




Artificial Intelligence: An Agenda for 2021

*Academy Election 2020 Project
Working Group:
Make Government AI Ready*





ABOUT THE ACADEMY

The National Academy of Public Administration (the Academy) is an independent, nonprofit, and nonpartisan organization established in 1967 to assist government leaders in building more effective, accountable, and transparent organizations. Chartered by Congress to provide nonpartisan expert advice, the Academy's unique feature is its over 950 Fellows—including former cabinet officers, Members of Congress, governors, mayors, and state legislators, as well as prominent scholars, business executives, and career public administrators. The Academy helps the federal government address its critical management challenges through in-depth studies and analyses, advisory services and technical assistance, congressional testimony, forums and conferences, and online stakeholder engagement. Under contracts with government agencies, some of which are directed by Congress, as well as grants from private foundations, the Academy provides insights on key public management issues, as well as advisory services to government agencies.

ABOUT THE ELECTION 2020 PROJECT

The Academy formed a series of Working Groups of its Fellows to address [Grand Challenges in Public Administration](#). These Groups were charged with producing one or more papers to advise the Administration in 2021 (whether reelected or newly elected) on the key near-time actions that should be taken to begin addressing Grand Challenges. This is a paper of the [Make Government AI Ready](#) Working Group. It includes these Fellows' recommendations for the U.S. to maximize the benefits of AI while protecting America through core values, safeguards, and collaboration.

Copyright © 2020 by National Academy of Public Administration.
All rights reserved. Published and hosted by the Academy.

ARTIFICIAL INTELLIGENCE: AN AGENDA FOR 2021

A REPORT OF AN ACADEMY WORKING GROUP

**NATIONAL ACADEMY OF PUBLIC ADMINISTRATION
ELECTION 2020 WORKING GROUP:
MAKE GOVERNMENT AI READY**

Working Group Members

Alan Shark, Chair
Erik Bergrud
Kaye Husbands Fealing
James Hendler
Theresa Pardo
Doug Robinson

Special Program Support

Tonya Thornton Neaves
Tyler Goodwin (GMU Intern)

THE CHALLENGE

Artificial Intelligence (AI) allows computerized systems to perform tasks traditionally requiring human intelligence: analytics, decision support, visual perception, and foreign language translation. AI and Robotics Process Automation (RPA) have the potential to spur economic growth, enhance national security, and improve the quality of life. In a world of “Big Data” and “Thick Data,” AI tools can process huge amounts of data in seconds, automating tasks that would take days or longer for human beings to perform.

The public sector in the United States is at the very beginning of a long-term journey to develop and harness these tools. Chatbots are being used in citizen engagement systems; AI technology is augmenting decision-making in the areas of cyber security monitoring, public policy modeling, database anomalies, and waste and abuse identification.

Although AI in the public sector can yield numerous benefits—including improving customer service and efficiency, while allowing employees to focus more on core agency missions—it also raises concerns about bias, security, and transparency. With biased data, AI systems will produce biased results. Cybersecurity will be more important than ever to protect against malicious actors that, by taking over AI systems, could do significant damage very quickly. Without transparency, the public may be confused about how key decisions were made. And governments may need to revamp their budgeting and procurement processes to be able to quickly acquire and deploy advanced technologies.

This panoply of issues is why the National Academy of Public Administration (the Academy) identified [Make Government AI Ready](#) as one of the [Grand Challenges in Public Administration](#). This paper by the Election 2020 Working Group for this topic provides information on the key elements of AI and its public administration and policy drivers. The Working Group has several recommendations for actions that the Administration in 2021 (whether reelected or newly elected) should take to address this Grand Challenge. These recommendations are designed to

maximize the benefits of AI while protecting America through core values, safeguards, and collaboration.

KEY DRIVERS OF AI

According to the Merriam-Webster Dictionary, AI is:

1: a branch of computer science dealing with the simulation of intelligent behavior in computers

2: the capability of a machine to imitate intelligent human behavior

Other definitions include this one: “Basically, AI is the ability of a machine or a computer program to think and learn.” Definitions of AI abound and these, like most definitions, have many different elements complicating many discussions about AI. But, there are some commonalities among these definitions. The most common and useful elements are drawn on in this paper.

The concept of AI is based on the idea of building machines that are capable of thinking, acting, and learning like humans. A more accurate definition might start by stating that AI is not a specific technology unto itself but is instead a broad concept whereby machines are programmed to perform tasks that one could call intelligent or smart. AI is often confused with machine learning, with the latter having been around for some 50 years. In reality, machine learning, at least today, is the most recognizable application of AI. It can be said that AI encompasses a collection of technologies that include machine learning as well as other technologies like natural language processing, inference

Drivers of AI in 2020

1. Advancements in complex algorithms.
2. Dramatic increase in computing speed and power.
3. Ability to digest data from various sources.
4. Ability to store and retrieve massive amounts of data.
5. Ability to “self-learn”.
6. Advancements in artificial speech and recognition.

algorithms, and neural networks. We are already active participants with AI every time we interact with Alexa or Siri. To more fully understand AI in its current form, it may be useful to recognize the specific drivers that make it all work (See Text Box).

From the early 1950s through the 70s, neural networks were developed, allowing machines to do things such as compete against humans in chess. From the 1980s through 2010, machine learning became popular. In our present day, deep learning has driven AI growth. Today, AI systems have been designed to interact with humans through speech recognition as well as through the written word. These systems have the ability to mimic the human voice. AI systems have had the most successful outcomes in recognizing patterns, anomalies, and thought process which includes the ability to contemplate, anticipate, and provide judgement. Some have stated that AI is more about “augmented intelligence” than “artificial intelligence”.

The Government Accountability Office (GAO) has introduced the concepts of the “three waves of AI.” The first wave encompasses AI in its simplest forms – expert knowledge of criteria transposed into logical reasoning to be encoded into a computer program (such as an online tax form). The second wave includes machine or statistical learning, which includes voice recognition, natural language processing, and computer vision technology. GAO claims even the most complex AI systems today are in the second wave. The third wave is reserved for the most sophisticated AI, of which most AI has not broken the threshold. The third wave combines the characteristics of the first two waves and is also capable of contextual sophistication, abstraction, and explanation. In terms of risks, GAO claims AI utilization could result in personal data being used against the individual from which it came, such as medical records being used to deny them insurance or employment.

Clearly, as we learn more about AI, more questions arise about the trustworthiness of AI systems and the quality of the data that drives those systems. There are legitimate concerns, for example, about machines circumventing human intelligence and action, poor or flawed decision

making, concentration of power in a few computer systems, the potential for algorithms that deviate from socially accepted and cherished norms, and decisions that are made too quickly. There are other concerns about the basic quality of data in terms of accuracy and completeness, as well as growing concerns about bias in data and the impacts of big data use in AI on privacy. With biased data, AI systems will produce biased results. These concerns are shared by computer and data scientist but also by social scientists and public administrators.

A majority of Americans see the need to carefully manage AI (Zhang and Dafoe, 2019), with the greatest importance placed on safeguarding data privacy; protecting against AI-enhanced cyber-attacks, surveillance, and data manipulation; ensuring the safety of autonomous vehicles; accuracy and transparency of disease diagnosis; and the alignment of AI with human values. Public trust in AI must be created; otherwise, useful AI products may be rejected, and government decision making may lose its legitimacy. Ethical frameworks, now plentiful, are rarely accompanied by tangible implementation in AI development (Hagendorff, 2019, Crawford, 2019). Indeed, there is little guidance on how to design policy and governance structures that implement even the most broadly accepted AI ethical principles.

AI holds great promise, but raises concerns about transparency. Without transparency, the public may be confused about how key decisions were made and therefore may not trust them. For example, if AI is applied in procurement, how will these processes be made transparent and understandable to the public? Accordingly, governments at all levels must work collaboratively to promote public trust in the development and deployment of AI tools; train an AI-ready workforce for both the public and the private sectors; and address the ethical concerns about AI's potential downsides in the areas of discrimination, civil liberties, and privacy.

Questions that policymakers and administrators must consider include:

- What are the conditions under which an AI system can be considered explainable?
- Is it appropriate (and, if so, under what conditions) to use unexplainable AI systems in ways that significantly impact human beings?
- How should departments and agencies adjudicate between different (and incompatible) conceptions of fairness, and which data-collection and data-interpretation practices are important for avoiding bias and discrimination?
- What are leading practices for developing AI systems that can be used by humans in ways that improve the human condition?
- What incentives should be established for data sharing?
- How can safe and security be improved?
- What reforms should be made to the current regulatory approach?
- What are acceptable risks?
- How can ethical decision-making be ensured?
- What is the impact of AI on jobs and training?

RECENT ADMINISTRATION EFFORTS

On February 11, 2019, President Trump signed Executive Order 13859 announcing the American AI Initiative — the United States’ national strategy on artificial intelligence. This strategy is a concerted effort to promote and protect national AI technology and innovation. The Initiative implements a whole-of-government approach in collaboration and engagement with the private sector, academia, the public, and like-minded international partners. It directs the Federal government to pursue five pillars for advancing AI: (1) invest in AI research and development (R&D), (2) unleash AI resources, (3) remove barriers to AI innovation, (4) train an AI-ready workforce, and (5) promote an international environment that is supportive of American AI innovation and its responsible use. The U.S. is also actively leveraging AI to help the Federal government work smarter in its own services and missions in trustworthy ways.

In February 2020, the White House released the American Artificial Intelligence Initiative: Year One Annual Report. In the year since the AI Executive Order was signed, the Administration called for record amounts of AI R&D investment, led the development of the first international statement on AI Principles, issued the first-ever strategy for engagement in AI technical standards, published the first-ever reporting of government-wide non-defense AI R&D spending, and released the first-ever AI regulatory document for the trustworthy development, testing, deployment, and adoption of AI technologies.

The White House also claims to be focused on introducing regulatory principles for AI in American industries and maximizing benefits for the American worker. They have introduced 10 principles on AI:

- Public trust in AI
- Public participation
- Scientific integrity and information quality
- Risk assessment and management
- Benefits and costs
- Flexibility
- Fairness and nondiscrimination
- Disclosure and transparency
- Safety and security
- Interagency coordination

In September 2019, the White House hosted the Summit on Artificial Intelligence in Government to spark ideas for how the Federal government can adopt AI to better achieve its mission and improve services to the American people. More than 175 leaders and experts from government, industry, and academia came together to identify best practices in the use of AI, opportunities to foster collaborative partnerships, and ways to develop a Federal AI workforce. The Summit highlighted innovative efforts at Federal agencies that have already adopted AI, and looked ahead to future transformative AI applications that will make government more effective, efficient, and responsive.

To improve the coordination of Federal efforts related to AI, the White House chartered a Select Committee on AI under the National Science and Technology Council. The Select Committee consists of the most senior R&D officials across the Federal government and represents a whole-of-government approach to AI R&D planning and coordination. This Committee advises the White House on interagency AI R&D priorities; considers the creation of Federal partnerships with industry and academia; establishes structures to improve government planning and coordination of AI R&D; and identifies opportunities to prioritize and support the national AI R&D ecosystem. The Select Committee also provides guidance and direction to the existing Machine Learning and AI Subcommittee, which serves as the Committee's operations and implementation arm. (www.ai.gov)

WHAT IS HAPPENING NOW: KEY FEDERAL AND STATE AI INITIATIVES

NOAA

NOAA has instituted its AI strategy goals, which include establishing an efficient organizational structure and processes to advance AI across NOAA, advance AI research and innovation in support of NOAA's mission, accelerate the transition of AI research to operational capabilities, strengthen and expand AI partnerships, and promote AI

proficiency in workforce. The goals align with the White House’s Executive Order on AI. NOAA is developing an AI Strategic Implementation Plan that defines detailed action items, deadlines, and responsibilities. Until then, its strategy is already improving performance in lifesaving and economically impactful missions, setting the course to strengthen environmental science and technology leadership.

- NOAA AI Strategy goals:
 - Establish an efficient organizational structure and processes to advance AI across NOAA
 - Advance AI research and innovation in support of NOAA’s mission
 - Accelerate the transition of AI research to operational capabilities
 - Strengthen and expand AI partnerships
 - Promote AI proficiency in the workforce

DoD

The Department of Defense (DOD) is using its AI strategy to “harness AI to advance our security and prosperity.” Realizing this vision requires identifying appropriate uses for AI across DoD, rapidly piloting solutions, and scaling success across the enterprise. The DoD’s AI strategy aims to drive the urgency, scale, and unity of effort necessary to navigate these transformations. To further accomplish these goals, the DoD is collaborating with those in government, academia, non-traditional centers of innovation in the commercial sector, and international partners. The DoD strategy echoes the importance of America being a leader in AI and having an AI-ready workforce. Additionally, DoD has committed to be a leader in military ethics and AI safety. There are four Strategic Focus Areas: delivering AI-enabled capabilities that address key missions, partnering with leading private sector technology companies, academia, and global allies, cultivating a leading AI workforce, and leading in military ethics and AI safety.

DoD's 5 Ethical Principles for AI:

- Responsible
- Equitable
- Traceable
- Reliable
- Governable

DoD – Military, Navy

The Navy Center for Applied Research in AI focuses on several areas of research: intelligent systems, adaptive systems, the interactive systems, and the perceptual systems. Currently, its research has led to projects like 3D Audio-Cued Operator Performance Modeling, Adaptive Testing of Autonomous Systems, Chat Attention Management for Enhanced Situational Awareness, Cognitive Robots and Human Robot Interaction, and several other topics regarding adaptive AI. The center is directed toward understanding the design and operation of systems capable of improving performance based on experience, efficient and effective interaction with other systems and with humans, sensor-based control of autonomous activity, and the integration of varieties of reasoning as necessary to support complex decision-making.

DoD – Military, Air Force

The Air Force's AI Strategy serves as a framework for aligning their efforts with the DoD's AI Strategy and the National Defense Strategy. It provides definition, context and purpose for AI in the Air Force. While echoing and complying with the DoD's AI strategy, the Air Force's strategy goes further by defining five specific focus areas: drive down technological barriers to entry; recognize and treat data as a strategic asset; democratize access to AI solutions; recruit, develop, upskill, and cultivate the workforce; increase transparency and cooperation with international, government, industry, and academic partners.

DOT

The Department of Transportation's (DOT) website says: "[DOT] is committed to safety and innovation and sees AI as a promising capability to help achieve these aims." DOT believes AI can be applied within transportation by enabling the safe integration of AI into the transportation system and adopting and deploying AI-based tools into internal operations, research, and citizen-facing services. Its research guidelines for AI support the development of regulations, policies, procedures, guidance, and standards for drone operations. DOT currently seeks data in critical areas such as detection and avoidance, communications, human factors, system safety, and certification to enable it to make informed decisions on safe drone integration.

NSF

The National Science Foundation (NSF) is a recognized leader across the US Federal Government in both advancing the use of AI in the public sector within NSF and across the Federal Government, and in advancing both fundamental and use-inspired AI research in the United States. NSF carries out this leadership role along three parallel paths: providing leadership on AI for the US Government, guiding AI use within NSF, and enabling and advancing AI research and development nationally.

NSF plays a significant role in advancing AI across the U.S. Government. The Director of NSF serves as co-chair, with the Director of Defense Advanced Research Projects Agency (DARPA) and senior leadership from the White House Office of Science and Technology Policy (OSTP), of the Select Committee on AI, a committee of the National Science and Technology Council. The Select Committee serves as the high-level budget authority setting the direction for USG investments for all facets of AI. NSF also co-chairs, with NIST and OSTP, the operational arm of the Select Committee on AI, the Subcommittee on Machine Learning and AI. These committees work together and across the USG to coordinate investments, the directionality of those investments, and what impact they are trying to have concerning AI research, education, and infrastructure. *For additional information on AI activities within NSF, please turn to the Addendum.

NASA

The National Aeronautics and Space Administration (NASA) is utilizing machine learning and applying it to space science. To make sure the agency is using machine learning in research, NASA's Frontier Development Lab brings together innovators for eight weeks every summer to brainstorm and develop computer code. They aim to advance machine learning techniques to quickly interpret data revealing the chemistry of exoplanets based on the wavelengths of light emitted or absorbed by molecules in their atmospheres.

USDA

The U.S. Department of Agriculture (USDA) is testing out FarmBeats, which collects data from multiple sources, such as sensors, drones, satellites, and tractors and feeds it into cloud-based artificial intelligence models that provide a detailed picture of conditions on the farm. This is designed to maximize agricultural outputs based on the weather inputs detected by the sources. USDA partners with public and private organizations to develop tools and practices like robots for agriculture, instruments for crop and soil monitoring, and predictive analytics. Their objectives include developing crop production systems to intensify plant and forest production with continuous improvements and adoption of new technology while reducing environmental impacts, advance science-based approaches to combat outbreaks of emerging pests and diseases, enhance plant product quality, and evaluate the adoption and use of enhanced technologies such as data analytics and precision agriculture. They are also using AI to enhance animal production, health, and genetics.

NASCIO

The National Association of State Chief Information Officers (NASCIO) has identified five emerging best practices and recommendations. These include: consider creating a framework for AI adoption; create multidisciplinary teams to address change management; assess data availability and capitalize on automation to make the data underlying AI more trustworthy; modernize legacy infrastructure with targeted technology investments; and choose AI projects where success can be clearly measured. NASCIO also identifies the obstacles to AI implementation. In short, the most difficult task is overcoming the constraints of current IT infrastructures; current designs are not equipped to handle the large volumes of data necessary for advanced analytics or AI applications. Other barriers include cultural concerns inside organizations, lack of necessary staff skills for AI, organizational data silos, and lack of executive support. “The roles and responsibilities for policy development are still being ironed out...” NASCIO claims that the focus of AI today

should be educating policymakers and increasing their confidence in governing and understanding AI.

RECOMMENDATIONS

The Working Group offers recommendations that are aimed at making government AI ready. This includes the federal government taking actions to set the conditions for success through data sharing, create a regulatory environment that fosters both innovation and citizen protections, enable public-private partnerships to advance both the state of the art and the deployment of practical tools, and address the significant issue of workforce readiness.

The Working Group's recommendations for making government AI ready are presented below organized in five themes:

1. Build Trustworthy AI
2. Use Ethical Frameworks to Identify and Reduce Bias
3. Build Intergovernmental Partnerships and Knowledge Sharing around Public Sector Uses of AI
4. Increase Investments in AI Research and Translation of Research to Practice
5. Build an AI Ready Workforce

1. Build Trustworthy AI

According to the National Science Foundation, "increasing trust in AI technologies is a key element in accelerating their adoption for economic growth and future innovations that can benefit society." However, NSF further points out that our ability to understand and analyze the decisions of AI systems and measure their trustworthiness is limited. Aspects of trustworthy AI include reliability and explainability.

Recommendations:

- Establish a single, authoritative, and recognized federal entity with a focus on trustworthy AI's long-range social, cultural and political effects (Crawford and Calo, 2016) that industry and government organizations can go to for guidance, to find solutions, or to propose challenges in Trustworthy AI.
- Leverage existing investments such as NSF's AI Institutes to create guidance and solutions including connecting AI research to public policy and governance, to ensure a broad social systems approach to trustworthy AI that will address a fundamental adoption barrier, namely, public trust in AI.

2. Use Ethical Frameworks to Identify and Reduce Bias in AI

Making government AI-ready necessitates a perpetual commitment to earning and sustaining community trust. In advancing an AI agenda, the Administration and Congress should consider a range of issues, including (a) ethical and moral questions; (b) greater public education about the benefits and risks of AI; (c) regulatory frameworks and guidelines; (d) legislation linked to current and future ethical issues; and (e) the proper relationship between technology, society, and public law. Introducing AI applications without an effectively structured system of oversight could lead to public concerns about “technological dystopia.”

Recommendations:

- Demonstrate a federal government-wide commitment to ethical principles and standards in AI development and use, such as those included in the American Society for Public Administration's Code of Ethics.
- Collaborate with the Association for Computing Machinery (ACM), the largest professional society for computing, in incorporating industry guidance in AI ethics training.

- Require departments and agencies to implement AI ethical frameworks, such as the AI ethical principles adopted by the DOD in February 2020.
- Institute a moratorium on the use of facial recognition technologies and conduct research on eliminating racial bias in AI applications.

3. Build Intergovernmental Partnerships and Knowledge Sharing around Public Sector Uses of AI

States and local governments are already utilizing basic forms of AI in practice. State and local governments have a keen interest in applying AI to other areas of government but often lack the resources and expertise to do more. This working group urges more intergovernmental sharing of information and research opportunities.

Recommendations:

- Establish an interagency and intergovernmental mechanism to develop a comprehensive AI strategy that addresses the following:
- The need to share leading practices for governance of AI with state and local governments. In a recent survey of state CIOs, 72 percent said that they do not have a policy regarding the responsible use of AI to ensure that it is governed by clear values, ethics, and laws. Leading practices regarding governance of AI can be applied across all levels of government.
- A call for broader use of AI in cybersecurity and sharing of federal developments with state and local governments. In the same survey, 78 percent of state CIOs chose cybersecurity as an area where AI could make the most measurable improvement and impact in their state. Finding a qualified cybersecurity workforce is very difficult for state and local governments who cannot compete with private sector salaries. AI solutions supporting threat detection can augment the workforce. The federal government should share with state and local governments any resources

developed for federal agencies using AI to protect from cyber-attacks.

- The need to incentivize and stimulate broader AI adoption in state and local governments through Federal agency programmatic funding. For example, appropriate use of AI and machine learning for fraud detection in federally-funded benefit programs. Recent UI claims fraud is a noteworthy example.
- Gaps in readiness to build an AI workforce that can work on all levels of government. In the same survey of state CIOs, 27 percent agreed that a lack of necessary staff skills was a significant challenge or barrier to AI adoption. Support and fund the AI Scholarship for Service Act (S3901) that contains provisions for eligibility by state and local governments.

4. Increase Investments in AI Research and Translation of Research to Practice

The White House's *National Artificial Intelligence Research and Development Strategic Plan: 2019 Update* ("2019 Plan") denotes the priority areas for federal investments in AI research and development and identifies specific goals under these key strategies: *make long-term investments in AI research; develop effective methods for human-AI collaboration; understand and address the ethical, legal, and societal implications of AI; ensure the safety and security of AI systems; develop shared public datasets and environments for AI training and testing; measure and evaluate AI technologies through benchmarks and standards; better understand the national AI R&D workforce need; and expand public-private partnerships to accelerate advances in AI.* Each of these strategies are designed to address the need for research into AI. Upon completing such research, the White House has made commitments to share the information to further the United States' advantages in AI. Executive Order 13859 requires agencies to 'increase public access to government data and models where appropriate'. Additionally, other

measures have been put in place to ensure the government research and information about AI is widely accessible, discoverable, and usable.

The current Administration has made several strides in maximizing AI utilization. A key component to achieving this goal is increasing public access to AI research. This will require that more research into AI is unclassified. In 2018, the White House committed to prioritizing Federal investment in unclassified research and development for AI and related technologies. In the same year, investment in such research was shown to have increased over 40% since the year 2015. This is absolutely necessary in advancing AI and other life-enhancing technologies; however, measures must be taken to protect individuals' personal privacy. Further research must be conducted to understand how AI can be utilized without relinquishing privacy at the individual level. Additionally, research into removing biases from AI must be prioritized. Presently, AI is allowed to perpetuate the injustices that plague human history, whether they are racial, misogynistic, or homophobic. These biases must be eliminated to reflect democratic norms and values. Life-altering technologies should not be exclusionary, and their objective should be to bring the American Dream closer to every American, rather than pushing it away. Without ethical principles to guide AI policy, the risk of AI being used to further oppress minorities increases. The key to preventing this lies within research and the communication of AI research.

Recommendations:

- Increase public access to federal government data.
- Increase by at least 50% investment into unclassified AI research.
- Ensure the protection of privacy at the individual level.
- Remove biases from programming to ensure equitable treatment.

5. Build an AI Ready Workforce

There is much concern that the workforce of today and possibly tomorrow lacks the skills necessary to fulfill the needs and requirements of an AI-Ready Government. With every advancement in AI, governments at all levels will continue to face the ongoing challenges of ethics, privacy,

human control, policy bias, predictive analytics, decision-making, citizen engagement, planning, and the future of work.

Those who desire to enter public service will need to have an entirely new set of competencies that cut across traditional program offerings found at most institutions of higher learning. Over the next decade, lawmakers and senior public managers will be making monumentally important decisions about the role of AI in society. It is therefore imperative that we provide them the skills, tools and knowledge in order to make the most effective and ethical use of AI in all levels of government.

Recommendations:

- The federal government should provide direct funding to support the growth of a federal workforce with AI competency, including through funds to train the next generation of faculty, researchers, and graduate students.
- Federal agencies must leverage expertise in the AI R&D workforce that spans multiple disciplines and skill categories to ensure sustained national leadership.
- Studies, supported by the federal government, are needed to increase understanding of the current and future national workforce needs for AI R&D. Data is needed to characterize the current state of the AI R&D workforce, including the needs of academia, government, and industry.
- The federal government must develop policies and fund incentives that encourage the AI R&D workforce to use multidisciplinary teams comprising not just computer and information scientists and engineers, but also experts from other fields key to AI and machine learning innovation and its application. These include cognitive science and psychology, economics and game theory, engineering and control theory, ethics, linguistics, mathematics, philosophy, and the many domains in which AI may be applied.

CONCLUSION

AI holds great promise for the public sector, but also raises important questions about bias, security, and transparency. The Working Group's recommendations are designed to help the Administration in 2021 (whether reelected or newly elected) maximize the benefits of AI while protecting America through core values, safeguards, and collaboration. By implementing recommendations to build trust AI, identify and reduce bias, build intergovernmental partnerships and knowledge sharing, increase investments in research and its translation to practice, and develop an AI-ready workforce, the Administration will ensure that AI is implemented for the benefit of all Americans.

Appendix

***AI Use Within the National Science Foundation**

This section was added as a supplement given its illustrative example and substance to the reports overall findings and recommendations.

NSF is exploring, and in some cases adopting, the use of AI for increased efficiencies in business processes and for increased effectiveness of program managers. Through the efforts of NSF Innovation Management Group, who developed a set of tools using NLP and a variety of algorithms, two of NSF's most critical and time-intensive business processes, suggesting reviewers and determining reviewer conflict of interest have been greatly enhanced. Of particular note is NSF's examinations of the potential of both Robotic Process Automation (RPA) and blockchain. Two key investments are being made in RPA; 1) providing access to tools for NSF staff to create RPAs to support their work and 2) to build capacity among NSF employees through training and the development of a community of practice. Eight RPAs are in place and operational at NSF. NSF is experimenting with the use of blockchain to optimize funding. The question being explored is can the use of blockchain in the proposal submission process make it possible for NSF to identify overlapping proposals both within NSF and cross-agency.

ENABLING AND ADVANCING AI RESEARCH AND DEVELOPMENT

NSF provides AI Research and Development (R&D) leadership in three ways: 1. Furthering research in and about AI, 2) Building an AI workforce, and 3. Ensuring an enabling infrastructure for AI.

Furthering Research In and About AI

Every year NSF invests about \$500m in research in AI, including investments in core areas of AI such as reasoning, language processing, knowledge representation, planning, machine learning, deep learning, and computer vision. About \$150m of this is core AI foundation research. The remainder supports use-inspired research that brings AI and other disciplines or sectors together to advance both areas. For example, in research that brings together biologists and computer scientists to extend new or existing techniques to solve biological problems. The most recent budget request for 2021 calls for increasing the budget from \$500M to \$868M. the overall budget is down, but the AI budget is up.

NSF also has several other programs and partnerships providing research support, including the AI Institutes Program, a year-on-year program funding Institutes that focus on one or more of six key themes. This program is envisioned to provide a blend of use-inspired and foundational research that is domain focused. NSF is seeking to scale this program to fund an Institute in every state so that each can tackle a range of domain-specific issues. For example, a single institute in the Southeast could focus on resilience capability with ML. The total cost of such a program, one institute per state, \$20 million per state, is approximately \$1 billion. This total represents a small percentage of USG's overall AI spend. Other programs include a collaboration with The Partnership on AI that brings together social scientists and computer scientists to support socio-technical AI systems and one with Intel focused on bringing AI and next generation wireless technologies and spectrum to explore the use of ML in assigning spectrum dynamically.

Building an AI Workforce

NSF's CS for US program, is focused on encouraging and exciting a diverse cadre of students to pursue advanced degrees in computer science and then specialize in AI. Today there are 35,000 undergraduate computer science majors in the US; this number rapidly falls off at the graduate level. New programs are being designed to incentivize US undergraduate computer science majors to pursue advanced degree programs. This program is a lever to address a whole host of science, security issues, and diversity issues and ensures the continuation of US leadership.

Ensuring an Enabling Infrastructure for AI

The NSF Office of Advanced Cyberinfrastructure plays a pivotal role in providing HPC, access to cloud computing resources and hybrid computing resources to specialize high-end computing for the real-time streaming data and machine learning and deep learning applications it is seeing today. For example, it funded a supercomputer, Frontera, to provide a leading-edge resource for AI discovery to the research community. Partnerships are leveraged in this area as well. For example, a cross-government partnership with DARPA is enabling research on real-time machine learning. Data of interest to DARPA, and others, streams from numerous sources simultaneously. To generate insights and predictions in real time, the hardware state must be reengineered. DARPA and NSF are working together on this challenge.

REFERENCES

ACM U.S. Public Policy Council and ACM Europe Policy Committee (2017, May 25), *Statement on Algorithmic Transparency and Accountability* https://www.acm.org/binaries/content/assets/public-policy/2017_joint_statement_algorithms.pdf

AI and Ethics, <https://www.springer.com/journal/43681/aims-and-scope>

American Society for Public Administration Code of Ethics, <https://www.aspanet.org/ASPA/Code-of-Ethics/ASPA/Code-of-Ethics/Code-of-Ethics.aspx>

Berg, Nate (2020, July 23) *COVID-19 has opened the floodgates for smart cities—whether we like it or not*, <https://www.fastcompany.com/90530580/covid-19-has-opened-the-floodgates-for-smart-cities-whether-we-like-it-or-not>

Brookings (<https://www.brookings.edu/research/what-is-artificial-intelligence/>)

Crawford, I. and Calo, R., “There is a blind spot in AI research,” *Nature*, vol. 538, no. 7625, pp. 311–313, 2016.

Crawford, K., “AI Now 2019 Report,” AI Now Institute, 2019

Executive Office of the President of the United States of America. “2016-2019 Progress Report: Advancing Artificial Intelligence R&D,” 2019. <https://www.nitrd.gov/groups/AI>.

Hagendorff, T, “The Ethics of AI Ethics -- An Evaluation of Guidelines,” *ArXiv190303425 Cs Stat*, Oct. 2019.

Investopedia (<https://www.investopedia.com/terms/a/artificial-intelligence-ai.asp>)

SAS (https://www.sas.com/en_us/insights/analytics/what-is-artificial-intelligence.html)

U.S. Department of Defense (2020, February 24), *DOD Adopts Ethical Principles for Artificial Intelligence*.

<https://www.defense.gov/Newsroom/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/>

Vought, Russell T. “MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES FROM,” January 2020. <https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american->.

Zhang, B. and Dafoe, A., “Artificial Intelligence: American Attitudes and Trends,” Social Science Research Network, Rochester, NY, SSRN Scholarly Paper ID 3312874, Jan. 2019

Working Group and Staff

Make Government AI Ready Working Group:

Alan Shark, Chair

Executive Director and Chief Executive Officer, Public Technology Institute (PTI); Associate Professor at George Mason University's Schar School of Policy and Government; Course Developer/Instructor at the Center for Government Services, Rutgers University; doctorate in public administration from the University of Southern California; MPA degree from Baruch College; U.S. Navy Seabee Veteran.

Erik Bergrud

Associate Vice President for External Relations, Park University, Parkville, Missouri. Former Positions with American Society for Public Administration: President, Senior Director of Program and Service Development, Senior for e-Organization Development, Director of Information Services, Director of Chapter/Section Relations.

Kaye Husbands Fealing

Professor and Chair, School of Public Policy, Georgia Institute of Technology; Senior Study Director, Center for National Statistics, National Academy of Sciences; Professor, Humphrey School of Public Affairs, University of Minnesota; Program Director, Social, Behavioral and Economic Sciences, National Science Foundation; William Brough Professor (also tenured full, associate, assistant professor), Economics Department, Williams College

James Hendler

Professor, Computer Science, Rensselaer Polytechnic Institute; Professor, Computer Science, University of Maryland; Program Mgr/Chief Scientist (IPA), Information Systems, Def Advanced Research Projects Agency (DARPA); Open Data Advisor, New York State (unpaid), NYS Government; Internet Web Expert, Data.gov project, IPA to GSA, working w/OSTP; Member Advisory Committee, Homeland Security Science and Technology Adv. Comm, DHS; Board Member, Board on Research Data and Information, Nat'l Acad Science, Engineering and Medicine; Director's Advisory Committee Member, Nat'l Security Directorate, Pacific Northwest National Laboratories

Theresa Pardo

Technical Information Analyst, Office of Computing Services, Union College; Coordinator of Academic Computing Services, Office of Computing Services, Union College; Assistant Director Academic Computing, Academic Computing Services, Union College; Director, Academic Computing Services, Siena College; Project Coordinator, Center for Technology in Government, University at Albany, State University of New York; Project Director, Center for Technology in Government, University at Albany, State University of New Yor; Deputy Director, Center for Technology in Government, University at Albany, State University of New York; Director, Center for Technology in Government, University at Albany, State University of New York; Professor, Rockefeller College, University at Albany

Doug Robinson

Executive Director, National Association of State Chief Information Officers and Account Executive, AMR Management Services. Former Executive Director, Policy & Customer Relations, Technology Office, Governor of Kentucky; Consultant, Information Technology Management; Former positions with Finance & Administration Cabinet: Executive Director, Kentucky Information Resources Management Commission; Executive Director, Kentucky Office of Geographic Information Systems. Information Resources Consultant, Commissioner's Office, Information Systems Department; Executive Director (Interim), Geographic Information Systems Office, Office of the Secretary, Natural Resources and Environmental Protection Cabinet; Former positions with University of Kentucky: Director (Acting), NASA Technology Applications Center; Technology Transfer Coordinator, NASA Technology Applications Center.

Special Program Support

Dr. Tonya Thornton Neaves

Director for Extramural Projects with George Mason University's Schar School of Policy and Government; Director for the Schar School's Centers on the Public Service; Assistant Professor and Coordinator for Emergency Management and Homeland Security graduate certificate in the George Mason University Master of Public Administration Program; Council Member for the George Mason University's Institute for a Sustainable Earth; Director for the Virginia Certified Public Manager Program; Visiting Professor at the National School of Administration and Magistracy in Cameroon; Director for the Mississippi Public Safety Data Laboratory at Mississippi State University's Social Science Research Center.

Tyler Goodwin

Masters of Public Policy student and Graduate Research Assistant at George Mason University. Bachelors of Arts in Political Science and Bachelors of Science in Economics. Former positions include Research Intern at The Institute for Human Rights and War Child USA.

Staff

Joseph P. Mitchell, III

Director of Strategic Initiatives and International Programs, National Academy of Public Administration; Member, National Science Foundation Business and Operations Advisory Committee; Associate Director, Office of Shared Services and Performance Improvement, General Services Administration; Director of Academy Programs, National Academy of Public Administration; Project Director, Senior Analyst, and Research Associate, National Academy of Public Administration.

James Higgins

Research Associate for Grand Challenges in Public Administration, National Academy of Public Administration; Researcher, Cohen Group; Extern, U.S. Patent and Trademark Office.

