2) a scalable process for performing impact assessments that identify and mitigate risks of an AI system. The governance framework is critical because it promotes collaboration between an organization's development team and its compliance team at key points in the design, development, and deployment of an AI system.

One way for companies to establish risk management programs is by using the AI Risk Management Framework (AI RMF), which was released earlier this year by the National Institute of Standards and Technology (NIST).¹² Creating the AI RMF was a significant achievement and builds on NIST's work creating frameworks for managing cybersecurity and privacy risks. For example, the NIST Cybersecurity Framework (CSF) is widely used by private and public-sector organizations worldwide; since 2017, it has been mandatory for federal agencies to use the CSF to improve their cybersecurity risk management programs.¹³ Like the CSF, the AI RMF is a voluntary tool. It helps organizations incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products.

The AI RMF can help companies establish risk management programs that guard against a range of potential AI-related harms. It focuses on corporate practices around four functions: govern, map, measure, and manage. The AI RMF is designed to be usable by large and small organizations alike and can be applied in varied circumstances, such as for procurement purposes, for organizations with existing governance programs, and for organizations just beginning to think about risk management. It also identifies several indicia of trustworthy AI, which include privacy, explainability, and fairness, and incorporates an assessment of these characteristics as part of the measure function. Importantly, the AI RMF acknowledges that tradeoffs among trustworthiness characteristics may exist. For example, the use of privacy-enhancing technologies may decrease accuracy, which could affect decisions about fairness in certain domains, and there may also be tradeoffs between interpretability and privacy.

The AI RMF encourages:

- Consultation with diverse stakeholders;
- Establishing processes to identify, assess, and mitigate risks;
- Defining individual roles and responsibilities to people throughout an organization;
- Identifying metrics for evaluation;
- Evaluating fairness and bias;
- Maintaining post-deployment feedback mechanisms; and
- Establishing incident response plans.

Ultimately, effective AI risk management programs should be underpinned by a governance framework that establishes the policies, processes, and personnel that will be used to identify, mitigate, and document risks throughout the system's lifecycle. The purpose of such a governance framework is to promote understanding across organizational units — including product development, compliance, marketing, sales, and senior management — about each entity's role and responsibilities for promoting effective risk management during the design, development, and deployment of AI systems.

C. Impact Assessments

Impact assessments have three purposes: (1) identifying potential risks that an AI system may pose, (2) quantifying the degree of potential harms the system could generate, and (3)

¹² NIST AI Risk Management Framework, *available at* <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

¹³ See NIST, Cybersecurity Framework, Questions and Answers, (discussing federal agency use of the NIST CSF), *available at* <https://www.nist.gov/cyberframework/frequently-asked-questions/framework-basics#agency>.

documenting steps taken to mitigate those risks.¹⁴ As noted earlier, performing impact assessments is a key part of creating a meaningful risk management program.

Impact assessments are already widely used in a range of other fields, including data protection, as an accountability mechanism that demonstrates a system has been designed in a manner that accounts for the potential risks it may pose to the public. Because impact assessments already exist today, they can be readily adapted to help companies identify and mitigate AI-related risks. For example, three state privacy laws already require companies to conduct impact assessment for specific activities, including processing sensitive personal data, engaging in targeted advertising, or selling personal data; seven more state privacy laws will soon do so.¹⁵ Globally, privacy and data protection laws worldwide use impact assessments as a tool for improving accountability.¹⁶

The use of impact assessments can be readily expanded to require assessments for high-risk AI systems. Conducting impact assessments can help companies identify and mitigate risks, including risks of unintended bias. In our view, when AI is used in ways that could adversely impact civil rights or access to important life opportunities, the public should be assured that such systems have been thoroughly vetted and will be continuously monitored to account for the risks associated with unintended bias. Requiring impact assessments for companies that develop and deploy these high-risk systems is an important way to do that.

D. Legislation Should Focus on High-Risk Uses and Create Role-Appropriate Requirements.

Legislation that requires risk management programs and impact assessment will create new safeguards for high-risk AI systems. Any legislation incorporating these requirements should:

- ***Focus on high-risk AI uses.*** Legislation should place guardrails around high-risk uses of AI, namely AI systems that determine an individual's eligibility for housing, employment, credit, education, access to physical places of public accommodation, healthcare, or insurance. These systems have the potential to affect important life opportunities — and are a key area for policymakers to address. In contrast, many everyday uses of AI

¹⁴ See BSA, Impact Assessments: A Key Part of AI Accountability, *available at* <https://www.bsa.org/files/policy-filings/08012023impactassess.pdf>.

¹⁵ Colorado, Connecticut, and Virginia already impose these requirements. See Colorado Privacy Act, Colo. Rev. Stat. Tit. 6, Art. 1, Pt. 13 §§ 6-1-1301–6-1-1313; Connecticut Data Privacy Act Conn. Gen. Stat. Tit. 42, Ch. 743jj, Sec. 42-515-525; Virginia Consumer Data Protection Act; Va. Code Tit. 59.1, Ch. 53, § 59.1-575-585. Recently passed state privacy laws in Florida, Indiana, Montana, Oregon, Tennessee, and Texas will also require impact assessments for certain activities.

¹⁶ Data protection impact assessments are an established accountability tool under privacy laws worldwide. Under the European Union's General Data Protection Regulation (GDPR), impact assessments are required for activities likely to result in a "high risk" to the rights and freedoms of individuals. Brazil's General Data Protection Law (LGPD) also allows the national data protection agency to require a company to prepare a data protection assessment, and in Singapore organizations are required to conduct impact assessments if they process data for certain purposes. There is also significant regulatory guidance on how organizations can conduct impact assessments to identify and mitigate privacy risks. See, e.g., Office of the Australian Information Commissioner, Guide to Undertaking Privacy Impact Assessments (Sept. 2, 2021), *available at* <https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/privacy-impact-assessments/guide-to-undertaking-privacy-impact-assessments>; European Data Protection Board, Guidelines on Data Protection Impact Assessment (Oct. 13, 2017), *available at* <https://ec.europa.eu/newsroom/article29/items/611236>; Singapore Personal Data Protection Commission, Guide to Data Protection Impact Assessments, (Sept. 14, 2021) *available at* <https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Other-Guides/DPIA/Guide-to-Data-Protection-Impact-Assessments-14-Sep-2021.pdf>.

present few risks to individuals and create significant benefits, like helping organize digital files, auto-populate common forms for later human review, or predict the fonts to use in a template document. Indeed, Americans have reported being most excited about the potential for AI to perform household chores and conduct repetitive workplace tasks, while expressing significant concerns about the use of AI to make important life decisions for people.¹⁷ Legislation should address these high-risk uses.

- *Recognize the different roles of companies that develop AI systems and companies that deploy AI systems.* Legislation should recognize the different roles of companies that create and use AI systems. Developers are the companies that design, code, or produce an AI system, such as a software company that develops an AI system for speech recognition. In contrast, deployers are the companies that use an AI system, such as a bank that uses an AI system to make loan determinations. Legislation should recognize the different roles of developers and deployers, because these two types of companies will have access to different types of information and will be able to take different actions to mitigate risks.

For example, the developer of an AI system is well positioned to describe features of the data used to train that system, but it generally will not have insight to how the system is used after it is purchased by another company and deployed. A company that deploys an AI system is well positioned to understand how the system is actually being used, what type of human oversight is in place, and whether there are complaints about how the system works in practice. Legislation should recognize these different roles, so that the appropriate company in the real-world AI supply chain can identify and mitigate risks. For the same reasons, this kind of distinction is considered best practice in privacy and security legislation around the world.¹⁸

Legislation that leverages these existing tools can create meaningful safeguards for the development and deployment of AI in high-risk uses.

E. Benefits of Legislation

If Congress passes a law that requires risk management programs and impact assessments, it can require companies to identify and mitigate across the lifecycle of an AI system. These requirements will have real benefits to consumers, by making companies evaluate and address potential risks that can arise across the lifecycle of an AI system. These risks can arise in contexts including:

- *Problem Formulation.* The initial step in building an AI system is often referred to as “problem formulation.” It involves the identification and specification of the “problem” the system is intended to address, an initial mapping of how the model will achieve that objective, and the identification of a “target variable” the system will be used to predict. Because many AI systems are designed to make predictions about attributes that are not directly measurable, data scientists must often identify variables that can be used as proxies for the quality or outcome it is intended to predict.

¹⁷ AI Index 2023, page 329.

¹⁸ See BSA, AI Developers and Deployers: An Important Distinction, *available at* <https://www.bsa.org/files/policy-filings/03162023aidevdep.pdf>.

While the use of proxy target variables can be entirely reasonable, the assumptions underlying the choice of proxies must be closely scrutinized to ensure that it does not introduce unintended bias to the system. The risk that can arise in problem formulation is exemplified by a study of a widely used healthcare algorithm that hospitals rely on to identify patients in need of urgent care. The research team concluded that the algorithm was systematically assigning lower risk scores to black patients compared to similarly sick white counterparts because it relied on data about historical healthcare costs as proxy for predicting a patient’s future healthcare needs. Unfortunately, because black patients have historically had less access to healthcare, the reliance of spending data painted an inaccurate picture and led to dangerously biased outcomes.¹⁹

- **Training Data.** The data used to train an AI system is a second major vector for bias. If the data used to train a system is misrepresentative of the population in which it will be used, there is a risk the system will perform less effectively on communities that may be underrepresented in the training data. Likewise, reliance on data that itself may be the product of institutional or historical biases can entrench those inequities in an AI model.
- **Labeling.** The process of “labelling” training data can also introduce bias. Many AI systems require training data to be “labeled” so that the learning algorithm can identify patterns and correlations that can be used to classify future data inputs. Because the process of labeling the data can involve subjective decisions, there is the potential for introducing unintended bias into the training data.
- **Deployment.** Even a system thoroughly vetted during development can begin to exhibit bias after it is deployed. AI systems are trained on data that represents a static moment in time and that filters out “noise” that could undermine the model’s ability to make consistent and accurate predictions. Upon deployment in the real world, AI systems inevitably encounter conditions that differ from those in the development and testing environment. Because the real world changes over time, the snapshot in time that a model represents may naturally become less accurate. If the input data for a deployed AI system differs materially from its training data, there is a risk that the system could “drift” and that the performance of the model could be undermined. For instance, if an AI system is designed (and tested) for use in a specific country, the system may not perform well if it is deployed in a country with radically different demographics. Bias can also arise if an AI system is deployed into an environment that differs significantly from the conditions for which it was designed or for purposes that are inconsistent with its intended use.

Congress should act now to require companies that develop or deploy high-risk AI to adopt risk management programs and to conduct impact assessments. These requirements will protect consumers and create meaningful rules that reduce risks and promote innovation.

* * *

We appreciate Congress’s leadership on the important policy issues raised by AI. We are ready to help as you craft and pass legislation. Thank you and I look forward to your questions.

¹⁹ Heidi Ledford, Millions of Black People Affected by Racial Bias in Health-Care Algorithms, *Nature* (Oct. 24, 2019), available at <https://www.nature.com/articles/d41586-019-03228-6>.

Annex: AI In Every Sector

Improving Healthcare and Quality of Life

The rapid digitalization of health information has created tremendous opportunities for AI to transform how clinicians care for patients, how consumers manage their health, and how researchers discover breakthroughs in the treatment and prevention of diseases.

Helping Pharmacies Redistribute Medication and Provide Personalized Advice to Patients

Walgreens uses the Databricks Lakehouse platform to run an intelligent data platform incorporating AI to forecast demand and redistribute medications across Walgreens' network of nearly 9,000 pharmacies, while delivering near real-time insights and recommendations for pharmacists to help provide more personalized advice to patients. This integrated AI-driven platform allows Walgreens' different data teams to work better together, create smarter algorithms and generate new types of reporting to help managers understand the supply chain, store labor and productivity, patient vaccine scheduling, and prescription pickup processes.

Advancing Accessibility

For people with visual impairments, AI is turning the visual world into an audible experience. Microsoft's Seeing AI app helps people who are blind or visually impaired recognize objects, people, and text via a phone or tablet's camera and describes what it recognizes to the user. With this new layer of information, users can navigate the world more independently.

Strengthening Security

Although data security is core to the management of most organizations, cyber threats continue to evolve at a breakneck pace. AI helps organizations stay a step ahead of hackers by predicting potential attacks, mitigating attacks in real-time, managing access to resources, and encrypting sensitive data.

Enabling Fast Action Against Security Threats

Palo Alto Networks' AI-driven Security Operations Center automation engine, XSIAM, is delivering never-before-seen cybersecurity outcomes. The company's own instance of this tool ingests 36 billion events every day from across all network layers and attack surfaces and triages just 8 of those for human analysis. This empowers their most precious resources — people — to focus on the most sophisticated attacks that uniquely require human analysis. Importantly, this AI-driven tool has reduced overall Mean Time to Detection (MTTD) to 10 seconds and Mean Time to Response (MTTR) to one minute for high priority incidents. This more resilient and automated cyber future would not be possible without AI.

Protecting Business Transactions

Splunk is helping financial institutions to leverage AI and data analytics to strengthen their cybersecurity and their ability to serve customers. For example, consumer report and risk scoring provider TransUnion uses data analytics and machine learning capabilities provided by Splunk to monitor customer traffic and transactions. TransUnion monitors and manages customer traffic to its website and detects when unusual activity takes place so it can alert customers about security concerns and ensure seamless customer experiences.

Building 21st Century Infrastructure

Whether it's creating smarter and safer cities by integrating sensors in bridges and highways to monitor their safety or increasing efficiency by cutting travel time and fuel expenses, AI plays an instrumental role in creating an infrastructure designed for the 21st century.

Optimizing Manufacturing

Generative design tools can optimize the manufacturing process to reduce waste and improve products. Autodesk teamed up with Michigan-based foundry Aristo Cast to develop an ultralightweight aircraft seat frame. The team used generative design, 3D printing, lattice optimization, and investment casting to ultimately create a seat frame that weighs 56% less than typical current models. For a 615-seat Airbus A380 plane, that would mean saving \$100,000 in fuel per year, as well as more than 140,000 fewer tons of carbon in the atmosphere.

Streamlining Building Projects

Companies are using AI to streamline the building design and construction processes. Bentley Systems has teamed with Hyundai engineering on an AI system that automates design processes for steel and concrete structures, reducing the time needed to create designs and the cost of building a structure.

Monitoring Vehicle Fleets

Oracle's anomaly detection software uses AI to monitor the operation of complex systems and detect potentially concerning incidents. Transportation and logistics company SS Global LLC uses Oracle's software to monitor their fleet of vehicles and get alerts when there are early signs of potential safety issues. By detecting the early onset of tire baldness and air leaks, the system helps SS Global perform predictive maintenance that keeps their fleet safer and more efficient.

Creating New Ways to Learn

AI applications are enabling personalized learning resources for every stage of life, including adaptive learning programs, digital tutoring, curriculum recommendations, and more. There are more digital resources available to instructors and students than ever before, and AI is affording them the ability to access relevant tools quickly and easily.

Enriching Math Education

Educators are using IBM's Teacher Advisor With Watson AI to access the math resources they need in seconds, including proven lesson plans, activities, standards information, and teaching strategies for students with varying degrees of preparation and ability. This can save valuable time for teachers throughout the school year.

Tailoring Workplace Learning

Employers are using Workday Learning, an application that uses machine learning to personalize workplace learning for individuals, to recommend professional development content and courses based on employee position, tenure at the company, interactions with the content, and other factors. This helps companies adjust learning strategies and programming to ensure employees learn new skills, continue to grow in their roles, and prepare for what's ahead.

Enhancing the Customer Experience

For businesses with large customer bases that are processing a high volume of purchases — such as banks, restaurant chains, and large retailers — analyzing the massive amount of data collected every day is impossible without the computing and predictive power of AI. By using machine learning tools, businesses across a wide range of industries can analyze customer preferences and their own business performance to improve end-user experiences and increase efficiencies. Software also helps businesses generate optimal product designs by using data to produce and analyze far more iterations than humans alone could create.

Customizing Care Experiences

Powered by Salesforce AI technology, Eli Lilly has reimagined patient care with its Patient Connect Platform app. The app helps customers learn to use products, access information about their medications, and record how well they are feeling. The desktop and mobile apps also allow patients to consult with a healthcare concierge — a specialist who provides one-on-one support to guide patients toward beneficial health outcomes.

Improving Customer Service Experience

Zendesk is using AI to improve the customer service experience for both customers and the agents that interact with them. Using Zendesk's intelligent triage functionality, a company can automatically detect a customer's intent (for example, whether a customer is making a return or checking on shipping status), the language the customer is using, and the customer's overall sentiment so that the inquiry can be quickly routed to the best agent for the job. Several of Zendesk's business-to-consumer customers are using this Zendesk AI feature to automatically classify and route incoming tickets to the right agents at the right time, which has resulted in higher customer satisfaction and more one-touch tickets.

Scaling Community Impact

Twilio provides AI chatbot services to help businesses interact with customers. The United Way Worldwide worked with Twilio to help scale and route inbound calls and texts to more than 200 agencies nationwide that use their 211 system to help people locate essential needs like housing, financial assistance, food, childcare, transportation, and more. Using the AI-assisted interactive voice response menu built with Twilio Autopilot, the United Way and Twilio built a system that enables a caller to access a single 1-800 number or be transferred by their local 211 to access assistance. The result is a centralized system that efficiently reduces the call volume nationwide but increases the time staffers are able to devote to mission critical calls.

Improving Business Operations

AI is helping to streamline business operations and increase productivity.

Enhancing Business Functions

SAP provides chatbot solutions that are seamlessly integrated into other business functions, giving customers, partners, and employees a bird's-eye view of business operations. For example, SAP provides software services to Hewlett Packard Enterprise Company, including an AI-based chatbot system that can reference serial numbers, packing slips, and shipment dates drawn from cloud services, thereby getting the right information to the right people at the right time.

Improving Contract Analysis

DocuSign has been helping organizations use AI-based technologies including natural language processing and rules-based logic to manage and analyze agreements for several years now. Using AI-powered contract analysis can increase productivity in the contract process by helping to speed up contract reviews, increase contract visibility, and identify opportunities and risks.

Empowering Creativity

AI and machine learning within Adobe's Creative Cloud tools help artists, photographers, designers, and content creators around the world handle the time-consuming aspects of their work that can easily be automated, so they have more time to be creative. From removing unwanted objects like mics and logos from videos in Adobe After Effects, to colorizing black-and-white photos in just a few clicks in Adobe Photoshop, to painting with digital brushes that look, feel, and act like the real thing in Adobe Fresco, Adobe's AI and machine learning features empower creators to focus their energy on what they love — ideating, experimenting, and creating.

Helping in Times of Crisis

In times of humanitarian crises, fast response is essential. Researchers are developing ways to use AI to help first responders in the critical hours and days after a natural disaster, and to track pathogens that could lead to outbreaks of disease and mitigate the spread.

Navigating the COVID-19 Pandemic

Siemens' Dynamic VAV Optimization (DVO) is a software solution for building management systems that uses machine learning and AI to configure HVAC controls according to a building's priorities, whether that's minimizing virus transmission or minimizing energy consumption. In direct response to the challenges of the pandemic, DVO was launched with a new operating Defense Mode in late 2020 to reduce the risk of viral spread in indoor spaces. DVO adjusts ventilation, temperature, and humidity conditions to minimize risk of viral spread indoors while also maximizing energy efficiency.

Enriching Our Lives

Leveling Up Gaming and Entertainment

AI can be used to create sophisticated 3-D environments and train autonomous characters in our favorite games and movies. Unity's AI products are used to develop video games, animations, and other detailed virtual environments. By training computer-based characters in Unity's software, game designers can create more realistic environments that capture a player's imagination and enhance the gaming experience.