Federal **AI** Adoption

Where It Stands, Where It Is Heading,
and the Opportunities and Challenges Ahead

President's Letter

The Professional Services Council Foundation is pleased to present this report assessing the state of adoption of Artificial Intelligence (AI) technologies and practices in federal government operations, and the challenges and opportunities that lie ahead. We have long observed the interwoven relationship between government technology and service delivery, by federal agencies, government personnel, and their contractor partners. Perhaps no area better exemplifies this interrelationship than AI. At the same time, rapidly evolving technologies present new questions and considerations with unique and significant implications for federal missions, functions, and service delivery. This report does not seek so much to answer these questions as to illuminate them, and to advance a conversation about applications of Artificial Intelligence by and for federal government.

We gratefully acknowledge the support and participation of all those who made this report possible.

David J. Berteau
President & CEO
Professional Services Council Foundation

May 2019

Table of Contents

Overview: Federal AI Adoption Opportunities & Use Cases

AI presents vast opportunities for the government

Reduce
costly administrative and manual burdens and free employees for more valuable work

Optimize
resource allocation

Accelerate
and improve decision making

Reduce
or eliminate backlogs or address rapidly increasing workloads

Extract
value from vast stores of collected data

Combat
inefficiency and fraud

Improve
programmatic performance across wide span of mission domains

Existing Government Use Cases

DoD uses AI to help plan deployments during crises.

GSA deploys robotic process bots to track invoicing and payments.

NASA uses bots to aid in its finance and procurement processes.

HHS employs AI to identify ways to consolidate contract vehicles.

Expected Future Government Use Cases

USDA is looking to employ chatbots to augment the human workforce at call centers.

NIST is looking at AI to assist with research and discovery into various science and technology domains.

VA is exploring ways AI can improve customer experience.

USPTO is examining AI to route patent applications more efficiently to examiners.

Where It Stands, Where It Is Heading, and the Opportunities and Challenges Ahead

Introduction

Artificial intelligence, or AI, is here. It is in our smart phones, our web browsers and search engines, our digital assistants, and in other digital tools we employ to conduct our daily work and personal tasks.

Although commercial companies are most often viewed as being at the forefront of AI adoption, the federal government is also making strides in its own AI adoption journey. Like commercial enterprises, government agencies are lured by AI's vast potential to reduce costs and improve performance, decision-making, service-delivery, and customer experience. The opportunities to exploit AI in federal agencies seem limited only by one's imagination and available funding.

But along with the promise of enormous benefits that AI can bring come potentially significant challenges. There are systematic challenges, such as workforce disruptions, ethical considerations, policy implications, and governance issues. There are also readiness issues, such as ensuring adequate data protection, sound data management and implementation strategies, requisite technical skills, and effective interfaces between machines and their human cohorts. All will require persistent attention from government leaders and decision-makers as federal AI adoption advances.

There has been much written already about AI and its potential to address many government needs. This paper does not aim to revisit those discussions. Rather, our intention is to focus on how government agencies today are actually incorporating AI technologies into their operations to improve mission and business performance and boost efficiencies, the approaches they employ to do that, as well as the lessons learned and suggested best practices they have accumulated along the way. In addition, we look at some of the resulting or anticipated implications and consequences arising from federal AI adoption and the questions and challenges they may pose for government decision-makers going forward.

To better understand the state of federal AI adoption today and where it is heading tomorrow, the PSC Foundation interviewed numerous federal AI project leaders and practitioners as well as industry experts who are working alongside federal trailblazers.[1]

This report is the product of those conversations as well the Foundation's independent research and analysis. All views presented are those of the PSC Foundation and do not necessarily reflect the views of any participants interviewed.

---

1 Our discussions adhered to Chatham House rules, meaning the information provided will not be attributed to any individuals by name.

Defining AI

When asked to define AI, one senior congressional staff member we interviewed replied half-jokingly, "I've been told by some very smart people to stay away from that. I think if we begin our discussion with that, we may never move on to the next step."

Stanford professor John McCarthy, who coined the term "artificial intelligence" in 1955, defined it as the "science and engineering of making intelligent machines, especially intelligent computer programs."

The Government Accountability Office (GAO) acknowledged in recent testimony that "there is no single universally accepted definition of AI, but rather differing definitions and taxonomies."[2] It adds that researchers have distinguished between 'narrow' and 'general' AI, with narrow AI referring to applications that provide domain-specific expertise or task completion, and general AI referring to "an application that exhibits intelligence comparable to a human, or beyond, across the range of contexts in which humans interact."[3] The AI we see in existence today, and that we refer to in this report, generally falls into the narrow AI category.

Terms that are often used interchangeably with AI are machine learning (ML) and deep learning. It is important to understand they are not the same thing, although they are related. Deep learning is a subset of machine learning, and machine learning is a subset of AI. While AI refers broadly to the use of computers to mimic human cognitive functions (for example, computer programs that play strategy games such as chess or Go), machine learning refers more specifically to the ability of machines to "learn" for themselves by altering their own algorithms as they are exposed to ever-larger sets of data. A common application of machine learning can be found in our email filters, which learn over time which emails to put in spam or other folders.

This report uses AI in its broadest sense, to include categories such as machine learning, natural language processing, expert systems, computer vision, speech, planning, and robotic process automation (RPA).

2 Testimony of Timothy M. Persons, Chief Scientist, Applied Research and Methods, GAO, before the Subcommittees on Research and Technology and Energy, Committee on Science, Space, and Technology, House of Representatives, June 26, 2018. Available at: https://www.gao.gov/assets/700/692793.pdf

3 ibid.

The State of Federal AI Adoption Today

DRIVERS AND USE CASES

One doesn't need to think long about AI and its vast array of capabilities to glimpse its potential for transforming federal government operations. Various industry reports and studies have already identified hundreds of potential AI use cases for federal organizations. Nor does one have to look far to find federal agencies incorporating AI into their plans and operations.

At its core, AI is about automating and augmenting tasks that would normally require some degree of decision-making or intellectual capacity so that enterprises are more empowered to advance their goals. So it is no surprise that many agencies are moving aggressively to explore and incorporate AI into their operations. As one agency technology official told us: "If you have a science, a technology, a development shop, odds are they're exploring some sort of machine learning type of algorithms or application today." Many federal organizations are motivated by pressing desires to:

Reduce costly administrative and manual burdens and free up employees for more valuable work;

Optimize resource allocation;

Accelerate and improve decision-making;

Reduce or eliminate backlogs or address rapidly increasing workloads;

Extract value from vast stores of collected data;

Combat inefficiency and fraud;

Improve programmatic performance across a wide span of mission domains;

and more.

Particularly notable is the wide variety of current AI use cases we already see in practice today across government. For example:

The Defense Department (DoD) uses AI tools to help plan force deployments during crises, recognize objects in video data, and perform predictive maintenance on Bradley Fighting Vehicles, among other things. It is also building an AI capability to check and verify security clearance applications against multiple data sources to speed up the process.

The General Services Administration (GSA) has deployed 10 robotic process automation (RPA) bots and plans to deploy dozens more. One bot, for example, identifies invoices coming due for payment and alerts those responsible to ensure compliance with the Prompt Payment Act requiring that vendors are paid within 30 days. Another bot extracts key data from its leases to ensure they are reconciled with data in other systems.

The National Aeronautics and Space Administration's (NASA) Shared Services Center has four bots running nine different processes, including distributing funds, procurement, documenting images, scanning files, and creating folders to establish grants packages.

The Health and Human Services Department (HHS) is employing AI to identify opportunities to consolidate contract vehicles.

Agencies are exploring many other potential applications for AI. For example:

The Department of Agriculture (USDA) is looking to employ chatbots or AI to augment the human workforce at its call centers.

The Veterans Affairs Department (VA) is exploring ways in which AI can improve the customer experience and service delivery for its veteran constituents.

The National Institute of Standards and Technology (NIST) is looking at having AI assist it with research and discovery into various domains of science and technology.

The Patent and Trademark Office (USPTO) is examining whether AI can help it route patent applications more efficiently to the correct examiners as well as to help examiners search for "prior art," which is evidence that a purported invention is already known.

WHERE WE ARE TODAY

Keith Nakasone, deputy assistant commissioner for GSA's acquisition operations, told Congress in 2018 that "from our unique government-wide perch at GSA, [we have seen] that agencies are tremendously interested in AI and other emerging technologies—not because they are the latest fad, but because people recognize the potential to transform and simplify the way Americans interact with their government."[4]

Our own research and interviews with federal leaders and practitioners confirm this. Yet, while interest in AI is high, most federal agencies today are still in the early stages of AI adoption. They are identifying existing problems and challenges that may be suitable for AI applications; preparing their data environments; launching pilot projects and refining them. Some are deploying AI solutions as fully-fledged components of their day-to-day operations.

Several federal officials we interviewed acknowledge the great potential that AI offers, but nevertheless view AI with a wary eye, questioning whether today's AI is ready for prime time and can be trusted to operate in their specific environments—or even whether government is ready to implement the capabilities that are available.

Agencies also are under severe and increasing budget pressures—particularly for development, modernization, and enhancement (DME) expenditures—so many investments in this area will require the promise of clear-cut and quantifiable returns on investment (ROIs). Consequently, agencies are choosing their opportunities to explore AI capabilities very selectively.

Conversely, there are a few large agencies—encouraged by recent dramatic advancements in autonomous systems, computer vision, natural language processing, and game playing—that have staked out far more aggressive paths.

In its *Augmenting Intelligence Using Machines (AIM) Strategy*, the Office of the Director of National Intelligence noted that "these [types of] AI systems can perform tasks significantly beyond what was possible only recently, and, in some cases, even beyond what humans can achieve."[5] The strategy offers a framework for incorporating artificial intelligence, automation, and augmentation technologies to accelerate mission capability across the 16 agencies that comprise the Intelligence Community (IC). "In light of these recent advances," the strategy said, "the IC is carefully considering methods for fully automating well-defined processes and augmenting human expertise with analytics or planning capabilities for their potential benefit."

Intelligence agencies and the Defense Department—which face intense modernization competition from China and Russia and are struggling to consume and leverage inordinate volumes of sensor and intelligence information—have identified AI as a core component of their future operations.

4 Statement of Keith Nakasone, Deputy Assistant Commissioner, Acquisition Operations, Office of Information Technology Category (ITC), U.S. General Services Administration, Before the Subcommittee on Information Technology of the Committee on Oversight and Government Reform, March 7, 2018. Available at: https://republicans-oversight.house.gov/wp-content/uploads/2018/03/Nakasone-GSA-Statement-AI-II-3-7.pdf

5 January 16, 2019. Available at: https://www.dni.gov/index.php/newsroom/reports-publications/item/1940-the-aim-initiative-a-strategy-for-augmenting-intelligence-using-machines

DoD's first big step occurred in June 2018 with the creation of the Joint Artificial Intelligence Center (JAIC) to accelerate delivery of AI capabilities, scale the impacts of AI capabilities across the department, and synchronize AI activities. DoD followed up with its first Defense Department Artificial Intelligence Strategy. A summary of that strategy, released in February 2019, noted that "AI is poised to transform every industry and is expected to impact every corner of the Department, spanning operations, training, sustainment, force protection, recruiting, healthcare, and many others."[6]

Likewise, many federal data, science, technology, and research-oriented agencies—including the Departments of Energy, Agriculture, Commerce, and Health and Human Services, NASA, and the National Science Foundation—are making significant investments in AI adoption and research, some of which are outlined in the February 2019 Executive Order on Maintaining American Leadership in Artificial Intelligence.[7]

POLICY

The PSC Foundation expects that as more agencies explore and experiment with AI capabilities—and as commercial AI capabilities continue to advance at an accelerating rate—we will see a more robust adoption of AI tools and approaches across government. Some recent policy developments and trends also lead us to anticipate a marked acceleration of federal AI adoption. These include:

→ The Modernizing Government Technology (MGT) Act, enacted in December 2017, which authorizes new working capital funds with which agencies can invest in modern technology solutions to improve service delivery to the public, secure sensitive systems and data, and save taxpayer dollars. Because any funds used for modernization must be repaid, priority is given to modernization projects offering a clear return on investment, which should favor the cost-saving potential of many AI applications.

→ The President's Management Agenda, unveiled in March 2018, which outlines many cross-agency performance (CAP) goals—as well as detailed quarterly measures of progress for each—that will likely encourage greater federal adoption of AI tools and approaches. CAP goals include: modernizing IT to increase productivity and security; leveraging data as a strategic asset; improving customer experience with federal services; shifting from low-value to high-value work; getting payments right; and leveraging common contracts and best practices to drive savings and efficiencies.

---

6 SUMMARY OF THE 2018 DEPARTMENT OF DEFENSE ARTIFICIAL INTELLIGENCE STRATEGY: Harnessing AI to Advance Our Security and Prosperity. February 12, 2019. Available at: https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMA-RY-OF-DOD-AI-STRATEGY.PDF

7 EO 13859: Maintaining American Leadership in Artificial Intelligence. February 11, 2019. Available at: https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/

- The Cabinet-level National Science and Technology Council (NSTC) Select Committee on AI, formed in May 2018, which advises the White House on interagency AI research and development priorities and coordinates those activities; considers new federal partnerships with industry and academia; and identifies opportunities to leverage federal data and computational resources to support AI research and development.

- Increasing bipartisan congressional interest in AI. There has been steadily increasing interest in and attention to AI promotion, research, and adoption in government. One example was the introduction by a bipartisan group of senators in 2018 of the Artificial Intelligence in Government Act (S. 3502, 115th Congress), that would promote the use of AI across federal agencies and the creation in the House and Senate of separate AI caucuses.

- Increasing attention to enabling AI procurement. In October 2018, GSA began providing agencies with automation technology to service their contact centers through a new Special Item Number on existing GSA schedules. The new Automated Contact Center Solutions SIN offers a range of automated and bot-technology solutions from more than 4,400 vendors. GSA estimates that the federal government receives 750 million public inquiries annually and spends $17.5 billion on contact center services to field those calls.

- Explosive growth in the use of Other Transaction Authority (OTA) purchasing arrangements across DoD to procure emerging technologies, including AI. According to market analyst Bloomberg Government, DoD spent $4.2 billion on OTA agreements in 2018,[8] an increase of nearly $2 billion from the previous year, and the PSC Foundation anticipates continued significant federal spending on these arrangements and growth in the deployment of AI solutions under those awards.

- The Foundations for Evidence-Based Policymaking Act, enacted in January 2019, which creates Statistical Officials and Chief Data Officers at each of the 24 CFO Act agencies, and requires that they prepare their agencies' datasets to be machine-readable, secure, protected, and accessible. This law, coupled with the required actions of the recently updated Federal Data Strategy, should enable agencies to more easily incorporate AI-based tools and approaches into their back-office, citizen-facing, and other mission-supporting operations.

- The Executive Order on Maintaining American Leadership in Artificial Intelligence, which notes that "artificial Intelligence will affect the missions of nearly all executive departments and agencies," and directs them to make sustained investments in AI research and development, promote openness to data for AI applications, lower barriers and set standards to promote greater use of AI, address security and protection measures to protect the United States' lead in AI, and other steps.

---

8 January 25, 2019. Available at: https://about.bgov.com/news/federal-market-grew-10-percent-fiscal-2018/
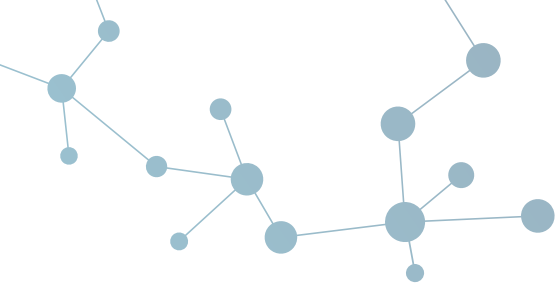
Challenges, Questions, and Considerations for **AI Adoption**

Harnessing AI to take on specific tasks requires a great deal of thought, teamwork, and preparation. For organizations new to AI, as many federal agencies are, this can present significant challenges. At every stage of the adoption process, federal agencies are navigating new learning curves and encountering new sets of questions and considerations.

In many respects, these questions and challenges reveal the contours of a typical AI learning curve for any large enterprise, whether government or commercial. In our conversations with numerous federal practitioners, we found that there were four specific aspects that often presented challenges to, and that must be considered in, agencies' AI adoption: creating the business case; growing AI competency; building an analytics culture; and ensuring ethical use.

CREATING THE BUSINESS CASE

Most agencies do not go out looking to buy an AI solution because of all the great things it can do, nor should they. They start with a problem to be solved and figure out how best to address that problem. As one agency CTO told us: "It's not about the technology—as a matter of fact, I shoo people away when they come and want to talk about their technology because I want to talk about my problem." Given the accelerating rate at which AI technologies are advancing, the best solutions increasingly include AI components.

> The objective should be to get started, find a "quick hit" that offers a clear ROI, get experience in deploying and working with AI, learn lessons and gain insights, and then expand further as experience, funding, and ROI are realized.

But introducing AI capabilities into a business operation often comes with a learning curve and a significant degree of transformation. In terms of initial learnings—understanding how AI might perform in your environment—innovation funding and similar sources can be particularly valuable. This is especially true as many of the tools are either open-source or cloud-based, meaning that even a limited budget can go a long way. The advantage of this approach is that you're not committing your full resources before you understand requirements and risks, which is how projects spiral out of control.

From there, you can apply traditional business case analysis to building the investment model for AI. Practitioners and experts advise starting with—but not stopping at—a relatively straightforward use case that can be easily automated and offers obvious potential efficiencies—in other words, the proverbial low-hanging fruit. Look initially for repetitive tasks that rely upon well-structured, transparent, rules-based decisional processes. Inspections,

case management, records and forms processing, waste, fraud and abuse activities, and call center support are among the government use cases ripe for AI solutions. Over time, tools like machine learning can be used to address more complex or abstract business rules to more fully take advantage of AI capabilities.

The objective should be to get started, find a "quick hit" that offers a clear ROI, get experience in deploying and working with AI, learn lessons and gain insights, and then expand further as experience, funding, and ROI are realized. This effort should occur within the context of a broader AI strategy that helps ensure each project is aligned with the agency's long-term roadmap. The longer-range goal is to develop a widening culture around analytics and intelligent automation so that AI is no longer viewed as an exotic tool or solution, but rather as an integral part of the agency's ongoing focus to improve operational effectiveness and efficiency.

Once a use case is found, the project team—in consultation with other stakeholders—must flesh out the project by defining its scope and success metrics. It must redesign the workflow as it will look within the embedded AI capability. And it must examine thorny questions such as "Will the algorithms in use stand up to legal, ethical, and technical scrutiny?" and "How are decision-making authorities and divisions of labor being delegated between machines and humans?"

GROWING AI COMPETENCY

Getting one or a few successful AI projects under the belt is great, but how does an agency mature as an analytics-driven enterprise? Like many maturity curves, this one takes a crawl-walk-run approach.

In the **crawl** stage, agencies are educating themselves about the capabilities and limits of the technology, exploring how they can apply it to their operational needs, and establishing proper data governance and data management strategies. This is where the agency embraces a more data-centric view of mission execution and the notion that data is a valuable strategic asset that requires proper handling and attention. This stage may include some routine task automation.

The **walk** stage is where agency data management policies and strategies begin taking shape and showing results. Priority data sets are prepared for processing, transitioned from closed to open environments so they can be leveraged, and matched up with suitable technologies and models that can deliver high-value analytics for the enterprise. Data and analytics begin driving decisions in a systemic way. Strong leadership drives the momentum throughout.

The **run** stage is where data management and governance regimes take hold across the enterprise and data emerges as a high-value strategic asset. Analytics is scaled and democratized throughout the enterprise. Decision-makers may even incorporate predictive analytics into their planning. Machine learning enables continuous optimization, helping to refine existing approaches to become even more efficient and effective. AI and intelligent automation are firmly in place and ingrained within the fabric of the agency culture and operation.

Throughout these stages, there will inevitably be challenges in getting agency leaders, managers, and workforces comfortable in working with AI. As one federal research agency IT official told us, "The biggest challenges have been the people challenges. [For example,] how do they trust the model?"

In many cases, agency staffs will want to "look under the hood" of the AI tools they rely upon so they can develop needed levels of trust in their capabilities and decision-making processes. This is especially critical for agencies with scientific, research, technical, law enforcement, or intelligence focused missions. A senior official at a large science and technology-focused agency told us that such scrutiny will be a critical part of federal AI adoption. "How did you get there? Why did it arrive at that? It's not just about delivering the end product—it's about having trust in that answer because you're going to have technical people challenging it."

The challenge of gaining transparency into AI algorithms is sometimes referred to as the "black box" problem. "The black box issue is a big issue in my world," an official at a federal law enforcement agency told us. "You need to understand what's sitting behind these complex calculations —and that understanding has to be built into this."

A related challenge is being able to explain how an algorithm reaches the outcomes it does. As a research agency CIO told us, "One of the problems that we found with deep neural networks is that you can't really explain how they work. They work, but they rely on a mathematical model of high complexity." This challenge of 'explainable AI,' as it is sometimes called, can be hard to address. Users increasingly desire greater understanding and confidence in decision-making, which require more transparent approaches. The CIO we interviewed came up with a work-around: "One of the things that we've been doing increasingly is giving people a confidence indicator, which says, 'The model is 95 percent sure that it's right,' or, for these more 'fluffy' ones, it may say, 'I'm only 50 percent sure'."

Leadership Needed to Change Culture, Accept Risk
Another aspect of the 'people challenge' is getting people comfortable with a greater degree of risk. "I work in a very risk-averse organization and we really have to do a lot of education … and push people," said one leader from a large financial services agency.

When it comes to shifting an organization's culture from being risk-averse to more risk-tolerant, strong leadership can be helpful. Said another agency CTO: "I worked very hard to build a culture that says, 'We're all human, we all make mistakes.' Those mistakes need to get to the top as fast as possible so we can deal with them—not so we can find blame and point fingers, but so we can learn from them, spread that learning across the organization, and become a better organization."

Many of the federal AI practitioners we interviewed agreed that having a champion for AI at a high level is perhaps the most critical ingredient to success. Said one computer scientist at a technical agency, "I have to say that since our director came in, and he has a penchant to use AI, he sees the big picture. We're able to break organizational boundaries that typically held us back. … You need someone [in] control at a much higher level, to put these higher teams together. I see how a champion can truly make the way."

BUILDING AN ANALYTICS CULTURE
Like any innovation, AI holds great potential, but realizing that potential depends on whether it is accepted and used effectively by the agency. Success starts with the agency's leadership. It is

important for leaders to think about their forays into AI as part of a broader goal of building an agency-wide culture around analytics and evidence-based decision-making.

Building that culture starts with engaged leaders who embrace analytics as a way to more effectively and efficiently deliver services, meet mission performance targets, and run the enterprise. Leaders can be effective if they convey the message that AI is not something to be feared, but rather something that will empower the workforce to deliver better results.

A key component in transforming a culture, according to many we interviewed, is to adopt a continuous and open communications approach. Discuss what is happening—the problems, the challenges, the progress, the goals, and the concerns—openly and often. Address any concerns and fears head on.

Another critical piece is reskilling and upskilling. As AI efforts allow people to move from low-value work to high-value work, they will need training and education to ensure they are successful in that transition. Those who are working with AI tools will need educational opportunities to improve their data literacy so they can be more effective in negotiating human-machine interactions and contribute to better outcomes.

ENSURING ETHICAL USE

Most federal missions—from national security, law enforcement and counter-terrorism, to public safety, healthcare, housing, education, and more—are freighted with considerable ethical dimensions. As AI capabilities propagate across these mission spaces to support decision-making, policy-making, and day-to-day program operations, there is a justifiable public concern over how and whether these tools will deliver outcomes that are consistently fair, ethical, accurate, transparent, accountable, and safe.

From a mission perspective, these are not trivial considerations. When an agency uses an AI tool to determine whether and how much of a government benefit to give someone, or to conduct a risk assessment that determines someone's criminal justice options, for example, that agency may be required to defend those algorithms in court.

Another big concern is human bias, which can, and does, creep into data and algorithms, usually unknowingly. On the data side, one can envision an AI concluding that a U.S. president must be male and very likely white. When that happens, outcomes can be skewed at best and harmful at worst. One way to mitigate against bias is to build multidisciplinary and diverse teams, supplemented with ethicists, for all AI efforts. Another tactic is vigilance—continually checking AI outcomes for signs of bias and solving for those biases quickly when found.

Data protection, security, and privacy are similarly major concerns for agencies today. Many agencies may realize they require more sophisticated data disclosure and management regimes to ensure that information is not unintentionally disclosed or used in unintended ways. As agencies appoint Chief Data Officers and Chief Privacy Officers, as required under the Foundations for Evidenced-Based Policymaking Act, these will be important concerns.

AI Transforms,
But It Also Disrupts

AI is inherently disruptive. When machines are inserted into a workflow to make decisions, even small decisions, that were once made by humans, the ripple effects can rattle the workplace, the culture, and even policies. So it is important that leaders of an AI implementation understand that disruptions—both anticipated and unanticipated—will occur as part of the implementation and be ready for them.

Anticipated disruptions refer, for example, to the changes that can occur when routine tasks are automated. Workflows, processes, roles and responsibilities, and needed skill sets may change or be replaced entirely.

Unanticipated disruptions, by their nature, are less well understood. But one frequent category of unanticipated disruptions involves cases where an AI reveals flaws in a process that require attention. This could involve, for example, a revelation that biases are present in a decision algorithm or in longstanding operating guidelines that unfairly skew outcomes, introduce inequities, or reduce effectiveness or efficiencies. When that happens, there is a need to review operating procedures and guidelines to address those concerns.

Another category of unanticipated disruption occurs when an AI tool is addressing situations for which it was not adequately "trained" or prepared. Operating scenarios can and do change. But AI tools do not naturally adjust along with them — they must be trained to adapt to ever-growing sets of variables. If not, they may produce adverse results or failures.

A variety of tactics are helpful in managing some of these disruptive impacts. One is to adopt an open communications strategy with affected workforces. Communicate often and candidly with them about what is going to happen, how the AI-enabled processes will work, what the benefits and expectations are, how anticipated disruptions will be managed, and, especially, what all of this means for those directly affected.

A second tactic is to embrace co-creation. Engage the workforce to be a part of the modernization process so their views and insights are considered and so that front-line employees are not only included, but invested, in how it takes shape.

The good news is that there is important work already being done in the ethics arena within the Executive Branch. The Defense Department has tasked the Defense Innovation Board to develop a framework and set of principles for the ethical use of AI and the board is soliciting public input as it develops those principles. The result of that effort may serve as a model for other agencies.

While the Defense Department is a leader within government when it comes to developing ethical principles for AI, other agencies will likewise need to give careful thought to the ethical dimensions of AI within the specific contexts of their missions and operations. Potential principles for an ethical AI program may include:

→ **Avoiding unfair bias** through vigilant oversight and awareness, as well as active measures to ensure equity in AI outcomes.

→ **Data privacy**, perhaps to include individual access controls so people have the option of protecting their own data.

→ **Accountability for AI creators and operators.** This could include legal frameworks for accountability and requirements for transparency and "explainability," which refers to the need to be able to explain how AI tools produce the results they do in terms that all stakeholders can understand.

→ **Communications requirements** to ensure people are informed about how their data is being used.

→ **Educational components** to help minimize the risks of misuse.

Congress is also engaged in the "ethical use of AI." Legislation has been introduced in 2019 in both the House and Senate to address guidance for ethical use of AI, and there are hearings and legislation relating to specific applications of AI, such as facial recognition.
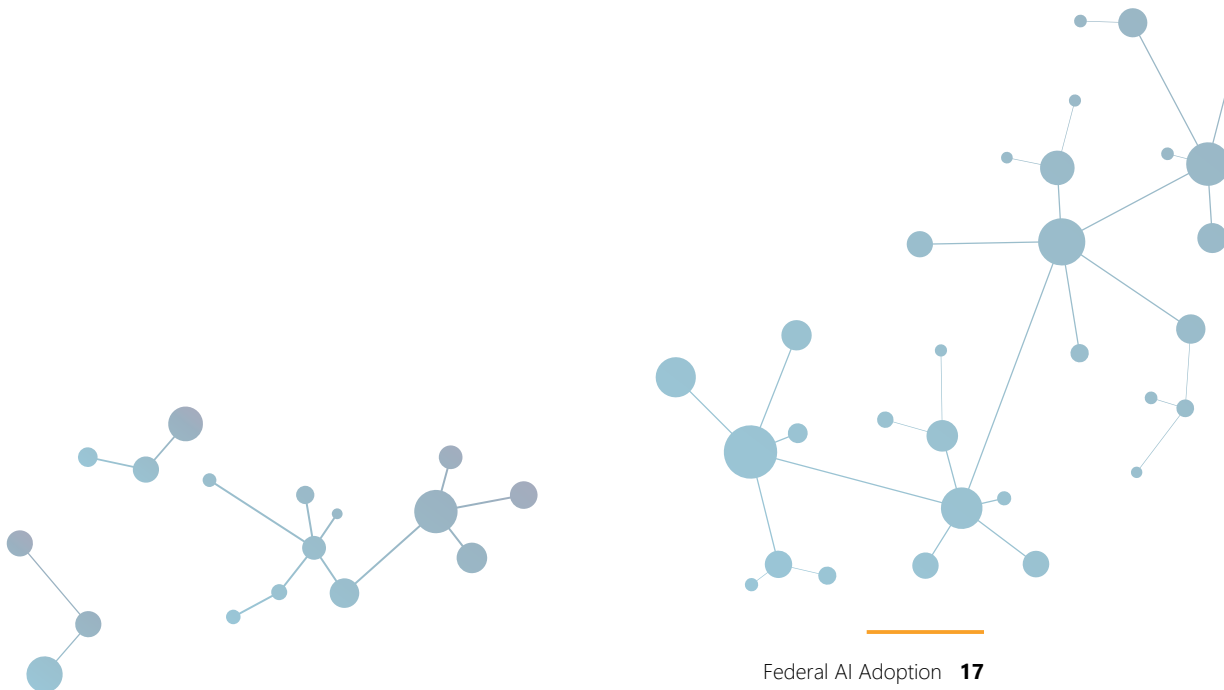
Conclusion

Today, we are beyond the point of asking whether AI and intelligent automation will catch on in government. The economic and mission incentives, use cases, policy directives, and advancing commercial capabilities are fast emerging to propel it forward.

The question now is how much progress we should expect to see. Each agency has its own culture, leadership style, mission portfolio, and operational DNA—and, consequently, each will necessarily chart its own path toward becoming a more data-centric enterprise and AI adopter.

Regardless of how they approach it, agencies will be well-served to be vigilant with their teams in key areas. They will need to stand up modern data management regimes that not only protect data but also recognize its potential value to the enterprise. They will need to bring active and effective leadership to bear. And they will need to navigate the shoals of transformation and disruption. But, as we have outlined here, there are workable strategies and tactics to overcome these challenges for those willing to make the effort.

It is increasingly clear that status quo approaches to service delivery and mission execution will not carry forward. AI and intelligent automation can no longer be viewed as optional for today's federal enterprises. Indeed, they represent an enormous opportunity for agencies to dramatically improve mission outcomes even as the growth in complexity of their mission challenges outpaces the growth of their resources.

The PSC Foundation intends to follow the federal agencies' journeys toward AI adoption, and to monitor the policies and lessons learned along the way. We welcome broad engagement with us in the new phase of our work.

Acknowledgments

Underwriter's Note:
Accenture

Accenture Federal Services is proud to partner with the Professional Services Council Foundation in publishing *Federal AI Adoption: Where It Stands, Where It Is Heading, and the Opportunities and Challenges Ahead*.

This report is timely, as federal leaders are ready to take Artificial Intelligence (AI) from the labs to the field. According to Accenture research, just 28 percent of these executives report having implemented AI within their agency at the beginning of the year. However, 53% say their agencies will be investing in the technology in 2019.

As detailed in this report, federal leaders have a strong sense for AI's overall potential. And there is genuine excitement about the variety of use cases for federal agencies that are beginning to emerge. At the same time, there is also real pragmatism around some aspects of adoption, given the need to maintain trust with citizens and other stakeholders.

With something as new as AI, open dialogue and a collaborative approach can be invaluable for envisioning how to apply the technology within an organization. What we have learned is that AI has broad applicability throughout government with real solutions ready to make a real impact today.

Most federal agencies have amassed a wealth of data. AI can help unlock this value, transforming their operations to become more intelligent, more intuitive and more efficient. By doing so, they will position themselves to meet the great challenges of the 21st Century.

Whether you are just getting started or already underway, I encourage you to lead this journey.

Dominic Delmolino
Managing Director & Chief Technology Officer
Accenture Federal Services

> accenture

PSC
PROFESSIONAL
SERVICES
COUNCIL
FOUNDATION

www.professionalservicescouncilfoundation.org