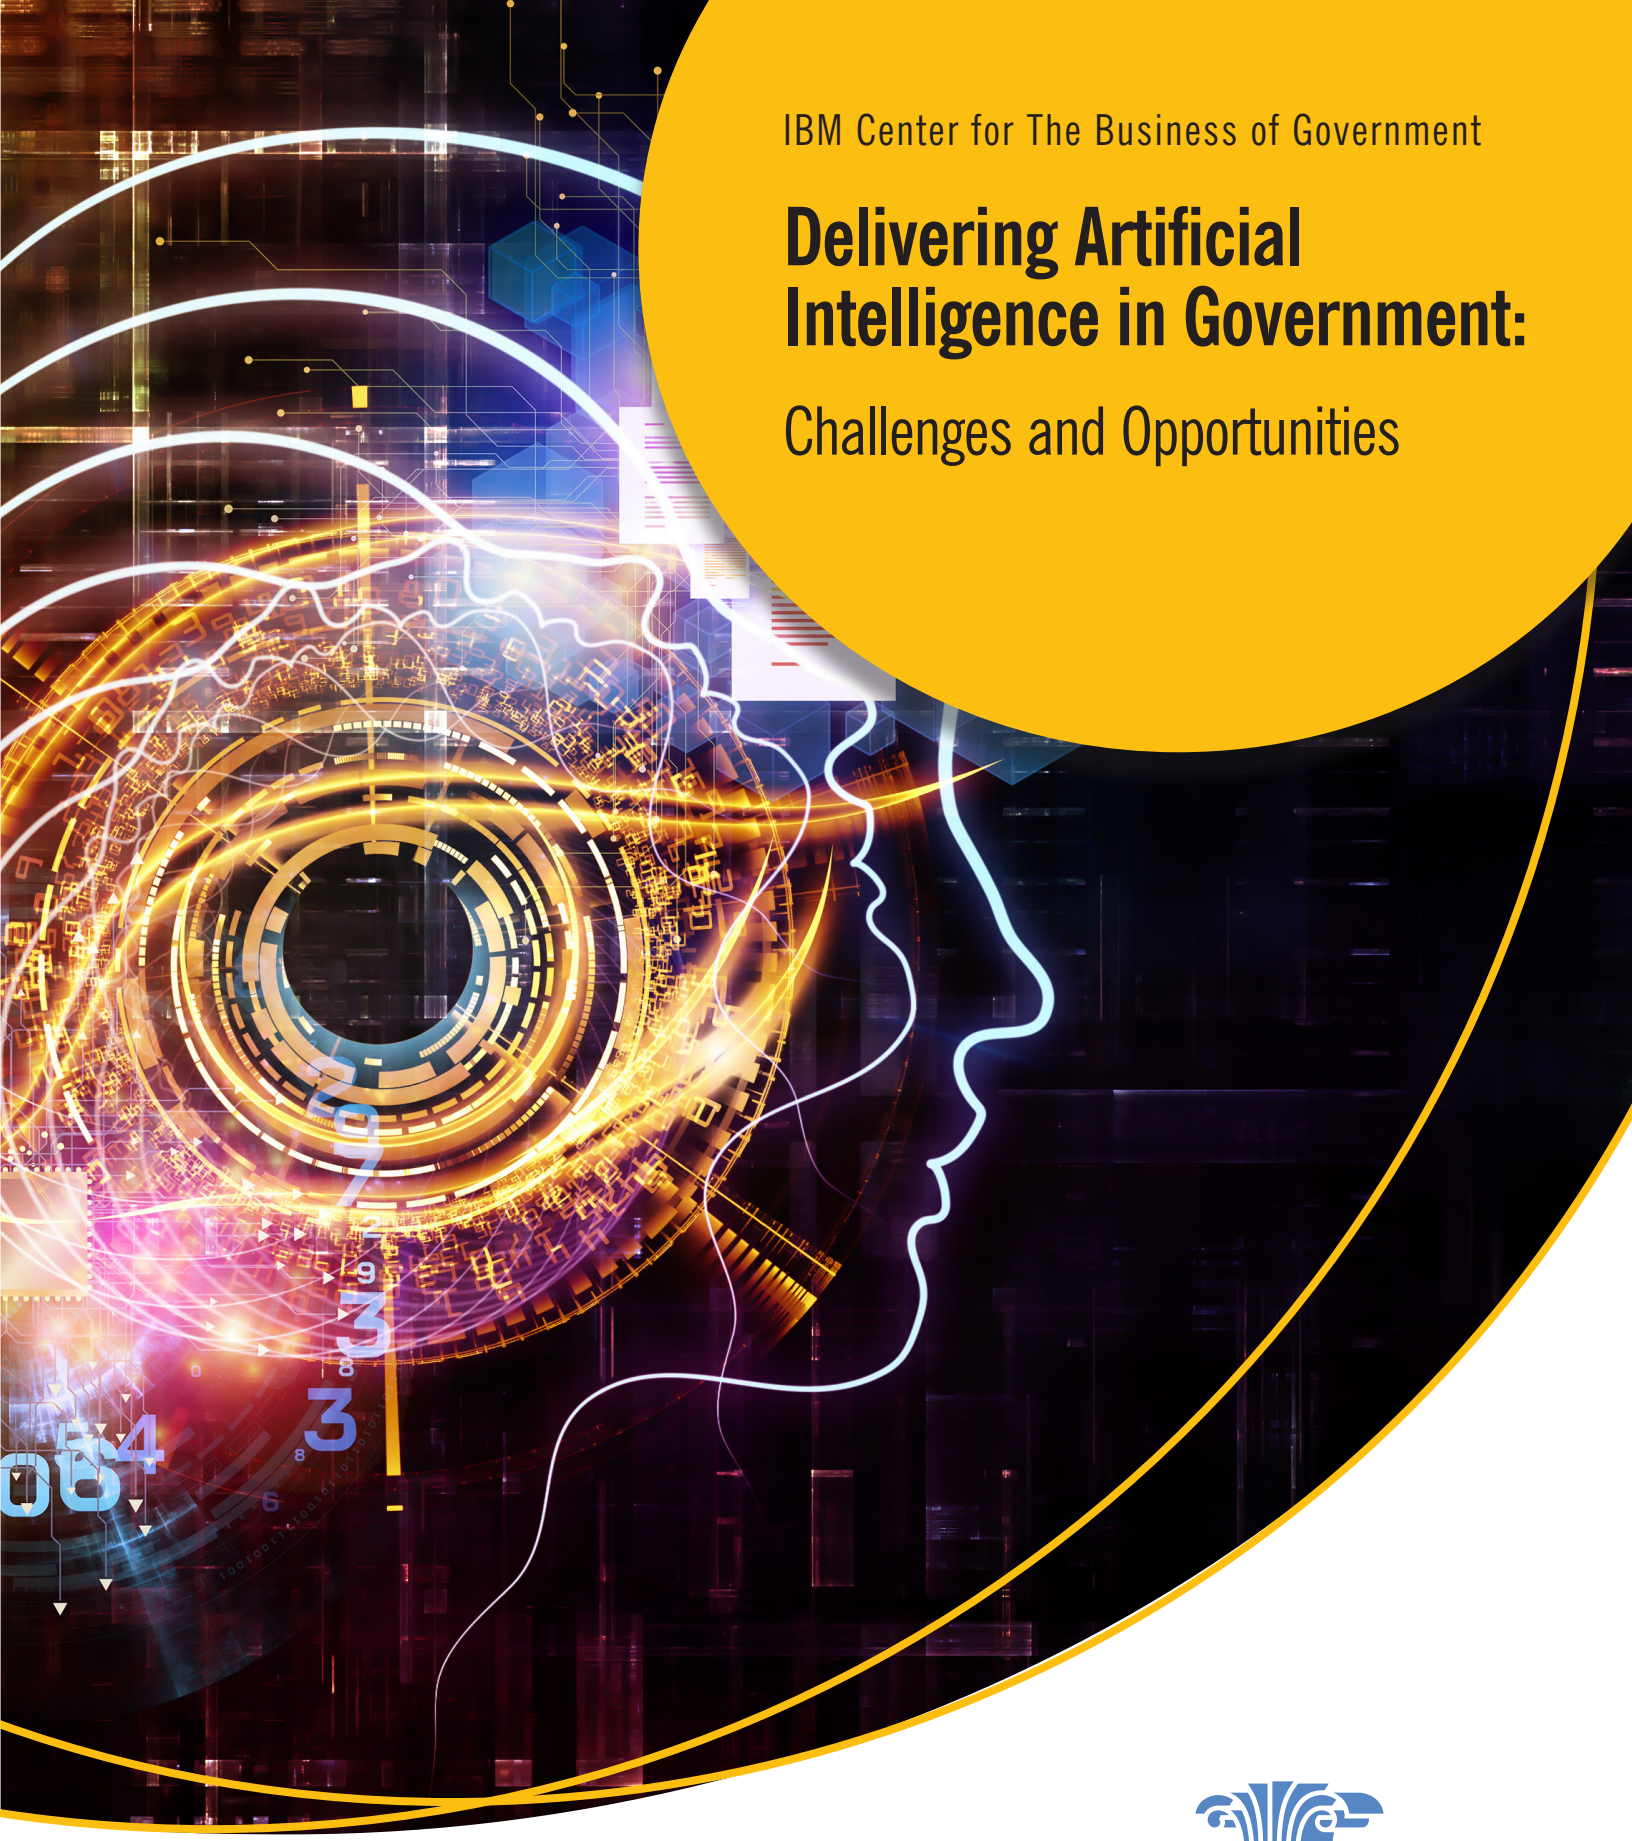

IBM Center for The Business of Government

Delivering Artificial Intelligence in Government:

Challenges and Opportunities

IBM Center for The Business of Government

Kevin C. Desouza
Arizona State University

2018

Delivering Artificial Intelligence in Government: Challenges and Opportunities

Kevin C. Desouza
Arizona State University

IBM Center for
The Business of Government

TABLE OF CONTENTS

FOREWORD

On behalf of the IBM Center for The Business of Government, we are pleased to present this report, *Delivering Artificial Intelligence in Government: Challenges and Opportunities,* by Kevin Desouza, ASU Foundation Professor in the School of Public Affairs at Arizona State University.

Professor Desouza reviews recent progress made in applying artificial intelligence to public sector service provision, drawing on lessons learned from commercial experience as well as burgeoning cognitive computing activity by Federal, State, local, and international governments. The author takes this real-world experience to set forth a framework for agencies to plan, develop, and deploy AI systems. He then puts forward a set of challenges for government leaders and innovators in this space, along with opportunities for agencies to act in addressing these challenges. Finally, Desouza outlines a maturity model for agencies to use in guiding their journey forward in applying AI to improve mission performance.

DANIEL J. CHENOK

Desouza frames his opportunities in three broad areas: technology and data, workforce, and risk management. In each area, agency leaders will find key factors that they can apply to increase the likelihood that emerging AI and cognitive applications will be implemented successfully. These factors include:

- Upgrading IT infrastructure to support AI systems, leveraging cloud computing technologies.
- Identifying data intensive applications that can benefit from AI, and establishing data governance to take best advantage of the benefits that AI can deliver.
- Enabling a skilled public sector workforce to use AI, including through agile implementations and redesigned work processes.
- Developing AI in a manner that augments human decisionmaking, and follows ethical imperatives around transparency, security, auditability, and citizen involvement.
- Working in partnership with government, academia, and industry.

CLAUDE YUSTI

The author builds on his prior work for the Center on IT management by CIOs and other public sector IT leaders, and follows on our recent report with the Partnership for Public Service that identifies case studies of government success in this space, *The Future Has Begun: Using Artificial Intelligence to Transform Government.*

We hope that agencies will find the practical and actionable steps offered in this report to be useful in capitalizing on the potential for AI to improve government.

Daniel J. Chenok
Executive Director
IBM Center for The Business of Government
chenokd@us.ibm.com

Claude Yusti
Partner
IBM Global Business Services
cayusti@us.ibm.com

EXECUTIVE SUMMARY

This report details the value proposition and success factors for enabling AI in the public sector.

Artificial Intelligence (AI) in government involves the design, building, use, and evaluation of cognitive computing and machine learning to improve the management of public agencies. To enable successful use of AI in government, leaders must design and implement governance and policy that promotes a skilled workforce that collaborates with academia and the private sector, risk management frameworks, secure systems, and modern technologies.

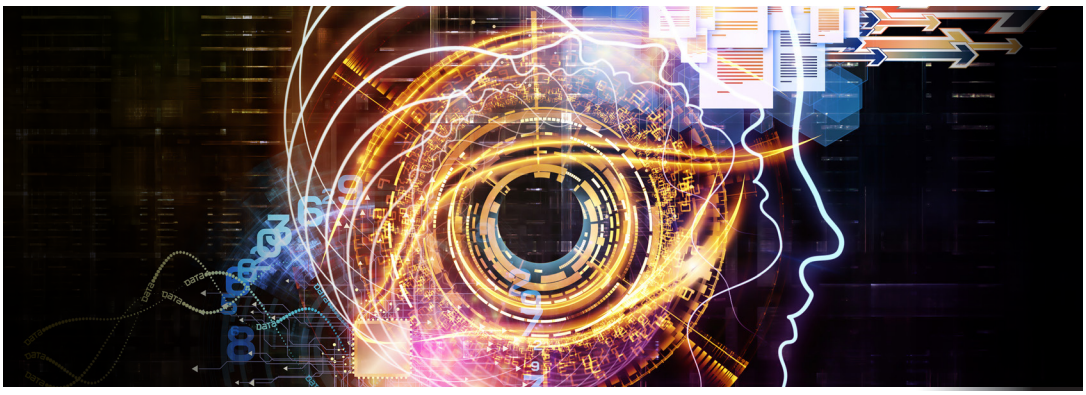

While public agencies are in their early days of experimenting with AI, these efforts are bound to intensify. To help government innovators progress in this area, the report outlines a maturity model for agencies to develop the capacity and capabilities to fully leverage AI's potential.

The report discusses the following challenges facing the public sector in implementing AI, with opportunities to address each challenge as summarized in the chart on the following page. It highlights practical steps that government leaders can take to leverage AI in ways that improves mission support and mission performance, addressing the subject in three sections:

- Planning, Developing, and Deploying AI
- Challenges and Opportunities in Delivering AI systems
- A Maturity Model for Assessing Public Sector Progress

Table 1: Challenges and Opportunities for Deploying AI in the Public Sector

Challenges	Opportunities
TECHNOLOGY AND DATA	
Legacy IT Infrastructure	• Invest in Upgrading and Modernizing the IT Infrastructure
Limited IT Interoperability	• Look at Cloud Computing and Open Source as Options to Get Started
IT Project Management Capabilities	• Use Agile Governance and Acquisition Practices for IT
Lack of Prioritization for Data-driven Solutions	• Identify Data Intensive Problems that Can Take Advantage of Machine Learning and Cognitive Capabilities • Invest in Data Governance • Deploy to Improve Service Delivery
WORKFORCE	
Public Workforce Management Practices	• Take Steps to Transform the Public Workforce to Take Advantage of AI Capabilities • Engage Human Experts in Designing, Testing, and Evaluating AI
Culture of IT Ownership	• Develop Collaborative Partnerships with Academia to Initiate AI Projects • Initiate Public-Private Partnerships to Design, Deploy, and Evaluate AI on Mission-Critical Priorities
Limited Capacity for System-level Redesign	• Redesign Work Processes to Increase Effectiveness and Efficiency • Employ AI to Augment Human Decision Making
RISK MANAGEMENT	
Securing Systems	• Develop Cybersecurity Capacity
Risk Aversion	• Promote Innovation Through Crowdsourcing Platforms • Increase Awareness of the Potential for AI Through Exploratory Projects • Develop Collaborative Partnerships Across Agencies on Cross-Cutting Issues
Ethical and Social Considerations	• Prioritize Value-Sensitive Design • Focus on Protecting Public Values and Common Good
Governance	• Proactively Monitor Systems to Track Unexpected Outcomes • Develop Robust Audit and Inspection Mechanisms

This report was developed by analyzing recent experience in the public and private sector in implementing AI systems, reviewing related literature, and offering a set of opportunities for action. The report generally references AI, the most widely recognized and adopted terminology, to reflect technologies that augment human abilities and in performing both simple and complex; related terms include cognitive computing, robotic process automation and machine learning.

INTRODUCTION

Artificial intelligence (AI) systems have come a long way from the days of residing exclusively in the domain of science fiction.

Key Advances in AI

Building on decades of progress outlined below, most technological systems today have some level of intelligence. Even now-ubiquitous mobile phones contain technologies such as voice recognition, can suggest responses to your incoming emails, and house social apps including Facebook and Twitter that use machine learning algorithms to decide what content to serve up.

AI-inspired systems are being deployed in almost every field, including healthcare, education, public safety, finance, international development, and the arts. Since 2012, more than 250 acquisitions have occurred in the AI space, including 37 in the first quarter of 2017 alone[1] (see next page). Interest in building better artificially intelligent systems is palpable.

AI systems apply a wide assortment of cognitive computing techniques to large volumes of data to augment human capacities and/or conduct autonomous decision making. These cognitive computing systems:

- learn in an iterative manner, from both data and human interactions, through the application of machine learning algorithms that continuously ingest data and build new knowledge and models based on the data
- interact with humans through natural language processing
- are context sensitive, having the ability to recall history and tailor outcomes in a personalized manner
- can provide confidence-weighted recommendations (outcomes) which can be acted upon by humans
- are adaptive (can adapt based on new data and information, and can tolerate ambiguity and incompleteness)
- are interactive (personalize interactions in natural language with users)
- have memory (remember previous interactions and can begin where something was left off)
- are iterative (work through a series of steps to solve a problem)
- are contextual (work within contexts and draw on environmental characteristics to deal with information requests; e.g., draw on user profiles and previous interactions)
- provide confidence-weighted responses (give people and systems a set of answers with weights)

1. CB Insights Research. "The Race For AI: Google, Baidu, Intel, Apple in A Rush to Grab Artificial Intelligence Startups." *CB Insights Research,* July 21, 2017. https://www.cbinsights.com/research/top-acquirers-ai-startups-ma-timeline/.

The public sector has started to engage with AI and cognitive computing. In October 2016, the US White House National Science and Technology Council (NSTC)[2] released a report that outlined the state of AI at that time, and its potential impact on society. The NSTC also published a companion report[3] outlining a research and development strategic plan for advancing AI. This report discussed several objectives to develop, track, and monitor AI-related R&D investments within the federal government. The report also discussed guidelines for managing federally funded research pertaining to AI but occurring outside the public sector, and outlined several recommendations to help federal governments develop near-term capabilities (e.g., consolidating databases, collaboration between humans and AI) and long-term capabilities (e.g., creating a sustainable tech-savvy workforce). NSTC provided guidelines to maximize the federal government's investment in AI and track progress to identify needs for developing an AI-driven society.

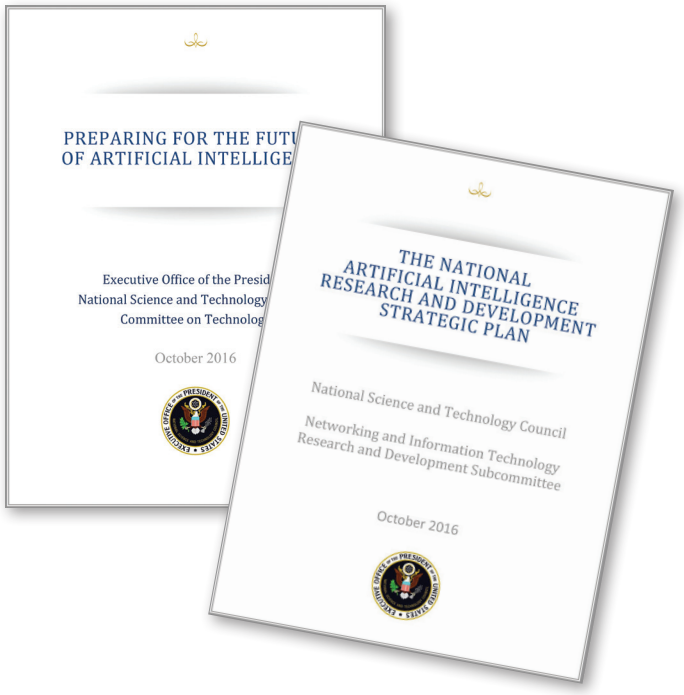

Source: *National Science and Technology Council (NSTC) 2016 report and companion report*

2. National Science and Technology Council. "Preparing for the Future of Artificial Intelligence." October 2016. https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.

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

MILESTONES IN AI

- In 1997, IBM's Deep Blue[4] defeated Garry Kasparov to become the first computer to beat a sitting world chess champion in a six-match game. Deep Blue could explore between 100 million and 200 million positions per second, running through possible moves and outcomes to figure out the best possible move. Deep Blue was trained to evaluate and learn—a significant improvement compared to its predecessor Deep Thought.[5] Computer scientists and grand masters trained Deep Blue to understand chess concepts and correct its weaknesses. They evaluated Deep Blue's weaknesses and developed solutions to improve its performance.[6]

- Stanford University's Stanley won DARPA's grand challenge for driverless robotic cars in 2005. The challenge required driverless cars to complete a 132-mile route that included dry lake beds, flats, and treacherous mountain passes.[7] In 2010, four driverless experimental cars completed an 8,000-mile test drive from Italy to China. These cars were equipped with solar-powered laser scans and cameras to detect obstacles and navigate traffic. The sensors also collected data for future research as these cars navigated through traffic from Italy to China.[8] The European Research Council funded this project with a €1.5 million (approximately $1.95 million) grant, in an effort to advance research on autonomous vehicles. Today, major companies[9] are experimenting with driverless and fully autonomous vehicles that require little to no human intervention.

- In 2011, IBM's Watson defeated two former winners on Jeopardy to win a $1 million prize. It took IBM more than four years to build Watson. Watson can process 80 million operations per second. Computer scientists fed 200 million pages of content, including structured and unstructured data, into Watson. To answer a question, Watson processes 6 million logic rules to identify the best answer.[10]

- Google's AlphaGo became the first computer to defeat Fan Hui, the European champion in the game of Go; the score was 5 to 0.[11] Go is a 2,500-year-old Eastern game that is considered more complex than chess. AlphaGo[12] employed a combination of supervised (where human experts trained the algorithm) and reinforcement learning techniques (where the algorithm learned from games of self-play). AlphaGo achieved a winning rate of 99.8 percent against other Go programs available in the market such as Crazy Stone and Zen.

4. Greenemeier, Larry. "20 Years after Deep Blue: How AI Has Advanced Since Conquering Chess." *Scientific American,* June 2, 2017. https://www.scientificamerican.com/article/20-years-after-deep-blue-how-ai-has-advanced-since-conquering-chess/.
5. A dual processor computer that took part in several chess competitions. Each processor of Deep Thought could analyze 450,000 positions per second; and 700,000 positions per second if the processors worked together. For more information see: Berliner, Hans. "Deep Thought Wins Fredkin Intermediate Prize." *AI Magazine* 10, no. 2, 1989.
6. Greenemeier. "20 Years after Deep Blue: How AI Has Advanced Since Conquering Chess."
7. Thrun, Sebastian. "Toward Robotic Cars." *Commun.* ACM 53, no. 4 (April 2010): 99–106, doi:10.1145/1721654.1721679.
8. Daily Mail Reporter. "Driverless Van's 8,000-Mile Test Drive from Italy to China." *Daily Mail Online*, October 28, 2010. http://www.dailymail.co.uk/sciencetech/article-1324515/Driverless-vans-8-000-mile-test-drive-Italy-China.html.
9. Marshal, Aarian. "Lyft Says, 'Me Too!' and Dives into the Self-Driving Game." *WIRED*, July 21, 2017. https://www.wired.com/story/lyft-self-driving-game/.
10. Jackson, Joab. "IBM Watson Vanquishes Human Jeopardy Foes." *PCWorld,* February 16, 2011. https://www.pcworld.com/article/219893/ibm_watson_vanquishes_human_jeopardy_foes.html.
11. Metz, Cade. "In a Huge Breakthrough, Google's AI Beats a Top Player at the Game of Go." *WIRED,* January 27, 2016. https://www.wired.com/2016/01/in-a-huge-breakthrough-googles-ai-beats-a-top-player-at-the-game-of-go/.
12. Silver, David, et al. "Mastering the Game of Go with Deep Neural Networks and Tree Search." *Nature* 529, no. 7587 (January 27, 2016): 484–89, doi:10.1038/nature16961.

MILESTONES IN AI (CONT.)

- In 2017, Carnegie Mellon University's AI Libratus[13] defeated four of the best poker players in the world and won $1,766,250 in chips. Libratus took part in a 20-day marathon poker competition—Brains vs. Artificial Intelligence: Upping the Ante, in Pittsburg. The algorithm was designed to analyze and learn from its mistakes. At the end of each day, the meta algorithm[14] analyzed weaknesses in its strategies that were exploited by poker players during the games. It learned from those mistakes and improved its performance. Libratus relied on Pittsburgh Supercomputing Center's Bridges computer that has 274 terabytes of memory[15] and is 7,250 times faster than high-end laptops.[16]

- Open AI's bot defeated professional Dota 2 player Danil Ishutin in 2017.[17] Dota 2 is a complex multiplayer online game, where teams use powerful characters (heroes) to battle each other. The game is over when a team destroys an opponent's Ancient structure. Open AI trained a bot with no prior knowledge of Dota 2 to learn the game through self-play. Within two weeks of training, the bot achieved an expert level equal to that of a professional Dota 2 player. AI's bot became the first bot to defeat a professional computer gamer.[18]

In June 2017, the ITU Telecommunication Standardization Sector (ITU-T) and XPRIZE Foundation, in collaboration with the United Nations (UN) International Children's Emergency Fund, the Office of the UN High Commissioner for Human Rights, the UN Global Pulse, and the UN Office on Drugs and Crime, hosted AI for Good Global Summit in Geneva.[19] This summit brought together public officials, UN agencies, businesses, civil society organizations, and AI experts to deliberate on guidelines and recommendations for promoting AI-based innovations that could contribute to the attainment of UN Sustainable Development Goals (SDGs) including poverty, hunger, well-being, quality education, gender equality, climate action, sustainable cities, etc.[20] Participants also discussed developing AI Guiding Principles[21] to address issues related to the ethical, social, and economic challenges of leveraging AI-based technologies.

The European Parliament's Committee on Legal Affairs issued a report on the policy implications of AI in 2017.[22] The report outlined a policy framework to govern the use of AI and how to protect society from AI-related issues. The report discussed issues related to data governance, liability, ethics, etc., as they are key components in the AI revolution. The European Union (EU) established a Public-Private Partnership in Robotics to provide a collaborative platform for academics and industries to develop a roadmap for leveraging and managing robotics

13. Popper, Ben. "AI Has Definitively Bested Humans at Poker." *The Verge,* January 31, 2017. https://www.theverge.com/2017/1/31/14451616/ai-libratus-beat-humans-poker-cmu-tournament.
14. Kurzweil AI. "Carnegie Mellon AI Beats Top Poker Pros — a First." *KurzweilAI,* accessed August 9, 2017. http://www.kurzweilai.net/carnegie-mellon-ai-beats-top-poker-pros-a-first.
15. PSC. "Bridges," *PSC,* accessed August 9, 2017. https://www.psc.edu/index.php/resources/computing/bridges.
16. Popper, "AI Has Definitively Bested Humans at Poker."
17. Weinberger, Matt. "Elon Musk's $1 Billion AI Startup Made a Surprise Appearance at a $24 Million Video Game Tournament — and Crushed a pro Gamer." *Business Insider,* August 11, 2017. http://www.businessinsider.com/the-international-dota-2-openai-bot-beats-dendi-2017-8.
18. Wattles, Jackie. "A Bot Just Defeated One of the World's Best Video Gamers." *CNNMoney,* August 12, 2017. http://money.cnn.com/2017/08/12/technology/future/elon-musk-ai-dota-2/index.html.
19. ITU-T. "AI for Good Global Summit." *ITU,* 2017. https://www.itu.int/en/ITU-T/AI/Pages/201706-default.aspx.
20. XPRIZE. "AI Solving Sustainable Development Goals." *IBM Watson A.I. XPRIZE,* May 8, 2017. https://ai.xprize.org/AI-For-Good/sustainable-development-goals; Frederic. "An Account of the AI for GOOD Global Summit, Geneva – CIONET | Blog." June 13, 2017. http://blog.cionet.com/2017/06/13/an-account-of-the-ai-for-good-global-summit-geneva/.
21. XPRIZE. "AI for Good." *IBM Watson A.I. XPRIZE,* May 8, 2017. http://ai.xprize.org/AI-For-Good.
22. European Parliament Committee on Legal Affairs. "European Parliament Resolution of 16 February 2017 with Recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL))." 2017. http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//NONSGML+COMPARL+PE-582.443+01+DOC+PDF+V0//EN.

in Europe.[23] This platform seeks to promote testing prototypes to further robotic research. Moreover, the EU launched Building a European Data Economy[24] to facilitate the free flow of non-personal data. Simultaneously, this initiative focuses on governing issues related to data liability, misuse, and ownership.

Applying AI to Help Government

AI in government involves the design, building, use, and evaluation of cognitive computing and machine learning to improve the management of public agencies, the decisions leaders make in designing and implementing public policies, and associated governance mechanisms. The public sector has begun exploring the use of AI across myriad domains. The report generally references AI, the most widely recognized and adopted terminology, to reflect technologies that augment human abilities and in performing both simple and complex; related terms include cognitive computing, robotic process automation and machine learning. Some examples follow.

Through AI and cognitive computing, government agencies will be able to augment the capabilities of their workforce by processing and learning from large amounts of disparate data across heterogeneous systems in near real time, and interacting with humans through natural language processing. Cognitive systems will allow public agencies to free up their workforce from tasks that are ripe for automation due to their structured and predictable nature. For example, in 2015, US Citizenship and Immigration Services (USCIS) launched a chatbot named "Emma" to handle a high volume of customer queries.[25] On average, USCIS receives about 14 million calls on immigration issues each year. Emma was designed to help visitors navigate the USCIS website and find relevant information easily. Emma can provide answers to routine questions. To answer a question, Emma searches through the USCIS website and provides short answers and links to contents. If Emma is unable to offer a reply, she directs the customer to a live agent.[26] Emma can type answers in English and Spanish and speak in English.[27]

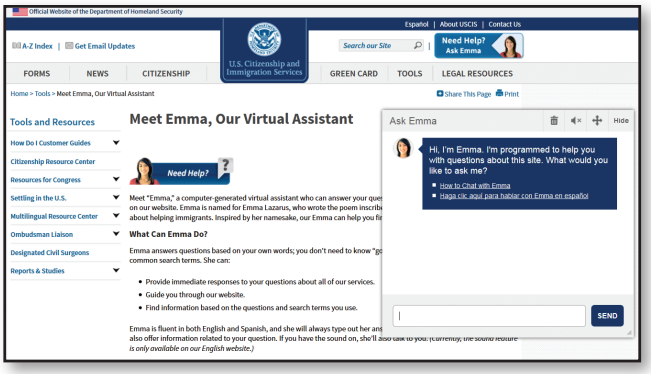

Source: US Citizenship and Immigration Services (USCIS) chatbot "Emma". https://www.uscis.gov/emma.

23. European Commission. "Robotics Public-Private Partnership in Horizon 2020." *Digital Single Market,* January 28, 2015. https://ec.europa.eu/digital-single-market/en/robotics-public-private-partnership-horizon-2020.
24. European Commission. "Building a European Data Economy." *Digital Single Market,* January 10, 2017. https://ec.europa.eu/digital-single-market/en/policies/building-european-data-economy.
25. USCIS. "USCIS Launches Virtual Assistant - Emma Gives Customers Another Option for Finding Answers." *USCIS,* December 2, 2015. https://www.uscis.gov/news/uscis-launches-virtual-assistant-emma-gives-customers-another-option-finding-answers.
26. Department of Homeland Security. "Privacy Impact Assessment for the Live Chat DHS/USCIS/PIA-065." May 19, 2017. https://www.dhs.gov/sites/default/files/publications/privacy-pia-uscis-livechat-may2017.pdf.
27. USCIS. "Meet Emma, Our Virtual Assistant." *USCIS,* 17 2017. https://www.uscis.gov/emma.

AI systems will also enable agencies to increase the speed with which transactions are processed, by reprioritizing the assignment of tasks between humans and machines. In another example, Hong Kong's Immigration Department was processing more than 4 million visa applications in 2004. The department processed about 100 different application forms related to visas, travel documents, identity cards, and other issues. To streamline the process, in 2007 the Hong Kong Immigration Department invested in developing an algorithmic system[28] which sorted passport applications and was trained to classify applications into three broad categories: approved, denied, and a gray area. Once the algorithm classified visa applications, the system would transfer the data to a visa officer, who reviewed the documents to make the final decision.

AI can also help accelerate innovation in public policy decision frameworks. Cognitive systems support complex decision-making scenarios in real time, due to their ability to work with larger amounts of data and connect dots at a far greater rate of effectiveness and efficiency. AI helps governments track the efficacy of policy interventions and recommend course corrections as needed. For example, the City of Pittsburgh collaborated with researchers from Carnegie Mellon University to develop SURTRAC (Scalable Urban Traffic Control) to manage traffic flow in real time.[29] SURTRAC manages multiple competing traffic flows shifts through multiple intersections, where each intersection has a virtual agent that controls traffic signals and dynamically coordinates with other agents in real time to monitor traffic flows. The deployment of SURTRAC resulted in a 25 percent reduction in travel time, 34 percent increase in vehicle speed, 31 percent decrease in traffic stops, 40 percent reduction in wait time, and 21 percent decrease in emissions.

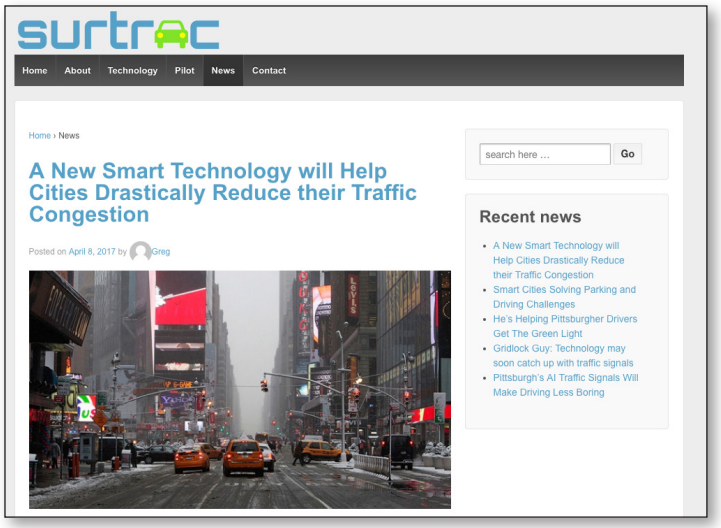

Source: *SURTRAC (Scalable Urban Traffic Control)* https://www.surtrac.net/news/

28. Chun, Dr. H. W., "AI Success Stories." accessed August 10, 2017. http://www.cs.cityu.edu.hk/~hwchun/AIProjects/stories/km/ebrain/.

29. Smith, Stephen F., et al. "Smart Urban Signal Networks: Initial Application of the SURTRAC Adaptive Traffic Signal Control System." *ICAPS* (Citeseer, 2013). http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.433.5935&rep=rep1&type=pdf.

Similarly, the Atlanta Fire Rescue Department (AFRD) collaborated with researchers from Georgia Institute of Technology, Emory University, and University of California, Irvine to develop an algorithm to predict buildings at risk of fire.[30] The algorithm was trained using data from 2010–2014 and included 58 variables (e.g., property location, fire incident, size, building structure, and year built) to predict fire risk buildings. Using this information, the algorithm forecasted fire risk scores for 5,000 buildings and identified 19,397 new properties to inspect. The algorithm accurately predicted 73 percent of fire incidents.

The IBM Center for The Business of Government recently collaborated with the Partnership for Public Service on an issue brief entitled *The Future Has Begun: Using Artificial Intelligence to Transform Government*,[31] which identified government innovators moving forward to implement AI their programs. Case studies highlighted in this report include the Department of Labor's Bureau of Labor Statistics to streamline tedious work, the Air Force's Procurement office to promote better decisions, and local and international examples of AI to improve law enforcement outcomes. These case studies distill the findings from over a dozen interviews with thought leaders applying AI in government, demonstrating that government can use AI to solve real issues in delivering on their mission.

Source: *The IBM Center for The Business of Government and Partnership for Public Service*

30. Madaio, Michael, et al. "Identifying and Prioritizing Fire Inspections: A Case Study of Predicting Fire Risk in Atlanta." 2015. https://www.cc.gatech.edu/~bdilkina/papers/madaio2015identifying.pdf.
31. http://www.businessofgovernment.org/report/using-artificial-intelligence-transform-government

Planning, Developing, and Deploying AI in the Public Sector

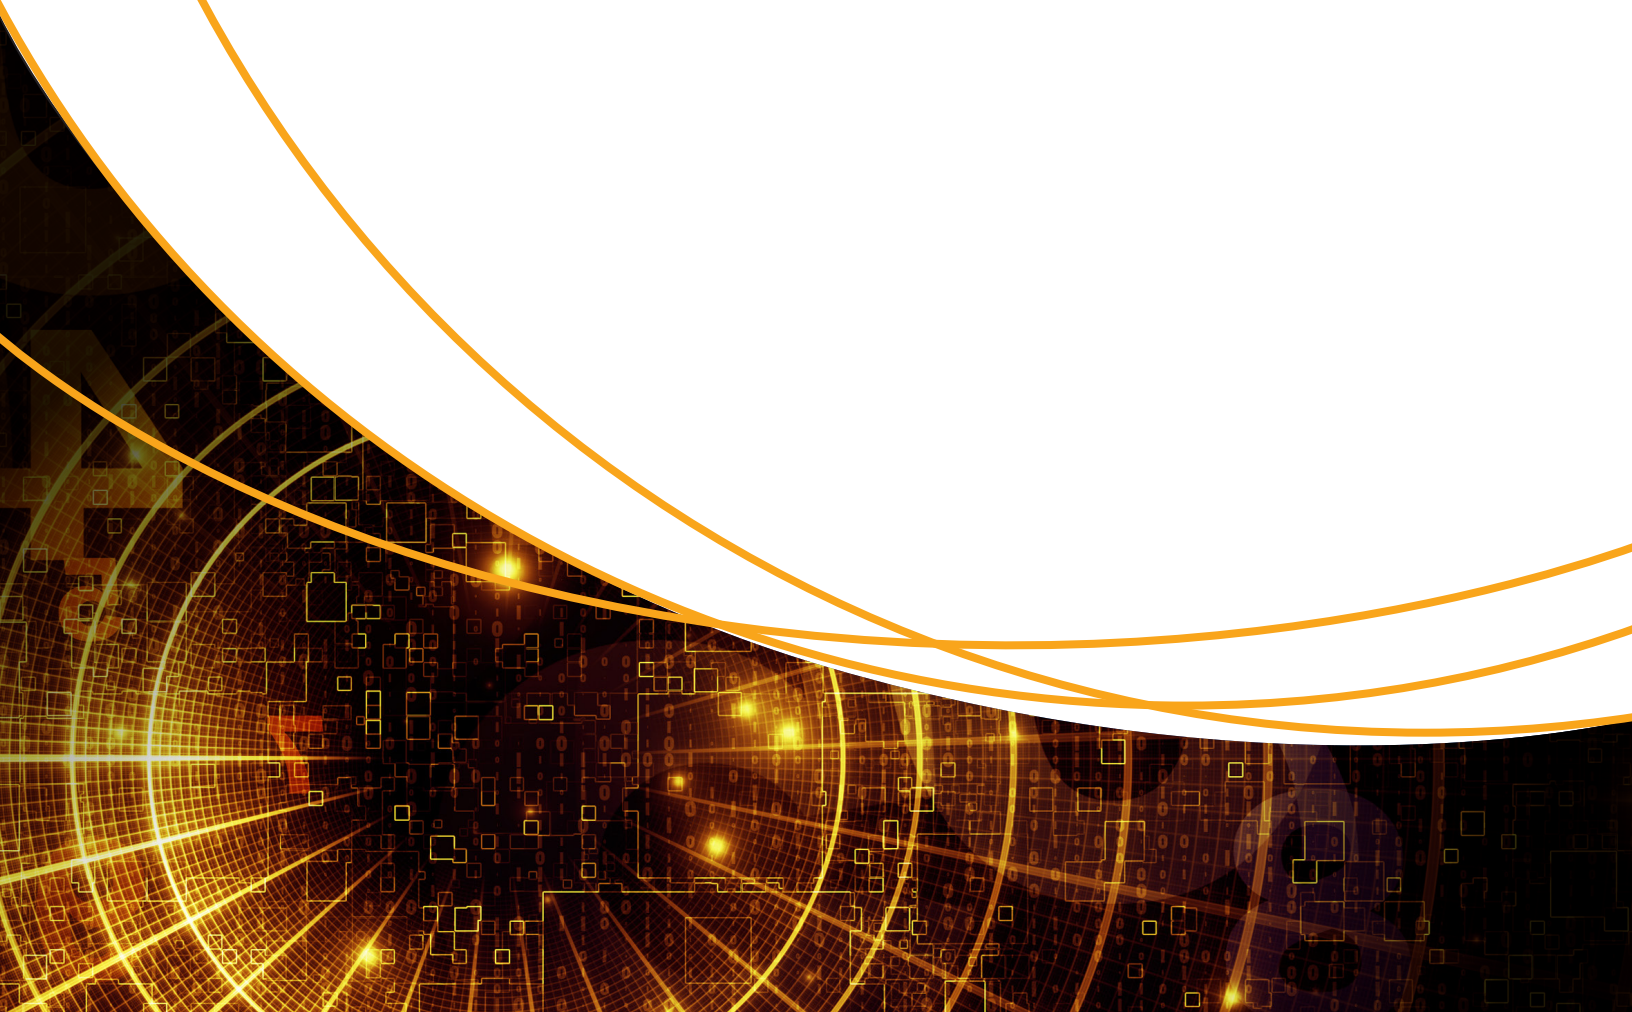

AI systems rely on cognitive computing capabilities to learn and develop even after they are deployed, through ingestion of more data and interactions with humans. In addition, these systems are deployed in various stages before they can function autonomously. However, like traditional information systems, organizations must think through the business case, manage the project carefully, and ensure that metrics are in place to track performance and outcomes so that refinements can be made over time. In this section, a few key issues that need to be considered as they pertain to AI systems are highlighted in three distinct phases—*planning, developing,* and *deploying.*

Planning

It is important to begin with a strong business case for AI. A high-level business case must outline:

1) how the systems will result in efficiencies and innovations in how work and processes are executed,

2) the data and analytical opportunities that make AI a viable solution, and

3) the key performance indicators that will be used to evaluate the systems and the impact on the organization.

AI systems leverage cognitive to tackle challenges and realize opportunities where computational tools can enable workers to analyze large-scale data meaningfully within the context of their work. Take the case of cancer research and treatment: The American Cancer Society projected that 1,688,780 new cancer cases would be diagnosed and 600,920 people would likely die of cancer in the US in 2017.[32] It estimated that 5 percent of patients who seek outpatient services receive incomplete or, worse, wrong diagnoses, which presents a significant blind spot in healthcare.[33] Moreover, it is projected that by 2020 healthcare literature will double every two months.[34] Reviewing and synthesizing this literature is beyond human cognition. At the same time, the large volume of healthcare literature offers a significant opportunity for developing high-quality and personalized healthcare products and treatments. In cases such as this, cognitive applications can be developed to learn from the data and to provide recommendations to physicians to augment their capabilities.

The availability of data is important, as these systems need to learn by building associations across concepts and elements found in these datasets. A lack of sufficient datasets will limit the ability of the algorithms to learn. Moreover, while having large-scale data is important, large-scale data in a structured or even semi-structured format improves the analytical benefits of cognitive systems.

Development

Developing AI systems should focus on organizing data, training systems, and testing performance. Organizations can create mechanisms to collect, combine, and prepare data across multiple sources to feed into the systems. Machine learning tools and algorithms can help identify patterns in the dataset and build predictive models that human experts can train and

32. American Cancer Society. "Cancer Facts & Figures 2017." American Cancer Society, Atlanta, 2017. https://www.cancer.org/research/cancer-facts-statistics/all-cancer-facts-figures/cancer-facts-figures-2017.html.

33. Institute of Medicine. "Improving Diagnosis in Health Care." September 2015. http://www.nationalacademies.org/hmd/~/media/Files/Report Files/2015/Improving-Diagnosis/DiagnosticError_ReportBrief.pdf.

34. Akhter, Shahid. "IBM Watson Health: Cognitive Computing Can Drive Tremendous Value into the Healthcare Industry: Robert Merkel." *ET HealthWorld,* September 22, 2016. https://health.economictimes.indiatimes.com/news/industry/cognitive-computing-can-drive-tremendous-value-into-the-healthcare-industry-robert-merkel/54457320.

test—for example, to assess the accuracy of solutions provided to known problem and solution pairs. In cases where the system gets the answer wrong, experts will provide the correct answer, which will then be used by the system in future iterations—improving overall accuracy.

In a recent case, the Las Vegas Health Department collaborated with researchers from the University of Rochester to experiment with nEmesis to address the problem of restaurant inspection.[35] nEmesis leverages machine learning and analyzes tweets to develop a list of problematic restaurants to inspect. The researchers conducted a controlled experiment where half of the inspections were performed using nEmesis and the other half were performed following traditional random inspections protocols. The results indicated that nEmesis was 64 percent more effective at identifying problematic restaurants than traditional random inspections protocols. nEmesis also identified restaurants that lacked permits or had contagious kitchen staff. The use of nEmesis helped the Las Vegas Health Department effectively use their limited resources to address their restaurant inspection problem.

In the development phase, it is important to feed the system sufficient and good-quality data that has balanced representation on the various elements of interest. A system to process visa applications, for example, needs datasets that are not skewed towards particular nationalities with a certain profile resulting in a bias toward a given outcome.

In addition, developers of cognitive systems should work with domain experts to embed expert knowledge. This normally happens by reviewing cases where the system has gotten decisions wrong or has provided useless recommendations, then identifying root causes and working with experts to tweak the system. This process is iterative, time consuming, and necessary.

Preparing to move from development to deployment requires defining an acceptable level of performance for a given topic. A chatbot to guide users through a government website may do fine with 80 percent accuracy, where a human expert can handle queries not answered or answered incorrectly. However, an AI-powered chatbot to support medical decision-making or assist with law enforcement calls for a much higher performance threshold due to the potential for harm caused by a faulty system.

Deployment

In initial deployment, cognitive systems support people who track, monitor, and test their performance and usability. Outputs are matched with the work of human experts to ensure systems meet performance expectations. The human decision maker retains control and can accept or reject recommendations—akin to having automated warning systems on an aircraft that alert a pilot about potential dangers and issues, but where the pilot remains the ultimate decision-maker.

Cognitive systems can then take an autopilot form within expected bounds. For example, the pilot may choose to handle the controls during takeoff, landing, or during turbulence, and have the system guide the plane when cruising. The pilot can override the system at any time and regain control of the aircraft. In another example, using AI to help process visa applications, the algorithm can sort applications into different buckets based on risk factors. Most applications in the low or medium risk categories can be processed by the system, while people have more time to inspect the remaining high-risk applications. If the system needs to be trained

35. Sadilek, Adam, et al. "Deploying nEmesis: Preventing Foodborne Illness by Data Mining Social Media." Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI Press, 2016, 3982–89. http://dl.acm.org/citation.cfm?id=3016387.3016466.

with new information for sorting applications, or there are changes to the environmental conditions (e.g., terrorist attacks), human experts can halt the use of the system as necessary.

In a final stage, cognitive systems can run autonomously and interact with humans or other systems. This level of operations requires designers, developers, and users to have high trust in the system and its capabilities. Human experts will regularly monitor, evaluate, and upgrade these systems to ensure that they stay current and perform at adequate levels, intervening when necessary. As discussed later, monitoring cognitive systems is critical to obtain early signals when performance falls below expected standards. Any system updates should be well documented and auditable. This enables sharing results and lessons learned from the deployment with relevant stakeholders.

Challenges and Opportunities for Government in Implementing AI

Public agencies face challenges when it comes to leveraging AI. Some of these challenges are specific to this new technology, while others are broader and are concerned with the general state of information technology capabilities in the public sector. Yet, these broader challenges are foundational, in that an agency has limited potential to enter the AI space unless efforts are invested to address them. We identify eleven key challenges for successful AI implementation, which can be grouped into three broad categories:

- Technology and Data
- Workforce
- Risk Management

Based on our assessment of public and private best practices, and lessons learned from implementing emerging IT systems in the public and private sectors, we offer opportunities to address each challenge area, totaling twenty recommendations overall. In addition, we propose an AI Maturity Framework for agencies to follow in implementing these recommendations as an appendix to this report; this Framework can serve as a guide to agencies regarding what actions to take based on their level of readiness to adopt AI systems.

Table 1 outlines these challenges and opportunities to help government move forward in this critically important technology arena.

Table 1: Challenges and Opportunities for Deploying AI in the Public Sector

Challenges	Opportunities
TECHNOLOGY AND DATA	
Legacy IT Infrastructure	• Invest in Upgrading and Modernizing the IT Infrastructure
Limited IT Interoperability	• Look at Cloud Computing and Open Source as Options to Get Started
IT Project Management Capabilities	• Use Agile Governance and Acquisition Practices for IT
Lack of Prioritization for Data-driven Solutions	• Identify Data Intensive Problems that Can Take Advantage of Machine Learning and Cognitive Capabilities • Invest in Data Governance • Deploy to Improve Service Delivery
WORKFORCE	
Public Workforce Management Practices	• Take Steps to Transform the Public Workforce to Take Advantage of AI Capabilities • Engage Human Experts in Designing, Testing, and Evaluating AI
Culture of IT Ownership	• Develop Collaborative Partnerships with Academia to Initiate AI Projects • Initiate Public-Private Partnerships to Design, Deploy, and Evaluate AI on Mission-Critical Priorities
Limited Capacity for System-level Redesign	• Redesign Work Processes to Increase Effectiveness and Efficiency • Employ AI to Augment Human Decision Making
RISK MANAGEMENT	
Securing Systems	• Develop Cybersecurity Capacity
Risk Aversion	• Promote Innovation Through Crowdsourcing Platforms • Increase Awareness of the Potential for AI Through Exploratory Projects • Develop Collaborative Partnerships Across Agencies on Cross-Cutting Issues
Ethical and Social Considerations	• Prioritize Value-Sensitive Design • Focus on Protecting Public Values and Common Good
Governance	• Proactively Monitor Systems to Track Unexpected Outcomes • Develop Robust Audit and Inspection Mechanisms

Technology and Data

Challenge 1: Legacy IT Infrastructure

The public sector continues to be plagued by an aging and outdated information technology (IT) infrastructure. Most agencies spend a significant portion of their IT budgets on maintenance of IT systems, and while modernization activities are increasing they are still comparatively limited. The Government Accountability Office (GAO) analyzed the operations and maintenance (O&M) spending of 10 federal IT investments with the largest budgets in 2012.[36] Eight agencies operate these 10 IT investments. Of the eight agencies, seven agencies reported that they invested $7.4 billion in O&M spending, which often occurred without developing policies to assess the performance of legacy IT systems. Several of the IT systems utilized in federal agencies are outdated, in some cases more than 50 years old, and incur significant maintenance costs.

Legacy system examples that GAO found include:

- The Department of Defense still uses 8-inch floppy disk systems, which run on a 1970s computing system to coordinate the operational functioning of nuclear forces.[37]
- The Department of Treasury uses assembly code language to assess and manage individual and business taxes.
- The Department of Veterans Affairs uses a Common Business Oriented Language (COBOL) system, about 51 years old, to track veterans claims for benefits, eligibility, and death records.[38]

The cost of continued maintenance for these legacy systems will continue to rise over time, while limiting how much budget can be spent on AI and other emerging technologies. In addition, these systems when built were specifically designed to process individual transactions rather than gather program data at scale, which limits an agency's ability to use these systems to deliver richer insights. Government needs more focus on conducting performance measures to understand the state of its current IT infrastructure, what needs to be modernized, and the cost of continued maintenance versus the cost of designing and implementing new modernized systems.[39]

36. US Government Accountability Office. "Agencies Need to Strengthen Oversight of Multibillion Dollar Investments in Operations and Maintenance." no. GAO-14-66, November 6, 2013. https://www.gao.gov/products/GAO-14-66.
37. US Government Accountability Office. "Federal Agencies Need to Address Aging Legacy Systems." no. GAO-16-468, May 25, 2016. https://www.gao.gov/products/GAO-16-468.
38. Ibid.
39. Greg Dawson. "Modernizing Government IT." IBM Center for the Business of Government, 2018 (forthcoming).

Opportunity 1: Invest in Upgrading and Modernizing the IT Infrastructure

Public agencies must inspect their IT systems to evaluate how to replace, modify, and retire systems to accommodate modern systems that provide a platform to develop and deploy AI. The US Federal Communications Commission (FCC), for instance, spent 80 percent of its IT budget on O&M of legacy systems in 2013. To modernize its IT systems, the FCC conducted an inventory to understand the health of these systems. Upon analysis, the agency realized that they had 207 IT legacy systems. The inventory of IT systems helped the agency decide how to consolidate these systems (e.g., retire or modify) and move to cloud computing to save costs. By 2015, the agency was spending only 50 percent of its budget on O&M of IT systems.[40]

In a survey of state CIOs, "90% of the CIOs considered at least 20% of their systems due for replacement or modernization, while nearly two-thirds of CIOs saw more than 40% of the systems as a legacy. As one CIO put it "We're coming out of a period of time where we weren't investing in technology; we have a lot of catching up to do."[41] And a recent NASCIO survey noted that 75% of the state CIOs surveyed reported they are seeing results in terms of lower asset maintenance cost when moving to cloud services, thereby providing resources to innovate.[42] Once public agencies understand the landscape of their IT infrastructure, they can then decide how to modernize legacy systems—replace, modify, and retire—for developing platforms to deploy AI systems.

While public agencies can deploy AI as an add-on to their existing portfolio of legacy systems, this is not the preferred option. Over time, the linking/adding-on of disparate new applications will limit interoperability and therefore impede an agency's ability to seek maximum value. The adoption of AI systems represents an opportunity to take a serious look at the current IT infrastructure and to think holistically about starting a modernization effort.

Challenge 2: Limited IT Interoperability

The siloed nature of IT systems in the public sector makes it a challenge to integrate data across systems on key thematic issues (e.g., homelessness at the local level, sustainability at the state level, or even STEM education policies at the national level). Agency-specific IT and data governance protocols are often misaligned. In addition, the nature of securing and utilizing public funding incentivizes each agency to focus internally and often compete with peers, rather than work collaboratively on larger system-wide issues. Many public agencies build stand-alone systems that have little potential for consolidation or, worse, are duplicative of other systems. For example, in 2013, the US GAO reported that the Department of Homeland Security had invested in two different IT systems that support immigration processing; the Department of Defense had invested in two systems for tracking the health status of warfighters and two others for maintaining dental care; and the Department of Health and Human Services has invested in six potentially duplicative systems—four to support the organization's information security and two to manage Medicare determination.

This lack of interoperability limits how agencies can integrate multiple databases that machine learning algorithms can then analyze and use to provide richer insights. Lack of interoperability also forces each agency to consider piecemeal solutions when working on acquisition of

40. Bray, David A. "Case in Point: Building an Agile Workforce and Enterprise at the FCC | LNW Community of Practice." *The 2015 Public Sector for the Future Summit,* 2015. http://lnwprograms.org/case-point-building-agile-workforce-and-enterprise-fcc.
41. NASCIO, Grant Thornton, and CompTIA. "The 2016 State CIO Survey." September 2016. https://www.grantthornton.com/~/media/content-page-files/public-sector/pdfs/surveys/2016/State-CIO-Survey.ashx.
42. NASCIO, Grant Thornton, and CompTIA. "The 2017 State CIO Survey." October 2017. https://www.nascio.org/Portals/0/Documents/2017/NASCIO_2017_State_CIO_Survey.pdf.

new technologies, increasing cost and maintenance over time. Even if agencies try to combine systems, there is often not enough common context to combine and utilize data from different sources. Agencies need to think about developing guidelines as to how data can be shared and used, and develop collaborative sharing networks. Governments can also give thought to collaborating with the private sector on analysis, as is now done with weather and public transportation agencies.

Opportunity 2: Look at Cloud Computing and Open Source as Options to Get Started

Public agencies can look at cloud computing and data management as starting points to deploy AI for their holistic data needs. In this process, open-source solutions and commercial platforms that can connect via APIs can provide a rapid way to move forward. The DoD, for instance, wants to further enhance its Biosurveillance Ecosystem (BSVE).[43] The DoD's Defense Threat Reduction Agency (DTRA) developed BSVE to integrate data sources for detecting and investigating biological threats. BSVE is a cloud-based open-source platform that allows users to customize their BSVE Analyst Workbench. Further, users can download Software Developers Kit (SDK) to integrate external data into BSVE for running analysis. To continuously improve BSVE, the DoD allows users to test BSVE's functionalities and provide feedback. In July 2017, DoD announced a request to solicit feedback for improving the capabilities of BSVE. The solicitation asks participants to suggest innovative ways to further improve BSVE's data framework, software maintenance, next-generation SDK, and to validate accreditation efforts.[44] DTRA continuously works on developing and expanding BSVE's capabilities to integrate and respond to new analytics and data sources.[45]

Challenge 3: IT Project Management Capabilities

Public agencies have a mixed track record when it comes to designing, managing, and executing long-term IT projects. Many IT projects in the public sector have run beyond schedule, cost more than planned, and have not met their requirements.[46] For example, the Department of Homeland Security (DHS) canceled the Secure Border Initiative Network (SBInet) program after investing about one billion dollars. GAO noted that the DHS lack of ability to effectively manage and plan long-term investments not only resulted in missed deadlines and cost overruns but also led to abandonment of the program.[47]

IT project management failures can lead to a culture of trying to play it safe, which is not conducive to leveraging emerging technologies like AI that require deep collaboration with external experts. These collaborations often call for a different perspective than traditional IT project management, because the tools are continually evolving in agile and iterative fashion. Cognitive systems learn while being deployed, from an infusion of more data and interaction with humans. As such, agencies need to develop more agile and nimble modes of shaping how they can collaborate with external parties on these kinds of projects.

Finally, even when agencies try to work collaboratively on small experiments, processes and systems can break down. Public agencies must pay especially close attention when working on interagency collaboration, to avoid financial redundancy as well as to increase data integration.

43. GCN. "DOD Plans for Better Biosurveillance." *GCN,* July 7, 2017. https://gcn.com/blogs/pulse/2017/07/biosurveillance-bsve-upgrade.aspx.
44. DRTA. "Biosurveillance Ecosystem Enhancement - HDTRA1-17-RFI-CBI-BSVE - Federal Business Opportunities: Opportunities." July 6, 2017. https://www.fbo.gov/index?s=opportunity&mode=form&id=1b59415c82c800984a7cab5126ed28aa&tab=core&_cview=0.
45. GCN. "DOD Plans for Better Biosurveillance."
46. Desouza, Kevin C. and Kendra L. Smith. "Mega-Scale IT Projects in the Public Sector." *Brookings,* May 28, 2015. https://www.brookings.edu/blog/techtank/2015/05/28/mega-scale-it-projects-in-the-public-sector/.
47. US Government Accountability Office. "Agencies Need to Strengthen Oversight of Multibillion Dollar Investments in Operations and Maintenance." no. GAO-14-66 (November 6, 2013). https://www.gao.gov/products/GAO-14-66.

Opportunity 3: Use Agile Governance and Acquisition Practices for IT

Modernizing legacy IT systems to adopt AI and other new applications involves considerable planning and managing unanticipated hitches. Larger projects have proven more likely to underperform. Public agencies must find ways to bring in new IT systems with clear planning, intermediate steps where necessary for large migration projects, and "quick wins"[48] to gain momentum and overcome resistance to change. In the case of the Maryland Total Health and Human Services Network, this meant starting with information sharing between two agencies, Human Services and Juvenile Services, rather than across multiple agencies and multiple datasets.[49]

Also consider the initiatives adopted by the FCC to retire its 207 legacy IT systems and migrate them to a cloud-based common data platform. The FCC, in an attempt to improve communications between the agency and its customers, decided to replace its aging Customer Help Center (CHC). The FCC realized that building a new CHC internally would cost $3.2 million and take about 2 years to complete. The FCC CIO brought in an outside change agent to build relationships with key FCC stakeholders and develop an alternative agile approach that employed a Zendesk cloud-based system to modernize its CHC. This new system cost about $450,000 and was completed in about six months.[50] The new system was built to make the FCC more user-friendly, interactive, accessible, efficient, fast, and transparent. The new system streamlined 18 outdated complaint forms into a single web portal.[51] Furthermore, the cost of maintaining a cloud-based CHC was about $100,000 compared to $600,000 for maintaining an in-house CHC model[52]—an 85 percent ongoing savings.

The process of moving CHC to the cloud via agile processes provided a quick win for the management team, which helped improve the agency's morale and overcome resistance to change.[53] Citizens and employees found the new system was easier to use, which further reduced concerns associated with the use of a cloud-based system and provided groundwork for modernization of other complicated legacy systems. In each case, ease of use, service improvements, rapid delivery, and cost savings delivered by cloud implementation generated quick wins that laid the groundwork for future success.

Challenge 4: Lack of Prioritization of Data-Driven Solutions

AI systems are data-intensive, best leveraged by organizations that prioritize data-driven identification of opportunities and build data-driven solutions. Public agencies have come a long way in recent years to prioritize data-driven (or evidence-driven) management and decision making. Yet most organizations have a long way to go in making data-driven identification of opportunities and solutions a priority.

48. Dawson, Gregory S. and James S. Denford. "A Playbook for CIO-Enabled Innovation in the Federal Government." IBM Center for the Business of Government, 2015. http://www.businessofgovernment.org/report/playbook-cio-enabled-innovation-federal-government.
49. Friedman, Sara. "Maryland builds cross-department cloud for data sharing." *GCN,* September 28, 2017. https://gcn.com/articles/2017/09/28/maryland-cross-department-cloud.aspx.
50. Fretwell, Luke. "Inside the new FCC consumer help center." *GovFresh,* January 14, 2015. http://govfresh.com/2015/01/inside-new-fcc-consumer-help-center/.
51. Monteith, Kris, Gigi B. Sohn, and Diane Cornell. "New Consumer Help Center Is Designed to Empower Consumers, Streamline Complaint System." Federal Communications Commission, January 25, 2015. https://www.fcc.gov/news-events/blog/2015/01/05/new-consumer-help-center-designed-empower-consumers-streamline-complaint.
52. Konkel, Frank. "Need a Use Case for Cloud? Look to FCC." Nextgov.com, July 22, 2015. http://www.nextgov.com/it-modernization/2015/07/need-use-case-cloud-look-fcc/118293/.
53. Dawson, Gregory S. and James S. Denford. "A Playbook for CIO-Enabled Innovation in the Federal Government." IBM Center for the Business of Government, 2015. http://www.businessofgovernment.org/report/playbook-cio-enabled-innovation-federal-government.

Several reasons exist for this. First, most public agencies struggle to keep pace with the current explosion of data in their midst. As such, most data go unanalyzed or unused. Second, government faces a shortage in a qualified workforce with a modern analytical skillset. Third, in most public agencies management is not often successful in infusing a data-driven mindset throughout the organization to meet mission objectives.

Opportunity 4: Identify Data Intensive Problems that Can Take Advantage of Machine Learning and Cognitive Capabilities

Public agencies need to think about challenges where they have large volumes of data in these cases, machine learning can be leveraged to automate or augment decision making in highly data-intensive tasks. Revenue and taxes are one such case. In 2012, the US federal government issued more than $5.2 billion in tax returns for fraudulent claims.[54] And as more taxes are filed online, governments are experiencing an increase in such claims.

Several states in the US have adopted data analytics systems to combat this challenge. New York, for instance, leveraged a fraud-spotting system to detect Internet Protocol (IP) addresses that looked suspicious. This new system flagged IP addresses that filed tax returns from other places such as Fort Lauderdale, Florida; tax auditors further analyzed these claims and detected fraudulent returns.[55] Similarly, Maryland utilizes analytics to detect anomalies in tax returns and reduce fraud. With the use of data analytics and algorithms to predict patterns, auditors can find tax scams and take appropriate action. The analytics-aided approach has increased tax returns in the state from about $10–20 million a year to $40 million a year.[56]

Opportunity 5: Invest in Data Governance

The public sector needs to modernize its data governance policies and practices to take advantage of data-intensive AI applications. Data needs to be liberated from systems, and data retention and sharing policies need to be standardized and modernized[57] to keep up with the needs for analysis and insights—keeping data in silos is no longer practical or popular, and prevents interoperability and connectivity of data repositories.

In addition to consolidation and modernization of data centers, building data maps—"a database inventory of what systems, applications, and repositories an agency has, where they are, and who is responsible for managing them"[58]—is a key step to take before beginning an operational phase of large changes to their data systems and implementation of AI systems.

Opportunity 6: Deploy to Improve Service Delivery

AI systems can use data to improve public and social services. The New York City Fire Department (FDNY) uses their Risk-Based Inspection System (RBIS) to predict outbreaks of fire sparks in the city. RBIS runs the FireCast algorithm that combines data from five different agencies and develops of a list of potential suspect buildings based on 60 risk factors. RBIS helps the FDNY's unit conduct coordinated inspections. Each of the FDNY units conducts 26 different types of inspections for 350,000 buildings in the city. Moreover, RBIS provides a

54. Framer, Liz. "A Victim Himself, Georgia's Revenue Commissioner Tackles Tax Fraud." *Governing,* February 2014. http://www.governing.com/topics/finance/gov-georgia-tackles-tax-fraud.html.
55. Newcombe, Tod. "States Use Big Data to Nab Tax Fraudsters." *Governing,* March 2015. http://www.governing.com/columns/tech-talk/gov-states-big-data-tax-fraud.html.
56. Caggiano, Kathleen. "Fighting the Rising Tide of Tax Fraud with Analytics: Q&A with Andrew Schaufele of Maryland's Bureau of Revenue Estimates." *Bloomberg BNA,* April 15, 2016. https://www.bna.com/fighting-rising-tide-n57982069917/.
57. Wood, Colin. "Data Governance: The Public Sector's Next Big Frontier." *gt,* April 29, 2014. http://www.govtech.com/data/Data-Governance.html.
58. Lewis, Michael. "Managing the risks of data sprawl." *FCW,* August 21, 2017. https://fcw.com/articles/2017/08/21/comment-data-sprawl-risks.aspx.

centralized data store for 49 fire units. The first version of the FireCast algorithm was launched in 2010; the current version of FireCast 2.0 is ten times more powerful than the first version. In 2015, the FDNY started developing FireCast 3.0 to combine data from 17 agencies to predict potential suspect buildings based on 7,500 factors.[59]

This experience shows that AI systems offer significant opportunities to improve the delivery of public services, and should not limit themselves to thinking about how to work within existing processes. The returns will be far greater if agencies build more sophisticated tools that take full advantage of emerging computational and analytical technologies to serve the public.

Workforce

Challenge 5: Public Workforce Management Practices

The growth in the use of cognitive systems will impact how the public sector thinks about its workforce. Automation will have an impact on how jobs are conducted; more and more low-skill and medium-skill jobs once done by human workers are now done by machines.[60] In addition, the augmentation of work and processes will call for agencies to rethink their overall human resource strategy. Transforming the public-sector workforce has not been easy for most governments, but ignoring the issue will only complicate matters down the road given how agency work will look in the future. AI and cognitive systems may result in some jobs being lost due to automation, but will also increase the quality of work being conducted and open up new work opportunities. Reequipping the current workforce to meet this challenge requires agencies to shift outdated processes for hiring into existing job positions, and even promoting personnel based on current guidelines.

Opportunity 7: Take Steps to Transform the Public Workforce to Take Advantage of AI Capabilities

Public agencies need to think about developing the capacity of employees to leverage AI. This can involve either automation (tasks that have backlogs can be automated) or augmentation (tasks that can be split up with human experts, or where AI can help with data analysis). While automation of some degree is inevitable over time, public agencies can take a proactive role in designing programs to move workers from one job to another.

59. Heaton, Brian. "New York City Fights Fire with Data." *Government Technology,* May 15, 2015. http://www.govtech.com/public-safety/New-York-City-Fights-Fire-with-Data.html.

60. Desouza, Kevin C., Gregory S. Dawson, and Bryce A. Santiago. "Disrupting work and workers in the age of cognitive computing systems." *Brookings,* April 26. https://www.brookings.edu/blog/techtank/2017/04/26/disrupting-work-and-workers-in-the-age-of-cognitive-computing-systems/.

For instance, Colorado Child Welfare County conducted a four-week study to examine the workload of child welfare caseworkers. The study found that time available from child welfare caseworkers was not sufficient to complete the mandated welfare activities. The study also identified that caseworkers spent 36 percent of their time on filling in documents and completing paperwork. They spent only 9 percent of their time with children and families. Similarly, North Carolina's Innovation Center (iCenter) found that IT departments often spend 80 to 90 percent of their time helping with password resetting related queries;[61] the iCenter is piloting a chatbot to free up IT employees time.

By automating routine, mundane tasks, companies can not only save staff time but also incur profits. Many organizations are using Robotic Process Automation (RPA) to transfer mundane tasks from employees to machines. RPA software is designed to mimic humans for carrying out functions. And, RPAs "can do repetitive stuff more quickly, accurately, and tirelessly than humans, freeing them to do other tasks requiring human strengths such as emotional intelligence, reasoning, judgment, and interaction with the customer."[62]

Augmentation is also a method to enhance the workforce. American Express and P&G are investing in AI systems and not planning to downsize their workforce while improving their productivity and creating new jobs.[63] The companies have adopted an approach to leverage cognitive computing that augments human intelligence, where machines and humans work side-by-side. Additionally, the companies took measures to encourage and educate employees about cognitive, particularly, asking employees to view these new systems as job extensions and ensuring employees that the systems will not replace them.

The US only spends about 0.1 percent of its GDP on programs related to helping workers cope with economic transitions.[64] In the early stages of AI deployment and automation, policy makers, businesses, and educational institutions need to collaborate and develop strategies that prepare the workforce for transitions. At all levels of government, a serious effort needs to be undertaken to look at the workforce and prepare for the future. AI systems are here, and they will fundamentally reshape how work gets done in the public sector. The public sector will need to modernize its human resource practices, from job classifications to training and development programs and even the overall strategy for the future of the public workforce. Agencies should take a long-term approach when contemplating options for workforce transformation, rather than simply trying to fill present-day roles and functions. Thinking through these issues now will enable agencies to make a smoother transition, rather than being disrupted by the advent of AI.

Opportunity 8: Engage Human Experts in Designing, Testing, and Evaluating AI

To develop AI and cognitive systems, public agencies must assign experts to train the systems in their decision support. Experts need to spend significant time teaching these systems, even though governments often have limited human resources to draw on. Experts will assess the accuracy of solutions provided to known problem and solution pairs. In cases where the system gets the answer wrong, experts will provide the correct answer, which will then be used by the system in future iterations, which improves overall accuracy. For example, at the US

61. Brown, Justine. "Chatbots Debut in North Carolina, Allow IT Personnel to Focus on Strategic Tasks." *Government Technology,* October 12, 2016. http://www.govtech.com/computing/Chatbots-Debut-in-North-Carolina-Allow-IT-Personnel-to-Focus-on-Strategic-Tasks.html.

62. Lhuer, Xavier. "The next Acronym You Need to Know about: RPA (Robotic Process Automation)." *McKinsey & Company,* December 2016. https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/the-next-acronym-you-need-to-know-about-rpa.

63. Davenport, Thomas H. and Randy Bean. "How P&G and American Express Are Approaching AI." *Harvard Business Review,* March 31, 2017. https://hbr.org/2017/03/how-pg-and-american-express-are-approaching-ai.

64. US Executive Office of the President. "Artificial Intelligence, Automation, and the Economy." December 2016. https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF.

Memorial Sloan Kettering Cancer Center and WellPoint,[65] clinicians and data scientists spent hours analyzing more than 2 million pages of medical data extracted from about 600,000 sources, including medical journals, articles, and clinical trials,[66] patient data such as diagnosis, medical notes, and treatment preferences. Clinicians and data scientists regularly corrected errors to improve system performance.[67]

Challenge 6: Culture of IT Ownership

Traditionally, agencies have preferred to "own" and even "build" their own systems. Owning is becoming an outmoded model of IT governance. Given the pace of technical innovations, public agencies may be better served by leveraging third-party vendor solutions to take advantage of the most recent advances in computation, including AI systems. The practice of government working in a closed loop and owning all systems made sense when government could clearly specify its needs and contractors simply had to design and build; the modern era demands a more open and collaborative process to keep up with the pace of innovation.

Governments can begin to utilize open-source approaches to generate cost savings, increase scalability, and remain up-to-date. Open software and code solutions can be maintained or updated by anyone, and thus agencies are not limited to one vendor for services or future innovation. At the federal level, the GSA's digital services group, 18F, is a team of designers, developers and product specialists that provides consultancy for government agencies to help with the rapid deployment of tools and services for more cost-efficient and reusable systems.[68]

18F has designed a new open-source experiment for agencies to better utilize their micro-purchasing authority: a simplified and informal procurement method for readily available goods and services that costs less than $3,500, run via reverse auction. After each new auction, 18F is learning more about ways to refine the process, and eventually this could be adopted by local, state, and federal government agencies. 18F's use of open-source solutions to facilitate micro-purchasing is just one example of how government can better innovate with technology and still retain quality and accountability. Creating opportunities for work to be completed quickly and cost-effectively, and in a way that is inclusive and gives qualified individuals a chance to work, could become a hallmark of government innovation for emerging technologies like AI systems.

There has been some movement in this direction with recent policies put in place by the US Government. The Federal Source Code Policy ("Achieving Efficiency, Transparency, and Innovation through Reusable and Open Source Software"),[69] provides guidance to covered agencies on software procurement considerations that must be made prior to acquiring any custom-developed software, and establishes policy requirements for government-wide source code receipt and reuse.

Opportunity 9: Develop Collaborative Partnerships with Academia to Initiate AI Projects

Academic institutions are valuable partners for high-risk exploratory projects and can also

65. IBM. "IBM Watson Breakthroughs Transform Quality Care for Patients." CTB10, February 8, 2013. https://www-03.ibm.com/press/us/en/pressrelease/40335.wss.
66. Ibid.
67. Best, Jo. "IBM Watson: The Inside Story of How the Jeopardy-winning Supercomputer Was Born, and What it Wants to Do Next." *TechRepublic,* accessed August 10, 2017. https://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next/.
68. Mergel, Ines. "Digital Services Teams: Challenges and Recommendations for Government." IBM Center for The Business of Government, 2017. http://www.businessofgovernment.org/report/digital-service-teams-challenges-and-recommendations-government.
69. Scott, Tony, and Anne E. Rung. "M-16-21 Federal Source Code Policy: Achieving Efficiency, Transparency, and Innovation through Reusable and Open Source Software." August 8, 2016. https://www.actiac.org/system/files/m_16_21.pdf.

bring critical capabilities to bear when tackling low-hanging opportunities. The Las Vegas Health Department collaboration with researchers from the University of Rochester to experiment with nEmesis, discussed under "Development" earlier in this report, is one such model to address the problem of restaurant inspection.[70] In addition, state and local governments can seek out academic institutions in their communities to work on AI development initiatives. Academic partners are often a rich source of specialized knowledge in a domain and can bring the necessary computational and analytical knowledge to bear.

Consider the case of the Charlotte-Mecklenburg Police Department (CMPD), which collaborated with researchers to develop a machine learning model to detect officers at risk of committing adverse events. The model was fed data from 2009–2015, which included information about employees (sensitive information was redacted), dispatch, arrests, traffic stops, and events recorded by CMPD to perform day-to-day actions. The machine learning model ran multiple evaluations on the data and finally aggregated the results to produce final statistics. The model predicted high-risk officers more accurately than CMPD's internal Early Intervention System that flags high-risk officers, increasing true positives by 12 percent and decreasing false positives by 32 percent.[71]

The City of Hampton, VA collaborated with researchers from the Virginia Institute of Marine Science to test its StormSense project. StormSense was created to predict the impact of floods and flag high-risk areas. StormSense utilizes street-level hydrodynamic modeling to analyze historical crowdsourced flooding data and real-time data from an ultrasonic sensor system deployed on the pilot streets.

The model can predict floods up to 36 hours in advance, which helps city government provide timely warning to citizens about floods and damages anticipated.[72] This system has proven more accurate than the flood predictions from the National Oceanic and Atmospheric Administration.[73]

Source: *StormSense Project. https://vims-wm.maps.arcgis.com/apps/MapJournal/index.html?appid=62c80853313743f3acf5a83ab420d015*

70. Sadilek, Adam. et al. "Deploying nEmesis: Preventing Foodborne Illness by Data Mining Social Media." Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI Press, 2016, 3982–89. http://dl.acm.org/citation.cfm?id=3016387.3016466.
71. Carton, Samuel. et al. "Identifying Police Officers at Risk of Adverse Events." *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,* KDD '16 (New York, NY, USA: ACM, 2016), 67–76, doi:10.1145/2939672.2939698.
72. "StormSense Project." accessed August 10, 2017. https://vims-wm.maps.arcgis.com/apps/MapJournal/index.html?appid=62c80853313743f3acf5a83ab420d015.
73. Lohmeyer, Suzette. "Faster and More Accurate Flood Prediction." *GCN,* July 21, 2016. https://gcn.com/articles/2016/07/21/storm-sense.aspx.

An important benefit of collaborating with academic partners is that one has the ability to engage the next generation of the workforce on AI systems. Students want opportunities to work on real problems where they can apply computational and system development approaches they are learning. Problems from public agencies can be used to create "living labs" that bring various stakeholders together to design use-inspired solutions and even conduct high risk exploratory projects.

Opportunity 10: Initiate Public-Private Partnerships to Design, Deploy, and Evaluate AI on Mission-Critical Priorities

Public-private partnerships are vital when it comes to AI implementation. Public agencies can conduct routine environmental scans to learn about current discoveries and advances in the AI space. Partnership arrangements of all sorts, from traditional outsourcing efforts to joint development initiatives and licensing agreements, should be explored. Partnership agreements for AI will need to be more collaborative than traditional information system sourcing projects, in which the agency specifies in detail what is required, the private sector responds with bids, and a vendor is chosen to work on the project.

It is unlikely that the agency will be able to specify AI requirements in as much detail as traditional systems. Moreover, given that an agency is looking for partners to work with after the initial system is deployed, the focus should be on building long-term collaborations—build version one and then keep refining the system to increase its capability and sophistication.

The launch of cross-sector AI initiatives is also a good opportunity to spur innovation. Cognitive computing can play a vital role in helping humans to solve significant challenges through assessing data housed across many agencies and systems. This bring agencies and industry together to link and fuse relevant data, thereby enhancing advanced learning systems. In addition, similar to the success seen with crowdsourcing efforts, competitions and challenges that invite stakeholders to build and evaluate cognitive systems can be organized.

Challenge 7: Limited Capacity for System-Level Redesign

AI systems bring unique challenges in that they require organizations to think deeply about how work and processes are executed to meet mission goals. The real opportunity is not to do things more efficiently than they are today by automating existing processes; rather, cognitive approaches allow for an opportunity to re-think how the underlying process is executed and find a more effective manner to complete the task.

Public agencies often lack an appetite to conduct fundamental and broad analysis of their work systems and process architectures. This has resulted in often incremental and patchwork technology solutions that limit future flexibility for system-level reengineering. In addition, the manner in which programs and activities are funded in the public sector limits the ability of an agency to spend resources on large-scale fundamental reform. It is easier to secure funding for existing programs than to propose new efforts, especially when such efforts are long-term and do not have immediate mission outcomes.

Opportunity 11: Redesign Work Processes to Increase Effectiveness and Efficiency

AI systems can have a large impact on the efficiency of public regulatory compliance. For example, over the past few years—particularly after hurricane Katrina—the City of New Orleans has struggled to address the problem of urban blight. The City's Code Enforcement Department faced backlogs of cases despite increasing the number of hearings from 270 per

month in 2012 to 416 per month in 2014.[74] In 2015, the City collaborated with Enigma.io to develop an algorithm to address this challenge. Several interviews were conducted with blight case reviewers to determine decision criteria. The algorithm dubbed as the Code Enforcement Abatement Tool was designed to sort building cases and recommend whether a structure needed to be demolished or could be sold. The algorithm considers dozens of variables and produces an abatement score between 0 (strongly recommend demolition) to 100 (strongly recommend resale). Then, an officer reviews the recommendation and makes a decision. The tool resolved more than 1,080 cases within the first 90 days.[75]

Cognitive systems can also augment human decision-making to develop more nuanced medical treatments, and to detect patterns in large volume of healthcare data. For example, researchers from Harvard University and Massachusetts Institute of Technology applied deep learning techniques to predict cancer.[76] The researchers trained a deep neural algorithm to distinguish tumor patches from normal patches. The algorithm aggregated patch-level analysis to create tumor probability heat maps, and utilized a 27-layer neural architecture to achieve nearly human-pathologist-level precision. The researchers also found that the algorithm errors were not correlated with human pathologist errors, and combined the algorithm's predictions with pathologist diagnosis to help reduce human error from 3 percent to less than 1 percent. The leveraging of terabytes of data presents an untapped potential that can transform healthcare.

Opportunity 12: Employ AI to Augment Human Decision Making

AI systems can also play a role in increasing objectivity for decision making. For example, courts around the world are often criticized as being partial. Theoretically, judges make decisions based on numerous factors. A judge deciding parole cases may consider factors such as the offender's criminal history, whether an offender is the main suspect or accessory to crime, type of offense, and whether the person committed a crime in duress. However, one study reported that a judge is more likely to grant parole after a meal break. The researchers followed eight judges serving on parole boards for four major prisons in Israel for ten months. They collected 1,112 judicial rulings. Data analysis revealed that judges ruled favorably in the early morning sessions (about 65 percent granted parole) and the number fell to zero with subsequent decisions. However, the judges start ruling favorably again after meal breaks, i.e., about 65 percent of cases are granted parole, and the number then fell again.

To the extent that biases may play a role in deciding court sentencing, AI systems can be deployed jointly with humans so as to analyze the results of decisions and use this data to identify any anomalies. Cognitive approaches can also be used to test out new approaches to problem solving, not just by mimicking human decision-making approaches, but also by finding novel ways to solve problems.

74. NOLA. "Code Enforcement Abatement Tool." accessed August 10, 2017. https://www.nola.gov/performance-and-accountability/reports/nolalytics-reports/nolalytics-blight-abatement-tool-brief/.
75. Wood, Colin. "Grounding AI: Artificial Intelligence Is Closer—and Less Awesome—than Most Realize." *Government Technology,* January 20, 2016. http://www.govtech.com/Grounding-AI-Artificial-Intelligence-is-Closer-and-Less-Awesome-than-Most-Realize.html.
76. Wang, Dayong, et al. "Deep Learning for Identifying Metastatic Breast Cancer." *arXiv:1606.05718 [Cs, Q-Bio],* June 18, 2016. https://arxiv.org/abs/1606.05718.

Risk Management

Challenge 8: Securing Systems

Recent breaches of public and private IT systems are well-documented. Security issues can be magnified by the development and deployment of AI applications, which often create highly connected systems-of-systems. For instance, as cities become highly connected through the deployment of AI, failure in one city system may ripple across others. One study that examined traffic lights and sensors found that these systems were highly susceptible to cyber-attacks and that traffic lights and sensors could easily be modified.[77]

In another example, the recent Equifax breach of consumer data highlights the importance of public (and private) entities securing large datasets, especially those storing sensitive information such as Social Security numbers or voting information. Public agencies need to find ways to bolster their cybersecurity capabilities before undertaking significant AI projects, as the risks from data compromises, algorithmic manipulations, and even systems-to-systems cascading failures will remain high. However, agencies can start with leveraging significant open data that is available, and working with less sensitive data (e.g., citizen queries submitted via websites) to build cognitive learning systems.

Opportunity 13: Develop Cybersecurity Capacity

Agencies must deal with complex issues of cybersecurity, and these issues will become more critical in the age of AI. Cybersecurity issues play out in numerous and complicated ways, especially as algorithms themselves are learning as they interact with the data and may eventually become hard to dissect and decode (e.g., one cannot specify the inputs and the weights like they do in a regression model). These systems might be compromised in ways that are hard to detect, and the manipulations could result in severe damage.

In addition, as public agencies are increasingly implementing chatbots to interact with customers, they need to think about security and privacy concerns. Chatbots could initiate contact with organizations as clients and access sensitive information. Similarly, citizens can become potential targets for bots designed by hackers and scammers. The cost of ransomware alone is high, with the FBI's Internet Crime Complaint Center receiving 7,694 ransomware complaints totaling $57.6 million since 2005.[78] Public agency executives should bridge security gaps for legitimate chatbots and similar applications in their people, data, applications, and infrastructure.[79]

77. Fadilpašić, Sead. "Smart Traffic Sensors Are Vulnerable to Cyber-Attacks." *IT Pro Portal,* April 19, 2016. https://www.itproportal.com/2016/04/19/smart-traffic-sensors-are-vulnerable-to-cyber-attacks/.
78. Lyngaas, Sean. "DHS: Over 300 incidents of ransomware on federal networks since June." *FCW,* March 30, 2016. https://fcw.com/articles/2016/03/30/ransomware-carper-hsgac.aspx.
79. Lainhart, John W., and Christopher Ballister. "Managing Advanced Threats in the Digital Age: Addressing security, risk and compliance for U.S. Public Sector executives." IBM Center for the Business of Government. http://www.businessofgovernment.org/sites/default/files/Managing Advanced Threats in the Digital Age.pdf.

There needs to be clear attention to the security and privacy details on all elements of AI systems, from data to the analytical engines and algorithms, computing platforms, networked systems, and devices. A comprehensive plan to monitor the system and ensure that it operates with integrity should accompany any AI initiative—one that includes tools, personnel, and resources.

Challenge 9: Risk Aversion

Public agencies are risk averse, which limits their ability to experiment and ingest emerging technologies. Even in cases where a technology has moved out of research and development, significant barriers exist when adopting innovations. Given that government is in the early stages of AI development, and even earlier stages when it comes to deploying full-fledged cognitive systems, agencies are often content to follow rather than lead.

This wait-and-see-approach is not helpful, given the rapid pace of innovation in the AI space. Innovation barriers that hold public agencies behind include:

- the prioritization of a closed approach to innovation—the default practice is generally to look inward for solutions and limit external engagement.
- the risk rationalization process—standard processes emphasize solutions that are mature and where technical and financial feasibility is calculated based on experiences from prior deployments.
- the procurement process—stagnant, outdated, and cumbersome buying rules make acquisition of innovative solutions difficult in a timely manner.

Successfully designing, developing, and deploying AI systems will require public agencies to rethink the notion of risk—specifically an understanding that AI projects do have risk but these risks can be managed over time, and that the cost of not getting engaged may severely limit future opportunities.

Opportunity 14: Promote Innovation Through Crowdsourcing Platforms

New agency platforms can promote experimentation and innovation with AI. Government should consider using competitions to engage external parties on developing solutions around salient challenges and issues.[80] Consider the following case: In September of 2015, a 19-year-old UK programmer launched a bot in London which allowed residents to contest their parking tickets.[81] By February, more than $3 million traffic ticket appeals were initiated through the bot. So far, the bot has successfully appealed 64 percent (160,000 of 250,000) of parking ticket cases both in London and New York. Agencies need to find ways to bring in outside talent and expertise like this when it comes to accelerating machine-learning applications.

Public agencies can leverage crowdsourcing platforms for key activities in moving forward with AI. For example, in June 2016, the US Office of Science and Technology Policy (OSTP) launched and RFI on Artificial Intelligence to solicit feedback from the public[82] and received about 161 comments from diverse stakeholders[83] (e.g., citizens, academic researchers, and

80. Desouza, Kevin C. "Challenge.gov: Using Competitions and Awards to Spur Innovation." IBM Center for the Business of Government, 2012. http://www.businessofgovernment.org/report/challengegov-using-competitions-and-awards-spur-innovation.
81. Muoio, Danielle. "Joshua Browder Bot Overturns 160,000 Parking Tickets." *Business Insider,* June 28, 2016. http://www.businessinsider.com/joshua-browder-bot-overturns-160000-parking-tickets-2016-6.
82. US OSTP. "Request for Information on Artificial Intelligence." *Federal Register,* June 27, 2016. https://www.federalregister.gov/documents/2016/06/27/2016-15082/request-for-information-on-artificial-intelligence.
83. Smith, Megan. "Public Input and Next Steps on the Future of Artificial Intelligence." *Medium,* September 6, 2013. https://medium.com/@USCTO/public-input-and-next-steps-on-the-future-of-artificial-intelligence-458b82059fc3.

non-profit organizations). OSTP encouraged people to respond and recommend issues on wide-ranging topics related to AI such as governance, workforce training, tools for social good, and economic impacts. This process helped the OSTP to understand and further examine key issues governing the use of AI in developing recommendations for policy.

Opportunity 15: Increase Awareness of the Potential for AI Through Exploratory Projects

By leveraging AI, public agencies can reduce risks and arrive at data-driven intelligence. This can help speed up the process of adoption of AI systems by increasing awareness of the benefits, and beginning projects at a small scale. Governments need to think about new ways to integrate AI by experimenting with sub-projects, which can become part of larger projects. Public agencies can also use phased approaches to integrating AI into the organizations. Once agencies feel comfortable, they can leverage new systems to address agency-wide problems, and even gain unexpected insights from the data analysis.

The Australian government, for instance, is investing about AUS$195 billion (approximately US$154.34 billion) to improve the capability of its defense in the next ten years.[84] Australia's Department of Defense will implement several initiatives to leverage AI for improving its surveillance, electronic warfare, cybersecurity capacity, and other programs.[85] Recently, Australia's Department of Defense piloted cognitive systems for conducting Psychological Operations (PSYOPs) analysis. As a proof-of-concept, the department chose Target Audience Analysis (TAA)—a detailed and systematic examination that helps select the audience for PSYOPs missions. The team completed TAA in two weeks compared to traditional analysis that takes three to six months to complete. The department also discovered several interesting insights from the analysis of unstructured data, and plans to use cognitive approaches to develop high-impact initiatives (e.g., warfare).[86]

In April 2017, the US Department of Defense (DoD) launched Algorithmic Warfare Cross-Functional Team (AWCFT) dubbed as Project Maven.[87] The primary goal of AWCFT is to leverage machine learning and detect patterns from DoD's massive amount of data. The AWCFT will integrate machine learning tools to augment (or automate) field technologies for Processing, Exploitation, and Dissemination (PED) for tactical Unmanned Aerial System (UAS) and MidAltitude Full-Motion Video (FMV) for the Defeat ISIS campaign. The efforts are likely to decrease the burden on employees and provide data-driven intelligence for improving DoD's decision making. The AWCFT will integrate computer vision algorithms to classify, manage, and sort data in 90-day sprints. Once the AWCFT trains the algorithms on computer vision, it will then move on to develop additional functionalities and features in the next 90-day sprint. The AWCFT will also consolidate algorithms-based initiatives at DoD.

Opportunity 16: Develop Collaborative Partnerships Across Agencies on Cross-Cutting Issues

Public agencies should collaborate with their peers on issues that neither agency can tackle on their own. Agencies can issue challenges on effective ways to address these issues, potentially as multi-stage competitions—where teams must meet certain deliverables during each stage. To solve the problem of overcrowded radio frequency (RF) signals, for example, DARPA

84. Australia Department of Defence. "2016 Defence White Paper." 2016. http://www.defence.gov.au/whitepaper/docs/2016-defence-white-paper.pdf.
85. Australia Department of Defence. "Integrated Investment Program." 2016. http://www.defence.gov.au/WhitePaper/Docs/2016-Defence-Integrated-Investment-Program.pdf.
86. Aiyaswami, Mohan. "Smarter Australian Defence Force the Information Advantage – CEBIT." Technology, May 2016. https://www.slideshare.net/CeBITAustralia/cebitaustralia2016mohanaiyaswamiegovernmenttransformationdefence.
87. DoD. "Establishment of an Algorithmic Warfare Cross-Functional Team (Project Maven)." April 26, 2017. http://www.govexec.com/media/gbc/docs/pdfs_edit/establishment_of_the_awcft_project_maven.pdf.

announced a Spectrum Collaboration Challenge[88] (SC2) In 2016. The SC2 aims to increase RF with machine learning technologies to support military and civilian wireless devices. However, unlike traditional challenges, the SC2 encourages teams to develop systems that collaboratively adapt in real time to a spectral environment. The three-phase competition will run for three years from 2017–2020. DARPA has selected 30 contenders that include eight individuals and 22 academic, small, and big businesses.[89]

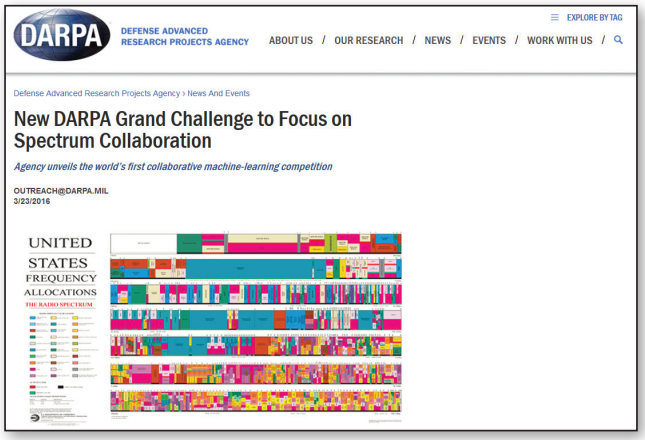

Source: *DARPA Grand Challenge https://www.darpa.mil/news-events/2016-03-23*

Similarly, GSA recently organized a Hackathon event to engage hackers and public agencies for developing AI-based solutions using government data.[90] GSA invited more than 100 hackers, 27 representatives from federal agencies, and businesses. During the event, many federal agencies shared information about their data and the respective formats. GSA aimed to foster discussion among diverse stakeholders to apply AI tools for making government data more accessible to spur innovation.[91] The hackathon was part of GSA's Emerging Citizen Technology program[92]—an open-source pilot that aims to help federal agencies leverage public data to make government services more citizen-friendly by leveraging AI and machine learning. The program goal is to create solutions that consider privacy, security, and accessibility challenges associated with the implementation of AI tools.

Challenge 10: Ethical and Social Considerations

Transparency and ethical use of data also become critical issues with the move towards AI, increasing the need to develop new ethical frameworks around algorithms that support decision making. Public agencies need to educate people about these systems, be transparent in their design, and make clear how to deduce and report errors. Additionally, cognitive computing raises the need to make new ethical judgments.

88. DARPA. "New DARPA Grand Challenge to Focus on Spectrum Collaboration." March 23, 2016. https://www.darpa.mil/news-events/2016-03-23.
89. DARPA. "The Spectrum Collaboration Challenge: Let the Games Begin!" January 30, 2017. https://www.darpa.mil/news-events/2017-01-30.
90. Douglas, Theo. "GSA Hackathon Sought Tangible AI Solutions for Making Government Data Accessible via 'Personal Assistant.'" *Government Technology,* May 23, 2017. http://www.govtech.com/civic/GSA-Hackathon-Sought-Tangible-AI-Solutions-for-Making-Government-Data-Accessible-Via-Personal-Assistant.html.
91. Herman, Justin. "Opening Public Services to Artificial Intelligence Assistants." *The GSA Blog,* June 6, 2017. https://gsablogs.gsa.gov/gsablog/2017/06/06/opening-public-services-to-artificial-intelligence-assistants/.
92. Herman, Justin. "AI-Assistant-Pilot: Inter-Agency Federal AI Personal Assistant Pilot." U.S. General Services Administration, 2017. https://github.com/GSA/AI-Assistant-Pilot.

Socioeconomic issues are also associated with AI systems. The development of predictive algorithms can also result in biases based on gender or race,[93] not because they are designed with biases in mind but because they may learn from data that is skewed or not representative of the population. It will be important to have frameworks in place to review and check cognitive systems, to safeguard against these issues.

Opportunity 17: Prioritize Value-Sensitive Design

Agencies can lead the conversation about how to build AI systems in a manner that preserves and advances human values. Public agencies should set up forums for dialogue on these issues at all levels—local, regional, and national. International bodies (as highlighted early in the report) have already begun such conversations, which should lead to the creation of conventions and conferences to develop global norms around these topics. This is also an area where new research can inform government.

Soul Machines, for instance, is developing customer service chatbots that can understand customers' facial expressions and respond with empathy.[94] These advances in chatbots will make them more human-like and increase their sophistication to improve human-to-machine interactions. Local Motors is experimenting on developing driverless electric shuttle bus—Olli—that can cater to people with disabilities. Powered by AI technologies, this shuttle can deploy machine vision to help a deaf passenger or use voice cues to help a visually impaired passenger find empty seats.[95]

Researchers are also working to build cognitive systems that comply with international humanitarian laws and learn human values. Quixote, for example, teaches human values such as right and wrong behaviors to robots.

The program teaches robots to read stories to learn about acceptable practices in human society.[96] Two professors—Susan and Michael Anderson—are working on building robots that care for seniors. Their model relies on teaching robots ethical behavior by training them with an ethicist.[97]

Taking a value-sensitive design perspective calls for ensuring that the voices of various stakeholder groups are adequately represented in the technologies being deployed. Doing so requires the public sector to be open and engaging with external stakeholders. The value of citizen engagement when it comes to the design of policies and initiatives is clear. Citizen engagement and, more broadly stakeholder engagement, will be vital for government's advance in building AI systems. Where possible, these systems should be designed *with* stakeholders rather than for stakeholders.[98] This approach will allow for the greatest chance that stakeholders have a voice and can contribute positively.

93. Desouza, Kevin C. and Rashmi Krishnamurthy. "How can cognitive computing improve public services?" *Brookings,* October 13, 2016. https://www.brookings.edu/blog/techtank/2016/10/13/how-can-cognitive-computing-improve-public-services/.
94. Simonite, Tom. "These Chatbots Have Realistic Faces and Can Read Your Expressions." *MIT Technology Review,* March 22, 2017. https://www.technologyreview.com/s/603895/customer-service-chatbots-are-about-to-become-frighteningly-realistic/.
95. Woyke, Elizabeth. "This Shuttle Bus Will Serve People with Vision, Hearing, and Physical Impairments—and Drive Itself." *MIT Technology Review,* April 13, 2017. https://www.technologyreview.com/s/604116/a-self-driving-bus-that-can-speak-sign-language/.
96. Georgia Institute of Technology. "Using Stories to Teach Human Values to Artificial Agents." *ScienceDaily,* February 12, 2016. https://www.sciencedaily.com/releases/2016/02/160212200239.htm.
97. Goldhill, Olivia, and Olivia Goldhill. "Can We Trust Robots to Make Moral Decisions?" *Quartz,* accessed August 10, 2017. https://qz.com/653575/can-we-trust-robots-to-make-moral-decisions/.
98. Desouza, K.C., Y. Awazu, S. Jha, C. Dombrowski, S. Papagari, P. Baloh, and J.Y. Kim. "Customer-Driven Innovation." Research-Technology Management, 51 (3), 2008, 35-44.

Opportunity 18: Focus on Protecting Public Values and Common Good

AI systems in the public sphere should, and will, be scrutinized on a regular basis. Unfortunately, too often in the press we hear of stories where AI has been misused or produced negative outcomes unintentionally. For example, as lenders use AI to determine which applications will receive loans, citizens can be victims even though they never had a choice in terms of opting in. Companies use algorithms to sort through a large volume of applications and decide on candidates to call for interviews. Insurance companies are using AI-based systems to determine differential insurance premiums. While these practices promote efficiency, they can also reduce consumer choice if not implemented with an eye toward the public good.

Developing rules and mechanisms to monitor algorithms may not be easy. Companies developing algorithms often guard the disclosure of their systems as trade secrets. Some researchers argue that current laws are insufficient to protect people against flawed algorithms. They suggest the establishment of a third-party arbitrator, who can investigate the decisions of algorithms[99] as AI algorithms evolve and change with more data.[100]

Given that public agencies are held to a higher standard when it comes to technical systems, agencies must ensure that they maintain values of transparency and fairness for algorithms. Stakeholders should have the right to receive an explanation so as to ensure fair process. It is also vital that governments pay particular attention to how these systems might introduce bias in an effort to increase efficiency. As AI systems are deployed, it is highly recommended that agencies find ways to "open up" the systems for inspection and audit, to ensure that the public understands how the system operates and can trust its outcomes (see discussion below).

Challenge 11: Governance

Emerging technologies are often introduced without standards in place. Many entities, including state and local governments, create regulations in rapid and reactive mode, as is currently the case with AI-powered drones.[101] The lack of consistent policy leads to two scenarios: Either organizations wait for guidance and do not move forward with new technology, or systems are deployed without a policy context for social and policy implications. Policy can help address data issues associated with technologies such as learning algorithms. The UK insurance company Admiral, for example, was blocked in October 2016 from using Facebook posts to categorize the risk level of potential drivers based on what their posts indicated about their personality.[102]

A key emerging area demonstrating the need for policy is in medical data mining. Started by one of the founders of the Beijing Genomics Institute, iCarbonX is "aimed at building an ecosystem of digital life based on a combination of individual's biological, behavioral and psychological data, the Internet and artificial intelligence."[103] iCarbonX would provide comprehensive health analysis and prediction through advanced data mining, machine analysis, and partnerships with medical entities across the spectrum—including research institutions, pharmaceutical factories, medical examination centers, and insurance companies. Policy can protect how a

99. Wachter, Sandra, Brent Mittelstadt, and Luciano Floridi. "Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation." SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, December 28, 2016). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2903469.

100. Lazer, David. "The Rise of the Social Algorithm." *Science* 348, no. 6239 (June 5, 2015): 1090–91, doi:10.1126/science.aab1422.

101. Desouza, Kevin C., Sabrina P.K. Glimcher, and David Swindell. "Drones and the 'Wild West' of regulatory experimentation." Brookings, August 17, 2015. https://www.brookings.edu/blog/techtank/2015/08/17/drones-and-the-wild-west-of-regulatory-experimentation/.

102. Ruddick, Graham. "Facebook forces Admiral to pull plan to price car insurance based on posts." *The Guardian,* November 2, 2016. https://www.theguardian.com/money/2016/nov/02/facebook-admiral-car-insurance-privacy-data.

103. iCarbonX, https://en.icarbonx.com/.

patient's data is collected, what it is used for, and what are the best systems to manage these new technologies.

Opportunity 19: Proactively Monitor Systems to Track Unexpected Outcomes

Cognitive systems differ from previous kinds of information systems in that they continuously learn and improve while on the job. This requires original thinking about how these systems are monitored and maintained, beyond initiatives such as ensuring things like up time or number of errors.

More broadly, AI-based tools designed to mine large volumes of data, but lacking thoughtful consideration about the environment in which they are deployed, can lead to failure. AI developers need to think critically about the social environment surrounding the application. Even AI-based tools with similar features and functionalities may succeed in one environment and fail in another.

Public agencies must invest resources to ensure that AI systems are monitored after deployment. Not doing so will result in unintended outcomes going unchecked until it may be too late; the performance of these systems needs to be evaluated on a regular basis against the intended set of outcomes. Agencies also need to plan for system downtime, should the system have to be decommissioned for a short duration or permanently. As such, it is important not to act with haste given initial positive results. Cognitive systems learn over time—governments need to provide that time, adapt to various inputs, and allow for a wide assortment of new scenarios while maintaining data integrity.

Opportunity 20: Develop Robust Audit and Inspection Mechanisms

Public agencies, and especially agencies that have specific audit and inspection mandates, need to develop capabilities to audit and inspect AI systems. While most of the skills that information system auditors have today are relevant, they are not sufficient. Given that these systems do not have agreed-on models and step-by-step functions that can be audited, innovations in auditing are required. Early stage work in how to open even complex cognitive systems is currently underway.[104] Government could fund further research and exploratory projects to advance the art and science of how to audit and inspect in this emerging area.

Several efforts in the legal space provide models that could inform this challenge. For example, researchers[105] from Cornell University, Stanford University, University of Chicago, and Harvard University trained an algorithm to predict judges' decisions to detain or release defendants. The algorithm was trained on a dataset containing 758,027 defendants arrested between 2008 and 2013 in New York City to forecast crime risk. The dataset included information such as current case, arrest history, age, and type of offense suspected. The study estimated that the use of the algorithm would help reduce crime rates with no change in jailing rates. Or, the algorithm would reduce the prison population with no increase in crime rates. The researchers replicated their analysis on a dataset covering 151,461 defendants between 1990 and 2009 in 40 large urban counties in the US. The algorithm consistently scored highly in determining whether a defendant is either a flight risk or likely to commit other crimes. For the new dataset, the researchers found that, on average, the algorithm would reduce the crime rate by 18.8 percent with no increase in release rates, or the algorithm would reduce the jail population by 24.5 percent.

104. Voosen, Paul. "How AI detectives are cracking open the black box of deep learning." *Science,* July 6, 2017. http://www.sciencemag.org/news/2017/07/how-ai-detectives-are-cracking-open-black-box-deep-learning.

105. Kleinberg, Jon, et al. "Human Decisions and Machine Predictions." National Bureau of Economic Research, 2017. http://www.nber.org/papers/w23180.

As use of algorithms such as these become more prevalent, lack of knowledge about their internal functioning will become a critical concern. Public agencies should begin a dialogue around how to ensure that systems being used can be audited. Toward this end, it will be important for public agencies to set up multi-sector partnerships to understand the various trade-offs when opening up AI applications, and begin to create policies to govern transparency for these emerging technologies.

CONCLUSION

Implementing AI in government involves the design, building, use, and evaluation of cognitive computing and machine learning to improve the management of public agencies, the decisions leaders make in designing and implementing public policies, and associated governance mechanisms.

AI systems offer significant value to the public sector, which is in early days of experimentation in this emerging field. While AI systems are being deployed first in the agencies that have traditionally been heavily data-intensive (e.g., agencies focused on energy, national security, defense, etc.), significant opportunities lie ahead for agencies that have not fully leveraged their data assets, including state and city social agencies. All agencies need to address several fundamental IT issues, from modernizing their IT infrastructure to upgrading their cybersecurity capabilities, before they can fully leverage cognitive systems.

While the optimal adoption and implementation of AI poses a challenge for public agencies, this report outlines a set of key recommendations to address these challenges. We also provide a maturity framework as an appendix to enable public agencies to identify their current level of knowledge and practice, and to progressively advance with AI adoption. Given the current early stage of knowledge and adoption of AI systems by public agencies, there is a large opportunity for government to improve their own functioning and serve citizens through adoption and implementation.

APPENDIX: MATURITY FRAMEWORK FOR AI

Maturity models are valuable tools to help organizations refine their processes, capabilities, and systems within domains.

Many fields such as education, quality control, software engineering, knowledge management, and security use maturity models. While maturity models might vary depending on the discipline and their intent, they frequently outline a series of evolutionary levels that an entity (e.g., individual, team, or organization) must progress through to increase their sophistication with processes and return on investments.

In the sections above, the recommendations outlined will enable agencies to move from one level to the next. In implementing AI systems, it is important for agencies to scan their environment to assess their current level. For each level, opportunities are noted that should be a priority.

Level 1: Unaware

Agencies at this level have not yet begun to explore the potential of AI. Much of this has to do with their lack of awareness about AI systems and/or the fear factor associated with introducing emerging technologies. Most agencies at this level spend a considerable amount of their IT budget trying to keep their current systems relevant, yet face significant challenges due to the rate of obsolescence. A fair amount of structured tasks are still conducted by their workforce, often resulting in significant backlogs and long processing times. The potential for technology-enabled innovation is stifled due to the lack of strategic foresight and direction.

Priority Opportunities

- Invest in Upgrading and Modernizing the IT Infrastructure
- Invest in Data Governance and Data Liberation Efforts
- Develop Cybersecurity Capacity

Level 2: Aware and Exploring

Agencies at this level are aware of the potential of AI. CIOs, and other senior leaders both within the IT function and beyond, have begun conversations on the role that cognitive systems can play within the agency. An initial exploration of key strategic areas where AI can be most beneficial is conducted. Exploratory initiatives are underway to get first-hand experiences with AI systems. The agency begins to look for external partners to collaborate with on initial cognitive pilots. The agency begins to take active steps to modernize the IT infrastructure and data governance mechanisms to ready themselves for AI systems.

Priority Opportunities

- Use Agile Governance and Acquisition Practices for IT
- Increase Awareness of the Potential with AI Though Exploratory Projects
- Develop Collaborative Partnerships with Academia to Initiate AI Projects

Level 3: Organized and Initial System Deployments

Agencies at this level have made a commitment to leveraging AI to optimize service delivery, foster innovation, and modernizing work practices. IT modernization efforts will be in full swing to liberate and organize data while at the same time transitioning to contemporary solutions (e.g., cloud computing). Senior IT leadership will have secured necessary commitments to develop an agency-wide strategic plan to guide investment efforts. Lessons learned from exploratory projects will be reviewed and used to shape initial system deployments. The agency will turn to their partners in industry and academia on system design, deployment, and evaluation activities.

Additionally, efforts will be underway to reshape how work is conducted in the agency. Staff will be called upon to share their expertise, train the systems, and monitor system performance when launched. Priority will be given to developing systems that can provide automated responses to standard queries and where cognitive systems can help triage incoming work so as to ensure that the most complex requests are handled by humans and the rest are processed through the system autonomously.

Priority Opportunities

- Look at Cloud Computing and Open Source as Options to Get Started
- Initiate Public-Private Partnerships to Design, Deploy, and Evaluate AI for Mission-Critical Priorities
- Promote Innovation Through Crowdsourcing Platforms
- Develop Collaborative Partnerships Across Agencies on Cross-Cutting Issues
- Identify Data Intensive Problems that Can Take Advantage of AI Capabilities

Level 4: Deployments at Scale

At this level, agencies will be active in deploying AI systems to scale, energized by the results from the initial system deployment efforts. Senior IT leaders will ensure that their efforts are given strategic-level attention. Progress on AI projects will be tracked and evaluated on a regular basis. An agency-wide transformation effort will be underway to embed AI into work practices. Efforts will be underway across the agency to redefine how work is conducted, administrative processes are designed, and resources are invested so as to take advantage the opportunities for innovation and efficiencies.

Robust collaborative networks with industry and academia will allow the agency to keep current on AI innovations and outsource system development efforts strategically. The agency will have a portfolio of systems deployed in an autonomous manner with experts playing a key role in auditing and updating them. Cognitive computing will also be embedded into work practices to provide the workforce with insights when dealing with complex problems. Systematic practices will be in place for the workforce to keep systems updated. Senior leadership will have access to metrics to track the performance and outcomes.

Priority Opportunities

- Deploy to Improve Service Delivery
- Redesign Work Processes to Increase Effectiveness and Efficiency
- Develop Robust Audit and Inspection Mechanisms
- Prioritize Value-Sensitive Design

Level 5: Refinement and Disciplined Innovation

Agencies at this level have a robust platform in place to deploy AI and leveraging cognitive computing. The agency considers cognitive systems as a strategic asset and they are embedded in almost all facets of the organization to drive value and achieve mission objectives. Agency leadership appreciate the value provided by AI and look for potential solutions when tackling complex problems. Metrics are used to continuously track the performance of systems and identify areas for refinement. The agency proactively engages in promoting innovation and works collaboratively with external partners in this regard. The agency is seen as leader in the AI space and lends its expertise to peer agencies. Efforts will be underway to tackle complex thematic issues that require inter-agency collaboration.

Priority Opportunities

- Take Steps to Transform the Public Workforce to Take Advantage of AI Capabilities
- Engage Human Experts in Designing, Testing, and Evaluating AI
- Proactively Monitor Systems to Track Unexpected Outcomes
- Focus on Protecting Public Values and Common Good

ACKNOWLEDGEMENTS

Rashmi Krishnamurthy served as lead research assistant on this project. This project would not have been completed on schedule without her dedicated work. I would also like to thank my research collaborators who provided valuable input, especially Gregory S. Dawson, Bryce Santiago, James Denford, and Lena Waizenegger. Rahul D'Silva edited several versions of this report.

ABOUT THE AUTHOR

Kevin C. Desouza is a Foundation Professor in the School of Public Affairs at Arizona State University.

KEVIN C. DESOUZA

Kevin C. Desouza is a Foundation Professor in the School of Public Affairs at Arizona State University. He served from 2012 to 2016 as the inaugural associate dean for research for ASU's College of Public Service and Community Solutions. He is a Nonresident Senior Fellow in the Governance Studies Program at the Brookings Institution. Immediately prior to joining ASU, he directed the Metropolitan Institute in the College of Architecture and Urban Studies and served as an associate professor at the School of Public and International Affairs at Virginia Tech. From 2005–2011, he was on the faculty of the University of Washington (UW) Information School and held adjunct appointments in the UW's College of Engineering and at the Daniel J. Evans School of Public Policy and Governance. At UW, he co-founded and directed the Institute for Innovation in Information Management (I3M); Founded the Institute for National Security Education and Research, an inter-disciplinary, university-wide initiative in August 2006, and served as its director until February 2008; And was an affiliate faculty member of the Center for American Politics and Public Policy.

He holds a visiting professorship at the University of Ljubljana. He has held visiting positions at the Center for International Studies at the London School of Economics and Political Science, the University of the Witwatersrand in South Africa, and the Accenture Institute for High Business Performance in Cambridge, Massachusetts. Desouza has authored, co-authored, and/or edited nine books, the most recent being *Intrapreneurship: Managing Ideas within Your Organization* (University of Toronto Press, 2011). He has published more than 150 articles in journals across a range of disciplines including information systems, information science, public administration, political science, technology management, and urban affairs.

His work has also been featured by a number of outlets such as *Sloan Management Review, Stanford Social Innovation Research, Harvard Business Review, Forbes, Businessweek, Wired, Governing, Slate.com, Wall Street Journal, USA Today, NPR, PBS,* and *Computerworld,* among others. Most recently, he curated the Ideas to Retire series for the Brookings Institution that examines practices that are stifling technological innovation in the public sector. Desouza has advised, briefed, and/or consulted for major international corporations, non-governmental organizations, and public agencies on strategic management issues ranging from management of information systems to knowledge management, innovation programs, crisis management, and leadership development. Desouza has received over $1.8 million in research funding from both private and government organizations. For more information, please visit: http://www.kevindesouza.net.

KEY CONTACT INFORMATION

To contact the author:

Kevin C. Desouza

Foundation Professor
College of Public Service & Community Solutions

Arizona State University

411 North Central Avenue

M/C 3520, Suite 750

Phoenix, AZ 85004

Phone: (206) 859-0091

kev.desouza@gmail.com

REPORTS FROM THE IBM CENTER FOR THE BUSINESS OF GOVERNMENT

For a full listing of our publications, visit www.businessofgovernment.org

Recent reports available on the website include:

Acquisition

Ten Actions to Improve Inventory Management in Government: Lessons From VA Hospitals by Gilbert N. Nyaga, Gary J. Young, and George (Russ) Moran

Beyond Business as Usual: Improving Defense Acquisition through Better Buying Power by Zachary S. Huitink and David M. Van Slyke

Collaborating Across Boundaries

Cross-Agency Collaboration: A Case Study of Cross-Agency Priority Goals by John M. Kamensky

Interagency Performance Targets: A Case Study of New Zealand's Results Programme by Dr. Rodney Scott and Ross Boyd

Improving Performance

Seven Drivers Transforming Government by Dan Chenok, Haynes A. Cooney, John M. Kamensky, Michael J. Keegan, and Darcie Piechowski

Five Actions to Improve Military Hospital Performance by John Whitley

Maximizing the Value of Quadrennial Strategic Planning by Jordan Tama

Leadership, Change, and Public-Private Partnerships: A Case Study of NASA and the Transition from Space Shuttle to Commercial Space Flight by W. Henry Lambright

Innovation

Tiered Evidence Grants - An Assessment of the Education Innovation and Research Program by Patrick Lester

A Playbook for CIO-Enabled Innovation in the Federal Government by Gregory S. Dawson and James S. Denford

Making Open Innovation Ecosystems Work: Case Studies in Healthcare by Donald E. Wynn, Jr., Renée M. E. Pratt, and Randy V. Bradley

Leadership

Best Practices for Succession Planning in Federal Government STEMM Positions by Gina Scott Ligon, JoDee Friedly, and Victoria Kennel

Risk

Risk Management and Reducing Improper Payments: A Case Study of the U.S. Department of Labor by Dr. Robert Greer and Justin B. Bullock

Ten Recommendations for Managing Organizational Integrity Risks by Anthony D. Molina

Using Technology

Using Artificial Intelligence to Transform Government by The IBM Center for The Business of Government and the Partnership for Public Service

Digital Service Teams: Challenges and Recommendations for Government by Professor Dr. Ines Mergel

Ten Actions to Implement Big Data Initiatives: A Study of 65 Cities by Alfred T. Ho and Bo McCall

The Social Intranet: Insights on Managing and Sharing Knowledge Internally by Dr. Ines Mergel

About the IBM Center for The Business of Government

Through research stipends and events, the IBM Center for The Business of Government stimulates research and facilitates discussion of new approaches to improving the effectiveness of government at the federal, state, local, and international levels.

About IBM Global Business Services

With consultants and professional staff in more than 160 countries globally, IBM Global Business Services is the world's largest consulting services organization. IBM Global Business Services provides clients with business process and industry expertise, a deep understanding of technology solutions that address specific industry issues, and the ability to design, build, and run those solutions in a way that delivers bottom-line value. To learn more visit ibm.com.

For more information:
Daniel J. Chenok
Executive Director
IBM Center for The Business of Government

600 14th Street NW
Second Floor
Washington, DC 20005
202-551-9342

website: www.businessofgovernment.org
e-mail: businessofgovernment@us.ibm.com

Stay connected with the IBM Center on:

or, send us your name and e-mail to receive our newsletters.

IBM Center for
The Business of Government