

THE FUTURE HAS BEGUN

USING ARTIFICIAL INTELLIGENCE TO TRANSFORM GOVERNMENT



PARTNERSHIP FOR PUBLIC SERVICE



**IBM Center for
The Business of Government**

The Partnership for Public Service is a nonpartisan, nonprofit organization that works to revitalize the federal government by inspiring a new generation to serve and by transforming the way government works. The Partnership teams up with federal agencies and other stakeholders to make our government more effective and efficient. We pursue this goal by:

- Providing assistance to federal agencies to improve their management and operations, and to strengthen their leadership capacity
- Conducting outreach to college campuses and job seekers to promote public service
- Identifying and celebrating government's successes so they can be replicated across government
- Advocating for needed legislative and regulatory reforms to strengthen the civil service
- Generating research on, and effective responses to, the workforce challenges facing our federal government
- Enhancing public understanding of the valuable work civil servants perform

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. Since its creation in 1998, the Center has awarded research stipends to public management researchers in the academic and nonprofit communities that have resulted in nearly 350 reports—all of which are available on the Center's website at businessofgovernment.org.

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.

PREFACE

There have been times in our history when new technological discoveries led to indelible changes in our lives. We are living in one of those times.

In hindsight, it is easy to identify Alexander Graham Bell's invention of the telephone in the 1870s as an instrument of marvel, eventually connecting people worldwide. And of course, there is the internet, which, although it burst into the public realm less than 30 years ago, is a technology and service that few can envision living without, whether we understood that in the 1990s or not.

Similarly, future historians may look back at the 2000s and 2010s as the point at which artificial intelligence forever changed how the world works, revolutionizing the way we perceive, think, reason, learn and make decisions. Additionally, AI has the potential to help address many of our country's pervasive problems and advance our safety, health and well-being.

That promise has great potential to transform government. The most challenging problems AI may help us solve—from fighting terrorists to serving vulnerable populations—will involve government. More immediately, though not less consequentially, AI will change the way public servants do their jobs.

Because of the impact AI could have on how governments ensure our safety and further our well-being, government leaders are likely to play a significant role in dictating future directions for AI.

The good news is that our federal government is in a position to do just that. The AI field is advancing rapidly at a time when government is looking for new ways to achieve its mission, and AI will change the rules of what is possible.

It is our hope that the broad ideas offered in the following pages, and the dialogue the Partnership and the IBM Center for The Business of Government plan to continue on this topic, spark a conversation on using artificial intelligence to modernize how government works, and improve the services it provides to citizens.



Max Stier
President and CEO
Partnership for Public Service



Daniel Chenok
Executive Director
IBM Center for The Business of Government

INTRODUCTION

With the promise of artificial intelligence no longer in some far-off future, federal leaders are delving into the possibilities. Agencies using it now range from the U.S. Citizenship and Immigration Services, which deploys an AI-based online virtual assistant to answer questions from citizens and immigrants, to the Intelligence Advanced Research Projects Activity, which has tapped the technology to improve facial recognition.

Other agencies are testing whether AI can improve the purchasing process or relieve employees of tedious work. And 77 percent of federal managers said their agencies will need to use more artificial intelligence over the next five years “to keep up with the increasing pace of work,” according to one survey.¹

AI has enormous potential for government. It can improve agencies’ effectiveness, make data more understandable and easier to use, and help citizens navigate government services. And it could save government up to 1.2 billion work hours and \$41.1 billion annually.²

No one sets out to use AI simply because it is available, according to the experts we interviewed. Government leaders working on agency issues, such as how to code data or buy goods and services more efficiently, turn to AI to streamline processes, relieve employees of tedious tasks and provide new insights into their agencies’ work.

In the pages that follow, we use the term artificial intelligence, the most widely recognized and adopted terminology, to refer to the use of computers that simulate human abilities and perform tasks that people typically do. Examples include reading documents to understand their meaning, looking at an image and recognizing the content, or making decisions. Related concepts and names include cognitive computing, predictive analytics, robotic process automation and machine learning.

One type of AI that three of the four organizations presented here used involves machine learning, which refers to systems that learn from data and improve over time. There is vast potential for machine learning in government, making the applications possible “too many to list,” said Alexander Measure, economist at the Bureau of Labor Statistics. The opportunities cover a wide range of activities, “from analyzing satellite images to processing Social Security disability claims,” he added.

1 ServiceNow, “Today’s State of Federal Automation.” Retrieved from <http://bit.ly/2gOmvHv>

2 Deloitte University Press, “AI-augmented government.” Retrieved from <http://bit.ly/2q2dfVp>

HOW GOVERNMENT IS USING ARTIFICIAL INTELLIGENCE

This issue brief describes the work of four very different organizations: two federal agencies—one defense and one civilian; one local government; and one university whose research was funded partly by two federal research and development agencies. Together, these four organizations highlight how AI technology assists agencies as they seek to transform the ways they work.

The cases presented here describe works in progress, not end results. Still, other agencies can learn, and likely benefit from these organizations' early experiences, particularly if these first stages end up being a springboard to significant shifts in agency practices. We also present insights, based on these stories, that can be guideposts for other government agencies interested in using AI.

The first case study shows how law enforcement agencies use or plan to use AI to outsmart criminals. Several U.S. law enforcement units, as well as wildlife rangers in two other countries, have used AI to plan their patrol routes to guard against two very different menaces: terrorism and poaching.

The second case study describes how the Bureau of Labor Statistics at the Labor Department plans to take away tedious, repetitive tasks from employees and save hundreds, even thousands of work hours. Employees' time can then be redirected to more important tasks.

The third case study illustrates how AI can help break down government silos, enabling agencies to serve citizens more effectively while also protecting vulnerable populations, such as substance abusers. A Kansas county is using the technology to identify people at risk of arrest, with the idea that getting people in these at-risk groups into county services could keep many of them out of jail.

The final case study explains how the Air Force plans to use the technology to make sense of complex acquisition regulations so it can speed the process of buying goods and services. Doing so could open government procurement to more small businesses and companies that have avoided working with agencies because the acquisitions process has been too difficult to navigate.

The researchers, AI experts and people described in these case studies, and numerous other people we interviewed, detailed their practical experiences with AI and what they learned. Their insights could help other leaders understand AI's potential and how it could be applied in their agencies.

METHODOLOGY

The pages that follow are based on in-depth interviews the Partnership for Public Service and the IBM Center for The Business of Government conducted from June to October 2017. We interviewed 14 people from 10 organizations ranging from government agencies to nonprofits to private universities. The professionals we interviewed were in charge of AI research, application or development.

We selected our case studies based on the availability of information and to cover a wide variety of government activities.

CASE STUDY #1

FIGHTING CRIME MORE EFFECTIVELY

Artificial intelligence is helping law enforcement agencies outsmart criminals by figuring out the best ways to deploy limited staff and other resources. AI makes it harder for terrorists and other criminals to predict where they will run into security personnel by adding unpredictability to homeland security and police activities. The same system helps wildlife rangers with another menace: poaching. Rangers in Africa use the technology to decide which wildlife areas to patrol on any given day to most effectively protect animals and plants.





A decade ago, a team at the University of Southern California created an AI system to help Los Angeles International Airport, and later the U.S. Coast Guard and the Transportation Security Administration, prevent terrorist attacks. Following the successful use of AI in homeland security, USC developed another version of the AI system to assist rangers with fighting wildlife poachers.

The Los Angeles Airport Police started using AI in 2007. An attack at Glasgow Airport that year prompted airports around the world to look for new ways to stop terrorists. One method was to make security schedules unpredictable, so terrorists could not anticipate where and when they would run into security checkpoints.

At the time, the Los Angeles Airport police did not have enough officers to staff checkpoints at all times on the eight roads that lead into the airport, according to Milind Tambe, engineering professor and AI expert at USC.³ Instead, airport police officials sought to allocate police in a way terrorists could not predict, according to the website of the USC team that created the software.⁴

That led to a partnership with USC to deploy software that can help determine potential targets, and then recommend randomized patrol routes and security schedules for police officers. In the years since, federal agencies have used different versions of the system. The Coast Guard uses it to randomize its boat patrol routes at major ports such as New York and Los Angeles, while the Transportation Security Administration uses it to assign air marshals to flights.

But the software does more than help law enforcement fight terrorism. In 2013, a presidential executive order on combating wildlife crime recognized illegal poaching as an escalating international crisis and directed the

departments of Agriculture, Treasury and State, among others, to help stop it.⁵

International wildlife crime is valued at \$8 to \$10 billion a year, according to the United Nations Office on Drugs and Crime.⁶ The U.N. ranks it among the most profitable illicit activities, along with the trafficking of humans, drugs and weapons. These crimes are pushing countless species to the brink of extinction, causing immeasurable harm to the Earth's biosphere and fueling instability that endangers U.S. interests at home and abroad, according to the executive order.

For example, between 2007 and 2014 alone, the African elephant population plummeted by 30 percent due to poaching, according to the Great Elephant Census, a project to get an accurate count of Africa's elephant population.⁷

Yet catching or deterring thieves who harm animals and plant species is typically a cat-and-mouse game, and it is not easy for wildlife rangers to figure out where the next hit will happen. "In most parks, ranger patrols are poorly planned, reactive rather than proactive, and habitual," said Fei Fang, a former University of Southern California Ph.D. student who worked on the team that developed the AI system, and current assistant professor at Carnegie Mellon University.⁸

When the rangers' rounds are predictable, poachers can avoid particular areas and commit their crimes far from the rangers' watchful eyes. Some parks are vast, and resources limited, and rangers fall back on their habits

3 Chris Lo, "Game theory: introducing randomness to airport security," *Airport Technology*, July 26, 2012. Retrieved from <http://bit.ly/2is7sBS>

4 University of Southern California Teamcore Research Group, "AI and game theory for public safety and security." Retrieved from <http://bit.ly/2gNSV88>

5 The White House, *Combating Wildlife Trafficking* (Executive Order 13648), July 1, 2013. Washington, D.C.

6 United Nations Office on Drugs and Crime, "Wildlife crime worth USD 8-10 billion annually, ranking it alongside human trafficking, arms and drug dealing in terms of profits: UNODC chief." Retrieved from <http://bit.ly/2ygrXJA>

7 Michael J. Chase et al., "Continent-wide survey reveals massive decline in African savannah elephants," *PeerJ* 4:e2354 (2016), accessed October 13, 2017, doi:10.7717/peerj.2354

8 National Science Foundation, "Outwitting poachers with artificial intelligence." Retrieved from <http://bit.ly/2oBRTLy>

when choosing where to patrol each day, said Tambe. That is an ineffective strategy, he added.

USC's development team recognized AI could strengthen wildlife conservation activities and save more animals by predicting where poachers will attack and enabling rangers to be there to thwart them. The National Science Foundation and the Army Research Office funded the university's research into modifying the AI software for this purpose. The result is the Protection Assistant for Wildlife Security software, or PAWS.

It has two components: The first uses historical data on poacher attacks, such as the location of snares, geographic and topographic data about a given area, and data on animal activity, to predict where poachers will attack next. The more data entered into the system, the better the software discerns patterns in how poachers operate.

Based on initial observations, the system indicated poachers seem to favor locations with certain topographical features, such as a specific slope of the land, or a set distance from roads, rivers or villages, places where they are less visible. It would have been time-consuming and difficult, perhaps impossible, for rangers to figure out the connection between topographical features in the region and the choices poachers were making about where to set up animal snares. And even though they now know some of what poachers look for, the AI technology assists rangers with deciding on the best patrol routes each day. "Machine learning is able to make excellent predictions on where poachers will set up traps," Tambe said.

The second component uses the predictions from the first component and applies game theory—the study of interactions between two decision-makers. The technology then randomly chooses patrol routes from several options, using its

calculations to recommend routes where poacher attacks are most likely to occur.

And hopefully, the route recommendations are highly unpredictable for poachers, who would have difficulty learning the rangers' patrol schedules and habits. Unlike before, poachers are less likely to figure out a rhyme or reason for the varying patrol routes.

Although PAWS has yet to be deployed globally, it was tested successfully in Uganda, where rangers found traps poachers set up to catch wildlife, and in Malaysia until funding ran out. "In a month, there were two locations where we thought rangers would find snares, and with PAWS we were able to show them dozens of antelope and elephant traps and other indications of active poachers there, potentially saving the lives of animals," Tambe said.

Although initial research found the wildlife security software to be "very effective" in making predictions about patrol routes, more research is needed to learn if it is reducing poaching in Uganda, the remaining pilot country, Tambe said.

Testing PAWS in the real world had its challenges due to the lack of large amounts of historical data on poaching. In general, the more data in an AI system, the better it will be at its task.

The team faced another challenge: difficult terrain, or "the street map of the jungle," according to Tambe.⁹ From their California lab, researchers initially failed to account for topography and terrain differences between sites in their two test countries.

The wildlife area in Uganda is mostly flat, whereas the wildlife area in Malaysia is rugged and mountainous, and has large changes in elevation. Rangers patrolling these areas

⁹ University of Southern California, "Artificial intelligence could turn poachers into prey." Retrieved from <http://bit.ly/2wR24yK>

"Machine learning is able to make excellent predictions on where poachers will set up traps."

MILIND TAMBE
ENGINEERING PROFESSOR
AND AI EXPERT, UNIVERSITY
OF SOUTHERN CALIFORNIA

need to use routes that preserve their energy so they can stay on patrol longer.

In the initial testing in both countries, the software would tell rangers to go directly from point A to point B without considering natural obstacles such as rivers, lakes or hills, and without accounting for paths rangers could take around obstacles. "We need to provide actual patrol routes that can be practically followed," Fang said.¹⁰

The testing in the actual patrol areas provided the topographical reality the team needed to improve the software. PAWS now comes up with "practical patrol routes that minimize elevation changes, saving time and energy," by using geographical and topographical data, even information on paths that animals seem to follow, and learning as it powers through more and more data.¹¹

¹⁰ National Science Foundation, "Outwitting poachers with artificial intelligence." Retrieved from <http://bit.ly/2oBRTLy>

¹¹ Ibid.

CASE STUDY #2

MAKING TEDIOUS TASKS A THING OF THE PAST

The Bureau of Labor Statistics at the Labor Department plans to use artificial intelligence to relieve employees of tedious, repetitive tasks and save hundreds of work hours. Bureau staff have to read and study hundreds of thousands of survey responses about workplace injuries and illnesses each year to understand and help prevent them. AI technology evaluates responses faster than a person can, and enables bureau staff to work on more important tasks, including follow-up on the survey results.



The Bureau of Labor Statistics collects data on workplace injuries from a sample of around 200,000 businesses through its annual Survey of Occupational Injuries and Illnesses. It then must assign codes that correlate to particular survey responses.

In 2015, there were 2.9 million reported private sector injuries and illnesses in the workplace and 752,600 reported by the public sector, according to the survey.¹² Understanding why and how these injuries and illnesses happen can help the bureau tell companies and governments how to prevent them. But reading hundreds of thousands of survey responses to pick out the important details to code is a repetitive, time-consuming process, and one that does not require critical thinking.

Employees of the bureau's Occupational Safety and Health Statistics program choose codes to assign to pieces of information, whether it is a code designating that a respondent is a nurse or one that indicates, say, an arm injury. They then read through the survey responses and assign the codes.

Before they could start doing this job, employees had to learn the many different codes and the procedures for using them. It was "very hard for people to remember all the various rules for how these codes are assigned," said Alexander Measure, a BLS economist.

Once they had the codes and the rules down, program team members had to go through about 300,000 incident narratives each year. It took about 25,000 work hours to read and code, and then more time to double-check the coding and correct mistakes—until the bureau found a way to automate the coding process and improve the quality and efficiency of the coding tasks.¹³

12 Bureau of Labor Statistics, "Employer-Reported Workplace Injury and Illness Summary." Retrieved from <http://bit.ly/2y9kPIZ>

13 Alexander Measure, "Automated Coding of Worker Injury Narratives." Retrieved from <http://bit.ly/2yiUz68>

In 2014, the bureau started using AI to code responses. It started small. The first survey year, computers assigned 5 percent of all codes, and they coded only the occupations in which the injury or illness occurred, as these are the easiest codes to assign. The agency called these "slam dunks," said John Ruser, the bureau's former assistant commissioner for safety, health and working conditions, and later associate commissioner for productivity and technology. For example, they assigned the code "registered nurse," no matter if someone put "RN" or "R.N." as the occupation.

By 2016, the most recent survey year, computers assigned nearly 50 percent of all codes, and these codes related not just to occupations, but to the nature of the injury or body parts affected.¹⁴

Even better, the bureau found from the start that the computer coded more accurately, on average, than a trained human coder, Measure said. The increase in quality and accuracy convinced the bureau that machine learning could make the process more effective. "We were able to improve the quality of coding we are producing, which means better estimates and better understanding of work-related injuries and illnesses in the United States," Measure said. And the work gets done faster because the computer can do in one day what it would take a human a month to do, he added.

To allay employees' apprehension about technology, in the form of machine learning, taking over some of their job responsibilities, the bureau's leaders played an active role in communicating the change and the benefits of using AI. The commissioner, assistant commissioner and associate commissioner—together with Measure, who was the first employee on the coding project

14 Bureau of Labor Statistics, "Automated Coding of Injury and Illness Data." Retrieved from <http://bit.ly/2wmKlQs>



team who had AI expertise—underscored to staff that the bureau could achieve its mission more effectively and benefit millions of workers around the country.

The technology "can provide more comprehensive, ongoing and timely surveillance to inform future injury prevention policy and practice," according to the bureau.¹⁵ It also could benefit bureau employees themselves. They could now focus on more complicated cases that require human judgment, shifting from mind-numbing to more interesting tasks, and increasing the "quality of their work," Ruser said. Staff also could do more follow-up work, such as calling respondents to get clarifications on their survey responses.

At the same time, the bureau held training sessions for coding staff on what machine learning is and what it does, to help them get a better understanding of its value to the bureau.

15 Kirsten Vallmuur et al., "Harnessing information from injury narratives in the 'big data' era: Understanding and applying machine learning for injury surveillance," *Injury Prevention: Journal of the International Society for Child and Adolescent Injury Prevention*, 22(Suppl 1) (2015), accessed October 26, 2017, doi:10.1136/injuryprev-2015-041813

CASE STUDY #3

HELPING THE NATION'S MOST VULNERABLE POPULATIONS

Artificial intelligence can help break down government silos, enabling agencies to serve citizens more effectively. A Kansas county is using artificial intelligence to protect its most vulnerable populations by merging and analyzing data from different county departments to determine who is most likely to be arrested and incarcerated. From this data pool, AI identifies at-risk residents who are not taking advantage of available services so the county can assist them in ways that could keep them out of jail.





“Government should be preventing crises, not just reacting to them.”

ROBERT SULLIVAN
CRIMINAL JUSTICE COORDINATOR
JOHNSON COUNTY, KANSAS

Johnson County, Kansas, was struggling to keep low-level offenders from incarceration when the White House launched the Data-Driven Justice Initiative in 2016. It is a partnership among federal, state and local governments, and the private and nonprofit sectors, to use strategies based on data to reduce incarceration rates, especially among vulnerable populations, such as those suffering from mental health, substance abuse and chronic health issues.¹⁶

Johnson County officials responded to the initiative by turning to the machine learning type of artificial intelligence, hoping to divert people from jail by providing them with county services.

In the area of mental health alone, data showed more than 50 percent of people in jails and prisons around the country had mental health issues, for example, according to Bureau of Justice Statistics data from 2005, the most recent year a comprehensive study was done on the issue.¹⁷

The county partnered with the University of Chicago to apply the technology to the data on people who

interacted with both the county’s mental health center, and jail and court systems, between 2010 and 2014. The data came from three sources: the emergency medical services department, the mental health center and the county’s joint database of law enforcement, court and corrections data.

For an initial pilot program, the team identified 200 people among those most likely to be incarcerated. The machine learning system based its predictions on 252 different types of information from the combined data set, ranging from demographics to enrollees in mental health programs to the number of times people had been arrested.

After the initial data analysis, the county used 2015 data to check whether the 200 individuals identified had indeed been incarcerated. The county at first was disappointed with AI’s success rate. Only a little more than half the people the AI system identified—102 of 200—ended up behind bars. In 2015, this group spent a total of 7,000 days in jail.

However, the system’s prediction was 500 percent more accurate than if people had been randomly identified as being at risk, according to the University of Chicago.¹⁸ And

the prediction is 25 percent more accurate than if case workers in the field had used one or two features to identify at-risk people, which is what typically happens.¹⁹ The results seem even more powerful considering the system used data from just three county service departments.

In the pilot program, the technology identifies people already in jail and, once they are released, county departments contact them to offer services. When the system is fully rolled out later in 2018, the county plans to do more to intervene and connect people with services, hopefully making them less likely to go to jail. “We were tired of reacting to problems,” said Robert Sullivan, criminal justice coordinator at Johnson County. “Government should be preventing crises, not just reacting to them.”

Initial success depended on trusting relationships among the three county departments participating in the project: mental health, emergency medical services and jails.

The county is now working to introduce data from more sources, including the two largest police stations and the county’s probation department. Officials hope this enables the county to help more citizens and prevent more family crises.

16 The White House, “Fact Sheet: Launching the Data-Driven Justice Initiative: Disrupting the Cycle of Incarceration.” Retrieved from <http://bit.ly/2zIHxor>

17 Bureau of Justice Statistics, “Mental Health Problems of Prison and Jail Inmates.” Retrieved from <http://bit.ly/2eRfbf2>

18 Erika Salomon et al., “Reducing Incarceration Through Prioritized Interventions,” *Proceedings of ACM SIGKDD*, Halifax, Nova Scotia, Canada (2017), accessed October 12, 2017

19 Ibid.

CASE STUDY #4

CONQUERING THE COMPLEXITIES OF FEDERAL PURCHASING

The Department of the Air Force plans to use artificial intelligence to help acquisition professionals make sense of complex acquisition regulations and speed the process of buying goods and services. The department will upload thousands of regulations, contract cases, acquisition training material and Defense Department policy. AI technology then will be able to answer queries from federal contract officials and contractors about acquisition rules and regulations, such as how to proceed with a contract, what procedures to follow or what contract a small business could bid on.

Every federal agency and branch has dealt with the complicated acquisition process for obtaining goods and services. Many federal workers and contractors report it is daunting to decipher the thousands of pages of intricate federal and defense acquisition regulations, let alone become familiar with them.

For the Department of the Air Force, a huge government purchaser, the challenge is exponentially larger than for many other agencies. In fiscal 2017 alone, the department spent around \$53 billion on products and services, or 11 cents of every dollar the federal government spent on acquisitions that year. AI could help the Air Force smooth the acquisition process by helping officials figure out rules and regulations and make good contract decisions more quickly and efficiently.

The department is now running a pilot project, working with two contractors, and hopes to unveil the AI system both online and as a phone application starting later in 2018.^{20,21}

Department employees and contractors are uploading a massive amount of data: information from Federal Acquisition Regulations and the Defense Federal Acquisition Regulation Supplement; other laws and rules; institutional knowledge in the form of past contracts and requests for proposals; Defense Acquisition University guides; and the requirements of the Joint Capabilities Integration and Development System process, which identifies the capabilities the Defense Department needs and how those capabilities will be evaluated.²²

20 In the interest of full disclosure: The Air Force is working with two small businesses on this project, both of which use IBM's Watson system.

21 Air Force Office of Transformational Innovation. "Cognitive Computing Application for Defense Contracting." Retrieved from <http://bit.ly/2iBn378>

22 Ibid.



This effort requires major effort before the system can be helpful. "It's not magic," said Frank Koniczny, the Air Force's chief technology officer. Once the Air Force completes this task, the system should be able to answer requests from department employees and contractors.

The goal is to enable a contract officer to query how a specific contract should be structured or if a particular contract type can be used for buying a particular product or service—whether that is an incentive contract, a fixed-price contract or some other contract type. "Contracting officers could quickly find answers to difficult questions so that they can focus on creating agreements and use the flexibility available in the procurement regulations," according to the Air Force.²³

And any business interested in contracting with the Air Force, from a large multinational corporation to a small startup, would be able to query the software about bidding on a contract. A company also could request a list of all the contracts it is eligible to bid on, options that

23 Ibid.

could simplify the process for those bidders.

At the same time, the AI technology can be taught to analyze text to get at the regulation's meaning. Human experts can help the system learn to discern the intent of regulations by correlating words, acronyms or phrases to an explanation. The acquisition process is confounding, in part, because it is difficult to understand the regulations. "Acquisition regulations and policies form an intricate maze that is impossible to navigate, raise costs, and prolong the time to deliver capabilities," according to an Air Force website.²⁴

With artificial intelligence, however, "you have a true representation of the terms and conditions" of a contract, Koniczny said. Acquisition professionals tend to use the contract type they are used to, rather than the most appropriate one—even if it is not the best or most efficient contract for their purchase—because they know how to follow the regulations for that specific contract. "They know it's right,"

24 Department of the Air Force, "AQ-Prime (AQ) - Cognitive Computing." Retrieved from <http://bit.ly/2y6MnEa>

“The challenge of operating in accordance with complex federal acquisition regulations discourage[s] small and innovative businesses from partnering with the government.”

GOVERNMENT ACCOUNTABILITY OFFICE

and “that’s why people go back to what they are currently familiar with,” Konieczny said. AI could help ensure Air Force professionals use contracts expeditiously and comply with all regulations, even when they use a new and unfamiliar contract type.

An acquisition process using AI software could also help company employees understand what they need to know for a contract application and what acquisition regulations apply when submitting a proposal, without the assistance of lawyers. It could also help businesses identify what regulations they must follow based on their type of work.

Additionally, making contracting easier could lead to new and innovative services and technologies for our country as more first-time contractors, small businesses and others become willing to bid for defense contracts. Small businesses and other companies can be deterred by the difficulties involved in federal acquisitions, leading government agencies to lose small-business bidders. “The challenge of operating in accordance with complex federal acquisition regulations

discourage[s] small and innovative businesses from partnering with the government,” according to the Government Accountability Office.²⁵


The federal government has long sought to expand its contract work beyond the large companies it usually works with, but has struggled to succeed. In fiscal 2017, for example, only 21 percent of federal contract spending, and 17 percent of the Department of Defense’s contract spending, went to small businesses. Expanding the number of potential bidders could help the department supply better products and services to the war fighter faster, and at lower cost.

A successful pilot would give other agencies an AI approach to emulate to simplify and speed up what is now a mystifying government necessity. The acquisition workforce no longer will have to spend hours trying to find and make sense of regulations, and employees would have more time to focus on other mission-critical work.

²⁵ Government Accountability Office, *Highlights of a GAO Forum: Managing the Supplier Base in the 21st Century*, GAO-06-533SP, March 2006, 7. Retrieved from <http://bit.ly/1bJ5CH0>

STEPS FOR GETTING STARTED

The people we interviewed shared several insights for government leaders seeking to use AI. Most of their ideas focused on the transformation their organizations went through when starting to use AI.



Not every task should be augmented by artificial intelligence.

Agencies and project teams should first discuss what role artificial intelligence could play in their work, what tasks could AI make easier and what outcomes they expect AI to help them achieve. AI is not a silver bullet, and it is not appropriate for every challenge.



Do not underestimate the upfront investment needed.

Once agencies and project teams identify areas where AI can help them achieve their missions, they need to consider the resources they will need, including experts with knowledge of AI systems and how to use them, and budgets to support implementation of the technology.

Agencies should also consider how much staff time will be necessary to get an AI system up and running, especially in cases where employees must upload a trove of data and information.



Start small.

Artificial intelligence, like most new technologies, is best tested on a small scale before it is deployed fully. Using a pilot program enables people to get familiar and comfortable with the technology and catch errors and correct course. And it enables the system to improve.

Johnson County, Kansas, for instance, started its work with a data set that contained only a couple different types of information on its target population, such as name and date of birth, and data from only three of several dozen departments. It is still in the pilot phase. “Start small. You don’t have to have a grand plan” at the very beginning, said Steve Yoder, GIS project developer at the county.



It is always about the data.

AI is data hungry. One of the most common challenges with using AI is data access, availability and quality. The more and better quality the data, the better its performance and accuracy. However, most government data and information is contained in separate agencies and, in many cases, the data is limited.

All agencies should ensure quality data and information are available for training, testing, using and refining AI systems. The Partnership and IBM's 2011 report, "From Data to Decisions: The Power of Analytics," highlighted the importance of data sharing and accessibility for making decisions.²⁶ This will become more crucial as AI spreads through government.

The first step could be adopting government-wide standards for different types of data, in the same way the Digital Accountability and Transparency Act of 2014 sets standards for government spending data to be more uniform and readily accessible.



Agency expertise in artificial intelligence could boost AI's potential.

Agencies will need a robust federal AI workforce to manage the growth and potential of these technology systems. These experts could serve as the repository of AI knowledge for agency programs, and could work directly on projects when teams lack AI expertise.

However, agencies likely will encounter challenges with attracting AI experts, as they have with cybersecurity experts. Therefore, they should prepare for a probable shortage of AI talent in government, and look for ways to work with AI experts in the private sector and academia.



Government could work with outside experts, particularly at colleges and universities.

Colleges and universities have a tremendous amount of artificial intelligence expertise and ongoing research and development programs and projects, and some have designated AI departments, such as the University of Southern California's Center for Artificial Intelligence in Society.

Agencies could conceivably realize an added benefit if, while working with AI departments, public service piques the interest of college and university students and researchers who could take their skills to the public sector.

²⁶ Partnership for Public Service and IBM Center for The Business of Government, *From Data to Decisions: The Power of Analytics*, November 2011. Available at <http://bit.ly/2hwvA7y>

CONCLUSION

Artificial intelligence has captured the imaginations of many people. For some, AI promises to solve significant and long-standing problems. For others, it presents a threat. For others still, the symbiosis between artificial intelligence and human beings is the next step in the evolution of our species.

But virtually all agree that artificial intelligence will have a lasting impact on our lives.

Therefore it is vital for government to make a strategic investment in understanding how to maximize AI's benefits and use it to improve agencies and government as a whole. Some federal organizations remain unaware of the opportunities AI presents, and how they can realize the possibilities of this growing field—a lack of understanding that is increasingly likely to put them at a major disadvantage.

Questions remain pertaining to information privacy and cybersecurity, AI systems' trustworthiness and reliability, and the role of these systems in operational and organizational transformation. Other issues to address include the prevention of unwanted bias in AI as well as the use of AI to counteract human bias. Perhaps the most complex question to deal with is what AI's effects will be on employment, and the jobs that AI is likely to both create and eliminate.

Before government can reap the benefits of AI, it should seriously consider unintended consequences, employees' apprehension about adopting AI and other topics beyond the scope of the research we did for this paper. As the Partnership and the IBM Center for The Business of Government continue the AI dialogue, we hope to present our findings on these additional issues of interest and concern.

Artificial intelligence is more than a technology. It is a road to transformation. In the coming years, AI will become more routine in government. Although there will be challenges and detractors, AI has the potential to make data more understandable and easier to use, help citizens navigate government services, allow agencies to respond to threats and crises more thoroughly, and improve agencies' overall effectiveness.

ACKNOWLEDGEMENTS

The individuals listed below generously offered their input on ongoing artificial intelligence projects at various agencies and organizations, and on how AI could improve government's effectiveness. We greatly appreciate their time and counsel. However, the contents of this report do not necessarily reflect the views of those we spoke with.

PRIVATE AND NONPROFIT SECTORS

Alison Hurst

Founder and Executive Director, Safe Place for Youth

Eric Rice, Ph.D.

Founding Co-Director, CAIS Center for Artificial Intelligence in Society, University of Southern California
Associate Professor, Suzanne Dworak-Peck School of Social Work, University of Southern California

John Ruser, Ph.D.

President and CEO, Workers Compensation Research Institute
Former Assistant Commissioner for Safety, Health and Working Conditions and Associate Commissioner for Productivity and Technology, Bureau of Labor Statistics

Milind Tambe, Ph.D.

Founding Co-Director, CAIS Center for Artificial Intelligence in Society, University of Southern California
Helen N. and Emmett H. Jones Professor in Engineering, University of Southern California

Jin Yao

Human Services Data Analyst

Steve Yoder

GIS Project Developer

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

Steve A. Chien, Ph.D.

Technical Group Supervisor, Artificial Intelligence Group,
Jet Propulsion Laboratory
Senior Research Scientist, Mission Planning and Execution
Section, Jet Propulsion Laboratory

OFFICE OF THE DIRECTOR OF NATIONAL INTELLIGENCE

Christopher Boehnen, Ph.D.

Senior Program Manager, Intelligence Advanced Research
Projects Activity

DEPARTMENT OF COMMERCE

Michael Garris

Senior Scientist, National Institute of Standards and
Technology
Co-Chair, Subcommittee on Machine Learning and
Artificial Intelligence, National Science and Technology
Council

DEPARTMENT OF DEFENSE

Frank Konieczny

Chief Technology Officer, Department of the Air Force

DEPARTMENT OF LABOR

Alexander Measure

Economist, Bureau of Labor Statistics

GOVERNMENT OF JOHNSON COUNTY, KANSAS

Chris Schneweis

Senior Management Analyst

Robert Sullivan

Criminal Justice Coordinator

PROJECT TEAM

PARTNERSHIP FOR PUBLIC SERVICE

Mallory Barg Bulman, Vice President, Research and Evaluation

Samantha Donaldson, Director, Communications

Kelsey Goetz, Associate Manager

Peter Kamocsai, Associate Manager and Project Lead

Ariel Leuthard, Intern

Ellen Perlman, Writer and Editor

Max Stier, President and CEO

Jaimie Winters, Associate

David Xue, Intern

IBM

Daniel Chenok, Executive Director, IBM Center for The Business of Government

Haynes Cooney, Research Program Manager, IBM Institute for Business Value

Claude Yusti, Federal Cognitive Solutions Leader, IBM Global Business Services

The Partnership for Public Service's work, including this report, would not be possible without the generous support of corporations, foundations and individuals who share our commitment to more effective government. Corporations that support our research and thought leadership provide financial support and valuable expertise on a wide array of government management issues. By enlisting a diverse group of donors, the Partnership ensures that no single person, entity or industry can unduly influence our organizational body of research. The Partnership is actively committed to transparency about all of our funding relationships and retains editorial control over all its thought leadership.



**PARTNERSHIP
FOR PUBLIC SERVICE**

1100 New York Avenue NW
Suite 200 East
Washington DC 20005

(202) 775-9111
ourpublicservice.org
CFC# 12110



**IBM Center for
The Business
of Government**

600 14th Street NW
Second Floor
Washington DC 20005

(202) 551-9342
businessofgovernment.org