Applications of Artificial Intelligence (AI) in Government and Industry

PAUL H. O'NEILL SCHOOL OF PUBLIC AND ENVIRONMENTAL AFFAIRS

PAUL H. O'NEILL
SCHOOL OF PUBLIC AND ENVIRONMENTAL AFFAIRS

Applications of Artificial Intelligence (AI) in Government and Industry

A workshop organized by the National Academies Innovation Policy Forum in cooperation with the Manufacturing Policy Initiative at the Indiana University Paul H. O'Neill School of Public and Environmental Affairs

September 26, 2019
National Academy of Sciences in Washington, D.C.

TABLE OF CONTENTS

INTRODUCTION

Artificial Intelligence (AI) is a transformational technology. For example, the manufacturing sector is using AI to serve a wide variety of purposes: workforce training, production process improvement, quality control, predictive maintenance, supply chain optimization, distribution (both inside and outside the fence), product design, and new and improved products. The federal government is using AI to prevent terrorism, code surveys of occupational illnesses and injuries, answer queries from citizens and migrants on immigration matters, help farmers in developing countries forecast weather patterns more accurately, and help small businesses more easily navigate the federal procurement process.

On Thursday, September 26, 2019, a one-day workshop was held to highlight AI-enabled innovation. The workshop, "Applications of Artificial Intelligence in Government and Industry," was organized by the National Academies Innovation Policy Forum in cooperation with the Manufacturing Policy Initiative (MPI) at the Paul H. O'Neill School of Public and Environmental Affairs at Indiana University. It was held at the National Academy of Sciences in Washington, D.C.

The event featured presentations by government and industry officials detailing their experience in developing and applying AI. Panel discussions addressed challenges faced by government and industry in applying AI. The goal was to inform perspectives of AI and identify emerging themes. This document serves as a summary of the workshop.

The workshop consisted of four panels. Each panel was led by a moderator and included presenters or discussants. The first two panels addressed AI research and applications in government. The third and fourth panels addressed AI applications in industry.

PANEL 1:
GOVERNMENT APPLICATIONS

Moderator **Heather Evans** (National Institute of Standards and Technology) noted the recent release of the federal government's plan for federal engagement in developing technical standards and related tools for AI. This plan will help ensure U.S. leadership for AI in accordance with Executive Order 13859, issued in February 2019. She briefly described the range of activities within NIST with respect to AI (emphasizing research and interagency engagement) before framing the panel presentations in terms of research critical to governmental use of AI.

Cindy Goldberg (IBM) focused her remarks on research efforts to ensure trust in AI. She illustrated an increase in the number and scope of data sets. She distinguished between narrow AI (designed to complete a highly defined task) and broad AI (designed to execute tasks across a variety of different fields), and concluded that narrow AI is working well: tasks are able to be programmed well. She raised the question: How will the applications shift as society moves from narrow AI to general AI? She observed that the world now relies on AI applications. IBM, for example, works with manufacturers to employ and apply AI to optimize supply chains. Ensuring cybersecurity is a major focus—preserving privacy while allowing data to be shared across supply chains. She observed that for AI to go from narrow to general, it must be trusted. Researchers must seek to build trust as capacity and sophistication grows. Transfer learning—the building upon existing data—offers considerable savings in terms of time and effort. We must add reasoning—something simple for humans—to AI models. She observed that hardware currently used to apply AI is unsustainable over the longer term given the expectations of AI growth and the energy required to facilitate AI computations. IBM is spending much effort to ensure that hardware can support the next generation of AI.

John Reeder (Navy Information Warfare Center Pacific) provided an overview of the Defense Advanced Research Projects Agency (DARPA) and offered an overview of DARPA projects related to explainable AI. He observed that the need for explainability is apparent in many "narrow AI" applications: A user of AI does not necessarily know why AI reached a particular conclusion or made a particular determination or why some other conclusion or outcome was rejected. Explainable AI helps the user understand why/why not, the circumstances under which AI fails or succeeds, the circumstances under which AI decisions or determinations can be trusted, and why AI errs when it errs. The current DARPA project is focused on explainability in terms of (1) analytics and (2) autonomy. It aims to maintain a high level of AI performance while enabling analysts and operators to understand, trust, and manage AI systems. Reeder provided several examples drawn from published literature and conducted at various academic institutions across the United States (e.g., RISE—randomized input sampling for explanation). Thus far, DARPA has identified several promising approaches to explainability in terms of analytics, autonomy, or both. Initial results indicate that users preferred explanations, and correct explanations engendered trust.

Ricardo Arteaga's presentation focused on two inventions he led at NASA: (1) a collision avoidance system for supersonic aircraft (based on Automatic Dependent Surveillance Broadcast—or ADS-B, which is replacing radar for tracking aircraft worldwide) and (2) AI-embedded drone technology to improve search-and-rescue operations. For the first, AI (specifically, neural-network trajectory predictions based on training data with F-18 aircraft) was used to reposition an aircraft traveling at supersonic speed to avoid an impending collision with another aircraft. In the second example, AI was used, in conjunction with an unmanned aircraft system, or UAS (i.e., a drone), to video document a plane crash and rescue site, determine the boundary of the scene to ensure the security of the crash site, and stream the video to a central location. The machine vision technology—based on you-only-look-once (YOLO) object detection using Convolutional Neural Networks—correctly detected and classified objects (e.g., the "black box" of a downed commercial aircraft) in real time with an acceptable level of precision.

PANEL 2: CHALLENGES TO APPLYING AI IN GOVERNMENT

Moderator **Dan Chenok** (IBM) provided a brief description of IBM's Center for Business and Government, which has collaborated with other organizations (the Partnership for Public Progress) to highlight use cases of AI applications in government. He identified major challenges for the federal government in applying AI, including timely investing in technology, and acquiring and maintaining a workforce skilled in AI. He introduced the panelists, who represent a variety of federal agencies and missions.

Gil Alterovitz (Veteran's Administration) emphasized how collaborative efforts by government and industry can be fostered to improve health outcomes. As an example, he described a recent publication in *Nature* where AI was used to predict acute kidney injury 48 hours in advance, which enables time to take preemptive action. He explained how the government can design data formats to get the most out of AI, by creating an ecosystem of datasets, platforms, and tools to enable standardized AI piloting, deployment, and testing. As an example, he described the use of AI for novel therapeutics accessibility (NTA): an ecosystem of three datasets built for Health Tech Sprint around clinical trials in structured, machine-readable information. This project involved resources and personnel from a dozen agencies. Participation included federal agencies but also many private sector organizations and nonprofits. He then described an AI-able data ecosystem framework, which would allow bidirectional data and results sharing between government and industry. This will create a voluntary framework for incentivizing federal agencies to release data that is useful and usable by the private sector, etc. It will also incentivize industry to make AI/data results usable by others, including federal agencies, so they can leverage industry tools. He described parallel pathways for potential partnerships (PPPP): the result of which would be signed research agreements around a particular AI/ML problem/task with potentially mutual expectations. VA aims to create a system where smaller data sets are shared, and AI models can be created by industry.

Erwin Gianchandani (National Science Foundation) described the breadth and depth of activities within NSF that support AI and AI applications. He described the NSF mission as supporting fundamental research including machine learning and innovation. NSF is the key extramural funder of research including research in AI and in all areas of discovery. Some examples include a recently launched $5-million partnership effort on collaborative research, a recently announced collaboration with Amazon on issues of fairness and AI, and a new project with DARPA on real-time machine learning to advance hardware systems. Individuals don't necessarily need to understand how to code but rather understand the technology that is surrounding them from a computational basis. NSF is investing in K-12 education to incentivize a broad and diverse group of students who are interested in AI. Finally, NSF is also assisting with research accessibility to infrastructure and cloud computing resources.

Diane Furchtgott-Roth (Department of Transportation) provided an overview of the breadth of research activities across the Department of Transportation including those that support applications of AI. She mentioned making available unlicensed spectrum to create a hybrid communications infrastructure to address a wide range of transportation needs. She mentioned the safety benefits of autonomous vehicles. She mentioned the need to thwart GPS spoofing. She mentioned that DOT data could be used in numerous ways—some known and many unknown—that benefit society. For example, AI could utilize data on the behavior of drivers to help to diagnose Alzheimer's disease.

Brian Valentine (Department of Energy) described three major technology programs within the DOE: sciences (physical, mathematical, and biological), defense (responsibility for entire defense nuclear enterprise), and applied programs (electricity grid programs, nuclear, fossil, renewable R&D). He observed that DOE has been at the forefront of applying AI (for 60 years), which began in the National Labs (Sandia and Los Alamos). AI has been extensively applied to research across DOE—but only recently applied in cooperation with external partners. The major challenge is to coordinate all of the AI activities within the Department. To meet this challenge, DOE recently established an AI Technology Office. DOE capacity for AI applications lies largely within the national laboratory system. Among the recent AI initiatives within DOE (extracting knowledge from data) include electric grid resiliency, biomedical (neurosciences), partnerships with VA on veterans health, cybersecurity partnerships with various federal agencies, seismic research (to predict earthquakes), cancer research (partnership with the national Cancer Institute), and materials research.

During the Q&A, the observation was made that cybersecurity is a double-edged sword for AI progress: hackers use it to create problems and organizations are using it to identify and thwart the efforts of hackers.

PANEL 3: INDUSTRY APPLICATIONS

Moderator **Andre Gudger** (Eccalon) emphasized the importance of understanding how AI is being used in a manufacturing setting. The panel includes three individuals who have a breadth of experience in this subject, and will address both the range of AI applications but also the challenges facing companies that employ AI (machine learning).

Alp Kucukelbir's (Fero Labs) presentation focused on three lessons learned from applying machine learning to industrial process improvement. He explained that steel mills aim to minimize the scrap rate utilizing 12,000 sensors across its production plant. Typically, process improvement starts with identifying a few factors and gathering limited data, followed by application of traditional Six Sigma techniques. Fero takes observational data to build interpretable, expertise-based, and safe machine learning models to improve the production process. By collecting all the data, Fero can analyze all of it without the bias that comes from selectively choosing to look at only a subset of data. This approach allows manufacturers to adaptively adjust production for dynamic recipes, maximize uptime, minimize faults by discovering new product configurations, identify and address the root cause of complex quality issues, and reduce emissions with existing hardware. Among the lessons learned from applying AI in industrial settings, interpretability matters (leverage domain expertise of factory personnel), humans are part of the loop (expertise-based machine learning based on custom compiled models), and statistics is still cool (factories must know when to trust AI predictions through statistical guarantees). According to Dr. Kucukelbir, the meta lesson is that to create value, you must change human behavior—and to change human behavior, you must empower users.

Jake Bartolone (Uptake) presented use cases of AI in industry. He framed his presentation by relating the growth of the Internet of Things (IoT) to AI and machine learning, and that the IoT is growing most rapidly in factories. Certain portions of manufacturing—such as fleet management and on-time, in-full shipping are expecting strong growth within the next five years. He provided two use cases of AI in industry: failure prediction (utilizing AI and machine learning to detect impending equipment failure) and transfer learning (applying learning from related but not identical tasks to make predictions on previously unseen data sets). In the case of failure prediction from a railroad locomotive, engine coolant and oil temperature led to replacement of a coolant fan that prevented engine shutdown. Savings are estimated at $20K per locomotive per year. In the case of transfer learning, neural networks can be supplemented by information from previous tasks to speed parameter search. Bayesian transfer, based on Bayesian learning, creates theoretical bounds on convergence and prediction errors. The value of these AI applications is significant. An average Class I railroad could realize annual savings of $86 million if 10 percent of unplanned maintenance was converted to planned maintenance. In mining, reducing tire repair on mining equipment can reduce the cost of daily downtime ($1.5 million per day). Dr. Bartolone listed challenges for applying AI in industry, including recognition that not all anomalies recognized by AI are meaningful. For example, AI identified elevated railroad locomotive engine temperature in specific locations, which turned out to be normal train behavior when going through tunnels. The solution was to add tunnel locations to the AI model. He also emphasized that not all failures are predictable (and not all sensors are reliable), such as a transformer failure due to squirrel-based outage.

Mark Johnson (Clemson University) focused on how AI improves energy management in manufacturing. Johnson showed how energy is used across society. He pointed to the DOE Advanced Manufacturing Office multi-year strategic plan, which linked 14 specific manufacturing technologies (e.g., combined heat and power, additive manufacturing, composite materials) with emerging and crosscutting areas (energy efficient advanced computing) and energy systems (transportation). There is substantial opportunity to maximize energy use through AI: energy savings from the application of IT and through the discovery of new materials that are energy-advantaged. Major challenges include ensuring privacy and security of data throughout a manufacturer's organizational boundaries. There are also challenges in transfer learning: two different researchers working on the same topic but not sharing best practices. However, the biggest challenge is trust—thinking that the status quo is working, thinking that established processes are working. Executives also have a poor general understanding of AI and what it offers.

During the panel discussion, it was noted that institutional knowledge is being lost due to retirement of baby boomers, who have considerable domain knowledge. AI is seen as a solution to mitigate against this loss in institutional knowledge.

PANEL 4:
CHALLENGES TO APPLYING AI IN INDUSTRY

Moderator **Keith Belton** (Indiana University) framed the panel discussion by describing the breadth of AI applications in manufacturing. He noted that while the media fixates on the potential job losses that AI might bring, history and recent experience suggest a net increase in jobs and greater demand for digital skills, including skills in data analytics and AI/machine learning. He observed that the assembled panelists—while all focused on U.S. manufacturing—represent organizations with very different business models and products/service offerings.

Stephen Ezell (ITIF) presented the results of a recent survey of 70 global manufacturers regarding the future of work in the age of AI. The manufacturing sector is entering what some claim to be a new industrial revolution, one that involves digitalization and data analytics driven by AI. As such, manufacturers are undertaking a digitalization maturity journey, a six-step process that begins with computerization and connectivity and ends with predictive capacity (nearly all manufacturers have reached these stages) and adaptability/autonomy (very few manufacturers have reached predictive capacity and none have reached adaptability/autonomy). Given these developments, AI is transforming the manufacturing workforce, slowly right now, but expected to increase significantly within five years. Leading barriers to AI adoption among manufacturers include (1) lack of data resources needed to enable AI solutions, (2) uncertainty on how AI can help solve specific challenges, and (3) lack of sufficient workforce digital skills. Developing and acquiring AI talent is a critical need and involves hiring new talent and reskilling existing workers. There are significant AI skill gaps among executives and front-line workers. Dr. Ezell identified barriers to finding employees with requisite skills and approaches that companies are taking to acquire AI workforce skills (relationship with local educational institutions was the top mechanism). He provided recommendations for business (establish an AI governing coalition for the enterprise) and policy makers (commit to a massive expansion of U.S. AI talent via NSF fellowship program) to ensure manufacturers remain competitive.

Sudhi Bangalore (Stanley Black & Decker) described his experience in leading a global manufacturer into the digital age (also known as "Industry 4.0"). His efforts include connecting factories around the world with data and AI. He noted that collaboration between humans and robots (cobots) result in better performance that simply relying on one to the exclusion of the other. He said that much of his job is persuading the existing workforce to believe in AI and become managers of the machines that can help them accomplish their goals. He gave examples of using data analytics to detect and minimize defective products that are off-spec. He has emphasized upskilling and retraining to develop workers at the same time that AI/machine learning is being adopted.

Ernest Begin (Kaman Corporation) referred to Dr. Ezell's report to frame the application of AI in his company—which produces helicopters and aerospace parts and is similar to most manufacturers in its "digital journey." He emphasize the difficulty in investing heavily in AI and machine learning when the ROI (e.g., for predictive maintenance) cannot be guaranteed. Small and medium enterprises (SMEs)—their suppliers—are particularly challenged to invest in AI and connectivity because they do have the resources to lead in these areas. His company is using robots and cobots. The scarcity of skilled workers is a major impediment to adopting AI/machine learning. They have a focused upskilling program—it poses significant costs as well as benefits. He recommended that manufacturers invest in AI applications where domain knowledge is greatest and purchase AI solutions from vendors in other cases.

Young M. Lee provided a brief description of Johnson Controls (JCI), a global company that began in the 1880s with invention of the electric thermostat. Today Johnson Controls develops technologies and services to manage energy use and security in buildings. Such technologies and services leverage large amounts of data on the internal and external environment of buildings, and utilizing a wide variety of algorithms (including

AI) to serve customer needs. For assets and operations, JCI employs integrated AI solutions, starting with operational data (energy optimization), computing the optimal operating conditions, detecting early indication of pending equipment failure, estimating the probability of failure, and optimize least-cost maintenance actions. Dr. Lee identified four key AI capabilities that together create a smart building: (1) sense (accessing, collecting, and cleansing data from the building systems and equipment), (2) learn (the evolution and patterns of building states), (3) predict (of various building states in the future and enabling intelligent decisions of control actions), and (4) control/optimize (changes to make a building more energy efficient, productive, and safe). He described four use cases: chiller shutdown prediction modeling (for predictive maintenance), chiller vibrational analysis using machine learning to filter out normal operating chillers, energy prediction modeling for whole buildings to reduce energy costs (typically by 10 percent), and optimal control of HVAC systems to reduce energy costs (typically by 20 percent). He reflected on these uses cases to suggest five attributes of successful AI applications: large amounts of high-quality data, significant domain knowledge (of the building, HVAC, security, etc.), competency in AI algorithms (data science), an Internet of Things computing platform (i.e., digital vault), and testbeds (for the building and for HVAC).

Bill Cronin (Xometry) described custom manufacturing and a new business model based on AI that is increasing the efficiency of U.S. custom manufacturers. Most U.S. manufacturers are small firms, with fewer than 20 people. This is certainly true for custom manufacturers (including machine shops), which involves producing in batch (versus continuous) mode, greater R&D expenditures, and shorter time to market. The process of finding and securing a custom manufacturer can be a time-intensive process, often taking months. Xometry was created to utilize AI in order to link buyers and sellers of custom goods quickly and efficiently. Xometry now has the largest on-demand manufacturing platform, which benefits both buyers (optimal pricing, one-click ordering, and reliable delivery) and suppliers (low-cost distribution, improved profitability, better cash flow). Xometry utilizes AI to match customers with suppliers instantly and generate an optimal price. Its pricing model was trained on data from more than 2 million produced parts catalogued on 657 features. By reducing friction in the market, Xometry's AI platform enables transactions to occur in seconds versus months. To be a supplier to Xometry (>3,000) requires information that is used to improve the AI algorithm.

Panelists were asked to identify the biggest obstacles or limitations to realizing more and deeper applications of AI in industry. Panelists offered a range of answers, including an insufficient pipeline of skilled workers, lack of imagination among executives for what is possible with AI, lack of connectedness across factories or supply chains, difficulties in scaling AI at the factory level and across factories, lack of digitization of data, and lack of confidence in AI results. The answers receiving most support across the panel included limitations in data (quality and quantity) and limitations in worker skills.

Panelists were asked what government do to foster AI applications in industry. No single answer generated consensus, but the resulting list includes incentives for small business investments; research on responsible AI; developing a national AI strategy; creating collaborative centers for research/innovation between industry, government, and academia (Holland is an example); expanding the Manufacturing Extension Partnership (MEP) to include AI applications for SMEs; and increasing the pipeline of skilled workers.

OBSERVATIONS AND THEMES

The workshop brought together 19 speakers from government and industry to describe applications of AI across a variety of organizations. And although the speakers/organizations were not chosen randomly, the presentations allowed for a comparison of opportunities and challenges facing the public and private sectors in applying AI.

Every organization represented at the workshop is leveraging large data sets meeting minimum quality standards. Xometry's pricing model was trained on data from more than 2 million produced parts catalogued on 657 features. Fero utilizes continuous, real-time data from more than 1,200 sensors at a typical integrated steel mill. Johnson Controls leverages large amounts of high-quality data on equipment and buildings to optimize its energy management services. Government agencies are making available large data sets in machine-readable form. The VA is designing data formats that allow for bidirectional data flow and the sharing of data sets for use by multiple modelers. Several of the speakers utilize transfer learning to leverage data from one application to inform similar applications.

Furthermore, the tasks being tackled are central to the organizational mission. The VA is using AI to improve health outcomes. NASA is applying AI to minimize the mid-air collisions of aircraft. Energy-intensive manufacturers are using AI to reduce energy costs. Railroads are minimizing unplanned maintenance for locomotives.

A skilled workforce is also critical. For the federal government, attracting and retaining technical expertise is a particular challenge in a transformational field like AI. The ITIF survey showed that skilled labor is a top priority need. Stanley Black & Decker is investing heavily in upskilling and retraining workers at the same time that AI applications are being developed. The scarcity of skilled workers is a major impediment to adoption of machine learning at Kaman Corp. Competency in AI algorithms as central to Johnson Controls' success.

Aside from these common themes, there were some differences in emphasis between the public and private sector panelists. Government speakers emphasized explainability and trust, which are central to ensure public accountability. In gearing AI applications to analysts and operators, DARPA found that users preferred explanations, and correct explanations engendered trust. Other workshop participants noted that for AI to go from narrow to general, it must be trusted. Therefore, researchers must seek to build trust as capacity and sophistication grows.

Industry speakers emphasized the importance of domain knowledge. From applying AI in industrial settings, Fero ensures interpretability by leveraging the domain expertise of factory personnel. Uptake has learned, from domain experts, that not all anomalies recognized by AI are meaningful. Successful applications of AI by Johnson Controls has been predicated on substantial domain knowledge.

Each speaker from industry also emphasized the need for top-down executive commitment. The ITIF survey of manufacturers found that executives generally have a poor general understanding of AI and what it offers. Mark Johnson believes the biggest challenge to applying AI is trust—thinking that the status quo is working, thinking that established processes are working. Yet once executives commit, workers can be persuaded. According to Sudhi Bangalore, much of his time is spent persuading the existing workforce to believe in AI and, once they do, they realize that AI helps them do their job better.

PATH FORWARD/POLICY IMPLICATIONS

The workshop illuminated additional research and policies needed to increase the breadth and depth of AI-enabled applications in the future.

First, the need for a skilled workforce is paramount. This is not a task left to the private sector. Government can play a role here—in the allocation of education and training funds. For example, NSF is investing in K-12 education to incentivize a broad and diverse group of students who are interested in AI. But Stephen Ezell thinks more is needed: a massive expansion of AI talent is called for.

Second, more research is needed to ensure explainability and trust. According to Cindy Goldberg, for AI to go from narrow to general, it must be trusted. Researchers must seek to build trust as capacity and sophistication grows. We must be able to add reasoning to AI applications. According to John Reeder, explainable AI helps the user understand why/why not, the circumstances under which AI fails or succeeds, the circumstances under which AI decisions or determinations can be trusted, and why AI errs when it errs.

Third, major advances in hardware will be needed to support next-generation AI. Current hardware cannot manage the growth of AI without more efficient computing power. And access to infrastructure and cloud computing resources will continue to be a high-priority need.

Fourth, improvements in AI are dependent on continued growth in connectivity and ancillary technologies and associated policies (e.g., cybersecurity). Diana Furchtgott-Roth mentioned the importance of unlicensed spectrum to create a hybrid communications infrastructure to address a wide range of transportation needs. Mark Johnson's list of major challenges include ensuring privacy and security of data throughout a manufacturer's organizational boundaries. Young Lee stressed the need for an Internet of Things computing platform (i.e., digital vault), and testbeds (for the building and for HVAC) for AI applications. The ITIF survey showed that AI is just one of a suite of complementary technologies that comprise the next industrial revolution (Industry 4.0).

Fifth, partnerships between public and private sectors will be needed to advance AI. Many collaborative efforts were mentioned at the workshop: Amazon's partnership with NSF, DOE sharing its AI expertise (from national labs) with federal and non-federal partners, NIST working with federal and non-federal partners to set technical standards. Ezell recommended creating collaborative centers for research/innovation between industry, government, and academia (and pointed to Holland is an example) to promote innovation in pre-commercial technologies. A few panelists emphasized the need for industry to partner with local schools, including community colleges, to ensure a pipeline of skilled workers.

APPENDIX A: AGENDA

Applications of AI in Government and Industry

A workshop organized by the National Academies Innovation Policy Forum in cooperation with the Manufacturing Policy Initiative (MPI) at the Paul H. O'Neill School of Public and Environmental Affairs at Indiana University

September 26, 2019

Room 125
National Academy of Sciences
2101 Constitution Avenue, NW
Washington, DC

AGENDA

9:00 AM	**Coffee and Pastries**
9:25 AM	**Welcome** *Tom Guevara, Indiana University (Co-Chair, National Academies Innovation Policy Forum)*
9:30 AM	**Panel 1: Government Applications** Moderator: *Heather Evans, National Institute of Standards and Technology*
	Trust in AI Computing *Cindy Goldberg, IBM*
	Explainable AI and DARPA XAI Program *John Reeder, Naval Information Warfare Center Pacific*
	AI-Enabled Technologies *Ricardo A. Arteaga, NASA*
10:45 AM	**Coffee Break**
11:00 AM	**Panel 2: Challenges to Applying AI in Government** Moderator: *Dan Chenok, IBM Center for the Business of Government*
	Gil Alterovitz, Department of Veterans Affairs (VA) *Erwin Gianchandani, National Science Foundation* *Diana Furchtgott-Roth, Department of Transportation* *Brian Valentine, Department of Energy*
12:00 PM	**Lunch in the West Court**

1:00 PM **<u>Panel 3: Industry Applications</u>**
Moderator: *Andre Gudger, Founder and CEO of Eccalon (Member, National Academies Innovation Policy Forum)*

Production Process Improvement
Alp Kucukelbir, Fero Labs

Machine Optimization
Jake Bartolone, Uptake

Applications of AI in Energy Management
Mark Johnson, Clemson University

2:15 PM **Coffee Break**

2:30 PM **<u>Panel 4: Challenges to Applying AI in Industry</u>**
Moderator: *Keith Belton, Indiana University*

Stephen Ezell, Information Technology and Innovation Foundation
Sudhi Bangalore, Stanley Black & Decker (via Zoom)
Ernest Begin, Kaman Corporation
Young M. Lee, Johnson Controls
Bill Cronin, Xometry

4:00 PM **Closing Remarks**
Tom Guevara, Indiana University (Co-Chair, National Academies Innovation Policy Forum)

4:10 PM **Reception in the West Court**

APPENDIX B:
LIST OF PANELISTS AND MODERATORS

Gil Alterovitz

Gil Alterovitz is the inaugural director of artificial intelligence at the U.S. Department of Veterans Affairs. Based in the Office of Research and Development, Dr. Gil Alterovitz is working on leveraging health data for AI as well as building AI research and development capacity. He was one of the core authors of White House Office of Science Technology and Policy's *The National AI R&D Strategic Plan* from 2019. He is also a professor at Harvard Medical School and core faculty at the Computational Health Informatics Program, Boston Children's Hospital.

Ricardo Arteaga

Ricardo Arteaga is a research engineer with NASA, where he has worked on the Constellation, SOFIA, and UAS-NAS programs. He has more than 35 years of experience in government and aerospace engineering development programs, including the design and integration of navigation systems on the F/A-22, B-2, F117A, and ER-2 aircraft.

Mr. Arteaga received his engineer degree and master's degree from the University of Southern California in systems engineering and also received a Master of Science in Computer Engineering from Syracuse University.

Sudhi Bangalore

Sudhi Bangalore is Stanley Black & Decker's vice president of Industry 4.0. In this position, Sudhi manages the company's Advanced Manufacturing Center of Excellence and leads our automation efforts with technologies such as the internet of things, cloud computing, artificial intelligence, 3-D printing, robotics and advanced materials. Prior to joining the company in 2017, Sudhi was the global head of Smart Manufacturing and Industry 4.0 Solutions with WIPRO, and before that was the global practice head for Industrial Automation. Sudhi also held leadership roles with technology companies such as Danaher Corporation, Siemens, and Rockwell Automation.

Mr. Bangalore earned a bachelor's degree in electronic engineering from Bangalore University, a master's degree in industrial engineering from the University of Louisville and a Master of Business Administration from Kent State University.

Jake Bartolone

Jake Bartolone is data science lead with Uptake, a Chicago-based data analytics firm. He holds a Ph.D. in psychology from the University of Chicago. Before joining Uptake, Jake was director of the higher education analytics center at NORC at the University of Chicago.

Ernest Begin

Ernest Begin is the executive director of information security and governance at Kaman Corporation. In that role he is responsible for the information security stature of the company including development of and compliance with IT policy; detection, reaction and prevention of cyber threats; and the education of the workforce in regards to information security.

Prior to this role Ernie spent 13 years in progressive roles at United Technologies Corporation spanning IT security, compliance, audit and governance. He is CISA- and CISSP-certified and graduated from WPI in 2005 with a bachelor's degree in management of information systems.

Keith B. Belton

Keith B. Belton is director of the Manufacturing Policy Initiative (MPI) in the Paul H. O'Neill School of Public and Environmental Affairs at Indiana University. The Initiative is focused on public policies affecting the U.S. manufacturing sector. Prior to this appointment, Keith held several positions involving public policy analysis and/or government relations, including with the U.S. Office of Management and Budget and the Dow Chemical Company. He holds a Ph.D. in public policy from The George Washington University.

Dan Chenok

Dan Chenok is executive director of the IBM Center for The Business of Government. He oversees all of the Center's activities in connecting research to practice to benefit government, and has written and spoken extensively around government technology, cybersecurity, privacy, regulation, budget, acquisition, and presidential transitions. Mr. Chenok previously led consulting services for Public Sector Technology Strategy, working with IBM government, healthcare, and education clients. Before joining IBM, Mr. Chenok was a senior vice president for civilian operations with Pragmatics, and prior to that was a vice president for business solutions and offerings with SRA International.

As a career government executive, Mr. Chenok served as branch chief for information policy and technology with the Office of Management and Budget, where he led a staff with oversight of federal information and IT policy. Mr. Chenok began his government service as an analyst with the Congressional Office of Technology Assessment. Mr. Chenok earned a BA from Columbia University and a Master of Public Policy degree from Harvard's Kennedy School of Government.

Bill Cronin

Bill Cronin is the chief revenue officer of Xometry. He has more than 20 years of experience in senior sales and marketing roles at MasterCard, *USA Today*, and WeddingWire. At WeddingWire, he was the vice president of marketing, where he led consumer and B2B marketing as the vertical marketplace more than doubled revenues in two years. Prior to WeddingWire, Cronin served as vice president of marketing at *USA Today*, leading the relaunch of the brand across all channels. He also held a number of senior roles at MasterCard over a decade. Bill is a graduate of Dartmouth College.

Heather Evans

Heather Evans is a senior program analyst in the Program Coordination Office in the Director's Office at NIST. Her portfolio includes artificial intelligence, advanced manufacturing, data policy, and open innovation. She co-developed the NIST AI Community of Interest and helped develop NIST's Plan for Federal Engagement in Developing Technical Standards and Related Tools. Heather coordinates NIST's interagency engagements with the White House Office of Science and Technology Policy (OSTP) and participates on many National Science and Technology Council subcommittees. Prior to working at NIST, Heather held a position at OSTP and served as an AAAS Science and Technology Policy Fellow in the National Nanotechnology Coordination Office. Heather was a Humboldt Postdoctoral Research Fellow at the Max Planck Institute for Dynamics and Self-Organization in Gottingen, Germany, and earned her doctorate in materials science from the University of California Santa Barbara. She graduated magna cum laude from Macalester College in St. Paul, Minnesota.

Stephen Ezell

Stephen Ezell is vice president of global innovation policy at the Information Technology and Innovation Foundation (ITIF), a Washington-DC based technology and economic policy think tank, where he focuses on science, technology, and innovation policy as well as international competitiveness, trade, and manufacturing policy issues. He is the co-author with Dr. Robert Atkinson of *Innovation Economics: The Race for Global Advantage* (Yale, September 2012) and a co-author of *Innovating in a Service-Driven Economy: Insights, Application, and Practice* (Palgrave McMillan, November 2015).

Young M. Lee

Young M. Lee is a Distinguished Fellow and the Director of Artificial Intelligence Product Development at Johnson Controls and is leading a global team of AI scientists and developing Industrial IoT solutions using ML/DL/AI and optimization technology. Dr. Lee previously worked for 15 years at IBM T.J. Watson Research Center as a research staff member, a research manager, and an IBM master inventor, and developed industrial applications of mathematical modeling, optimization and machine learning. Dr. Lee also worked for BASF for 14 years, founded and managed the Mathematical Modeling Group, and led development of numerous optimization and simulation models for various manufacturing and supply chain processes. Dr. Lee received B.S., M.S., and Ph.D. degrees from Columbia University.

Diana Furchtgott-Roth

Diana Furchtgott-Roth is deputy assistant secretary for research and technology at the U.S. Department of Transportation. Previously, she served as acting assistant secretary for economic policy at the U.S. Department of Treasury. She has been Senior Fellow and director of Economics21 at the Manhattan Institute for Policy Research and an adjunct professor of economics at George Washington University. She served as chief economist of the U.S. Department of Labor; chief of staff of the President's Council of Economic Advisers; deputy executive director of the Domestic Policy Council; and junior staff economist at the Council of Economic Advisers. Ms. Furchtgott-Roth received her B.A. in economics from Swarthmore College and her M.Phil. in economics from Oxford University.

Erwin Gianchandani

Erwin Gianchandani is acting assistant director for Computer and Information Science and Engineering at the National Science Foundation. In this role, he contributes to all aspects of the management of the CISE directorate, including strategic and human capital planning, formulation and implementation of the directorate's more than $900 million annual budget, and oversight of day-to-day operations. In the last several years, he has led the development, launch, and implementation of several new NSF investment areas, including Smart & Connected Communities and Platforms for Advanced Wireless Research. Dr. Gianchandani earned his Ph.D. and M.S. in biomedical engineering and his B.S. in computer science from the University of Virginia.

Cindy Goldberg

Cindy Goldberg is the program firector of the IBM Research AI Hardware Center in Albany, NY. The Center is a global collaboration network between IBM, industry partners, academic partners, and the State of New York, launched in February 2019 to create and mature new AI hardware technologies for cloud and edge applications. The AI Hardware Center is also home for a first of its kind testbed for developing AI applications, with an ensemble of tools for simulation, emulation, and benchmarking against tailored infrastructure topology.

Prior to joining IBM, Cindy was director of engineering at STMicroelectronics and previously at Freescale Semiconductor. She earned her Ph.D. in physics from the State University of New York at Albany.

André Gudger

André Gudger is the founder and CEO of Eccalon with over 20 years of experience leading middle-market technology companies and senior positions within the federal government. From 2011-2017, he served as President Barack Obama's Deputy Assistant Secretary of Defense for Manufacturing and Industrial Base Policy (MIBP) and the Director of the Office of Small Business Programs. As the Deputy Assistant Secretary of Defense, André provided detailed analyses of the increasingly global and financially complex industrial supply chain and took appropriate actions to maintain the health, integrity and technical superiority of the industrial base.

Tom Guevara

Thomas Guevara is director of the Indiana University Public Policy Institute (PPI), appointed as its fifth director in April of 2017. He also holds a lecturer appointment in the O'Neill School of Public and Environmental Affairs. Mr. Guevara has served the public sector in a variety of local, state and federal government executive positions, and as a senior consultant to public sector clients. Immediately prior to joining PPI, Mr. Guevara served as a Deputy Assistant Secretary of the U.S. Department of Commerce's Economic Development Administration.

Mark Johnson

Mark A. Johnson is Thomas F. Hash Endowed Chair in Sustainable Development and Director of the Center for Advanced Manufacturing at Clemson University. He is the former Director of the Advanced Manufacturing Office (AMO) in the Office of Energy Efficiency and Renewable Energy (EERE), which enables development of new processes and materials technologies to reduce the energy intensity and life-cycle energy consumption of manufactured products.

Previously, Mark served as a program director in the Advanced Research Projects Agency–Energy (ARPA-E). He also served as the Industry and Innovation Program Director for the Future Renewable Electric Energy Delivery and Management (FREEDM) Systems Center. Mark has a bachelor's degree from MIT and a Ph.D., from NC State, both in materials science and engineering.

Alp Kucukelbir

Alp Kucukelbir is co-founder and chief scientist of Fero Labs, an industrial machine learning software company. Fero develops actionable ML software for the metals, chemicals, energy, and automotive sectors, who use Fero to minimize emissions, increase product quality, and diagnose root causes of production issues. Alp is also an adjunct professor of computer science at Columbia University. He has developed core machine learning algorithms that have been adopted into open-source software packages, such as Stan, PyMC, and Pyro. Alp received his Ph.D. from Yale University, where he won the best thesis award.

John Reeder

John David Reeder II is a senior research engineer for the Naval Information Warfare Center Pacific (NIWC Pacific). Dr. Reeder's research focuses on the application of Machine Learning techniques to the fields of Autonomous Systems, Intelligence Surveillance & Reconnaissance (ISR), and Electronic Warfare (EW). His research in autonomous systems concentrates on the robustness and trustworthiness of learned behaviors to improve human-autonomy teaming. His work in ISR and EW seeks to improve generalization and reliability of analysis tools by incorporating machine learning with existing tools, allowing them to work in scenarios beyond their original design. John received his Ph.D. in computer engineering from the University of Central Florida with a focus in Life-Long Machine Learning for Simulated Agents.

Brian Valentine

Brian G. Valentine is a technical project officer for the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy EERE. Brian has worked for EERE for nearly 30 years, as technical manager for the Solar Thermal Technology program, the Industrial Energy Management program, the Advanced Manufacturing program, and as consultant to DOE's Office of Basic Energy Sciences for applied engineering science. Brian worked for the U.S. government's Iraq reconstruction mission from 2004 through 2006. Brian holds a Bachelor of Science in Physics from Siena College and a Ph.D. in engineering from the Rensselaer Polytechnic Institute, and is a registered professional engineer in Virginia.