ARTIFICIAL INTELLIGENCE / MACHINE LEARNING PRIMER

Enabling AI/ML Innovation in the US Federal Government

EMERGING TECHNOLOGY COMMUNITY OF INTEREST
Artificial Intelligence Working Group

Date Released: March 12, 2019

Synopsis

This primer serves as a foundational tool in the understanding of artificial intelligence and machine learning (AI/ML) technology and, through the use cases, reveal how they can be applied to benefit the federal government and improve the lives of American citizens.

AI/ML has the potential to help government agencies operate more effectively, save money, improve their information security, and improve the lives of American citizens. The potential derives from its ability to increase the value of the government's information assets (by enabling more data to be analyzed more quickly and more robustly) and workforce (by providing government knowledge workers with enhanced analytic tools and assistants).

The technology can advance insights from petabytes worth of structured and unstructured data by significantly accelerating the assessment, clarification, and synthesis of government data. The adoption of AI services may also allow government agencies to better improve their prediction of new mission trends and stakeholder needs to effectively deliver services to citizens and achieve agency objectives.

This page is intentionally blank.

American Council for Technology-Industry Advisory Council (ACT-IAC)

The American Council for Technology (ACT) is a non-profit educational organization established to create a more effective and innovative government. ACT-IAC provides a unique, objective, and trusted forum where government and industry executives are working together to improve public services and agency operations through the use of technology. ACT-IAC contributes to better communication between government and industry, collaborative and innovative problem solving, and a more professional and qualified workforce.

The information, conclusions, and recommendations contained in this publication were produced by volunteers from government and industry who share the ACT-IAC vision of a more effective and innovative government. ACT-IAC volunteers represent a wide diversity of organizations (public and private) and functions. These volunteers use the ACT-IAC collaborative process, refined over thirty years of experience, to produce outcomes that are consensus-based. The findings and recommendations contained in this report are based on consensus and do not represent the views of any particular individual or organization.

To maintain the objectivity and integrity of its collaborative process, ACT-IAC does not accept government funding.

ACT-IAC welcomes the participation of all public and private organizations committed to improving the delivery of public services through the effective and efficient use of IT. For additional information, visit the ACT-IAC website at www.actiac.org.

Emerging Technology of Community of Interest

ACT-IAC, through the Emerging Technology Community of Interest, formed an Artificial Intelligence Working Group to give voice to and provide an authoritative resource for government agencies looking to understand and incorporate AI/ML technology and functionality into their organizations. This working group includes government and industry thought leaders incubating government use cases. The ACT-IAC Emerging Technology Community of Interest (ET COI) mission is to provide an energetic, collaborative consortium comprised of leading practitioners in data science, technology, and research, engaged with industry, academia, and public officials and executives focused on emerging and leading technologies which transform public sector capabilities.

Disclaimer

This document has been prepared to contribute to a more effective, efficient, and innovative government. The information contained in this report is the result of a collaborative process in which a number of individuals participated. This document does not – nor is it intended to – endorse or recommend any specific technology, product, or vendor. Moreover, the views expressed in this document do not necessarily represent the official views of the individuals and organizations that participated in its development. Every effort has been made to present accurate and reliable information in this report. However, ACT-IAC assumes no responsibility for consequences resulting from the use of the information herein.

Further Information
For further information, contact the American Council for Technology-Industry Advisory Council at (703) 208-4800 or www.actiac.org.

Table of Contents

Table of Figures

This page is blank.

Introduction and Executive Summary

"Artificial Intelligence" is the demonstration of intelligent behavior by human-made systems.

This primer is an introduction to the fundamentals of artificial intelligence (AI)/machine learning (ML) technology and how they can be applied to benefit the federal government and improve the lives of American citizens.

Just in the last few years, there have been dramatic advances in the use of artificial intelligence technologies in various areas ranging from robotics, machine vision, and speech recognition to future prediction. Trends in investment, commercial activity, and technical performance have increased. The improvements in technical performance of AI systems are being driven by multiple, reinforcing trends: the explosion of available data, Moore's law continuing to drive computational speed, network speeds increasing while latency is decreasing.

Artificial intelligence and machine learning has already revolutionized how we live, work, learn, and communicate. The significant impact on our economy, education, quality of life, and national security is continuing to grow.

According to Stanford's Human Centered AI Institute's AI Index 2017 Report [1], artificial intelligence tools have reached or exceeded human-level performance at narrowly defined tasks such as strategy games, visual image detection, and parsing natural language. The performance of current tools, however, may degrade dramatically if the original task is modified even slightly.

The federal government has recognized the importance of AI technology and has prioritized funding for fundamental research and computing infrastructure, machine learning, and autonomous systems. President Trump's FY2019 Budget Request was the first in history to designate artificial intelligence and autonomous and unmanned systems as Administration research and development priorities[2].

This primer gives a preview to artificial intelligence, provide fundamental understanding of the technology, its benefits to the federal government, present use cases, and the impact of this technology on the future of the government.

Primer Scope

This primer provides government and private sector audiences a basic overview of artificial intelligence/machine learning (AI/ML) technology and a set of definitions of basic terminology. The primer also provides examples of different needs AI can assist to solve, as well as example use cases for government organizations describing where AI/ML can help organizations automate operations and find new insights from their data.

This primer is not meant to be a technical document. This paper does not provide a deep dive into the mathematical principles behind AI algorithms, or the computational techniques and languages commonly used by data scientists.

This primer does not fully address the area of Intelligent Automation nor Robotic Process Automation. For further understanding of Intelligent Automation, please refer to the ACT-IAC Intelligent Automation Primer[3] assembled by the Emerging Technology Community of Interest Intelligent Automation Working Group.

Definitions

Artificial intelligence and machine learning technologies are not new. Having a common baseline for terminology used in the AI/ML space helps level set technologists new to this body of knowledge.

Derived from available literature, these are the most commonly accepted definitions of the terms of the terms related to AI technology.

Artificial Intelligence (AI): The field of computer science that deals with the simulation of intelligent behavior in computers or the capability of a machine to imitate intelligent human behavior. It is used broadly to refer to any algorithms, methods, or technologies that make a system act and/or behave like a human and includes machine learning, computer vision, natural language processing, cognitive, robotics, and related topics.

Robotic Process Automation (RPA): A technology platform that enables a software robot to interact with applications. RPA software can be easily programmed to do basic tasks across applications just as human workers do. The software robot can be taught a workflow with multiple steps and applications such as taking received forms, sending a receipt message, and checking the form for completeness. RPA software is designed to reduce the burden of repetitive, simple tasks on employees. RPA technology can automate many common data entry activities that can include invoice entry, human resources personnel action entry, and data verification.

Intelligent Automation (IA): Using AI techniques to improve business process automation by making cognitive decisions and taking actions. IA is an evolution of RPA. Where RPA automates processes that are relatively simple and static, IA discovers and applies intelligent business rules to prioritize, approve, route, and evaluate business process tasks and workflows. It is essentially a software that mimics the behavior of an end user by using existing enterprise application screens or web pages to find, evaluate, cut, calculate, transform, and enter data into existing enterprise application fields according to business rules.

Cognitive Computing: Technology that simulates the human thought process (how the human brain/mind senses, reasons, and responds to stimulus) and assists the human decision-making process. Cognitive systems understand, reason, and self-learn and are often based on data mining, pattern recognition, and natural language processing in order to mimic the way the human brain works. They are adaptive, interactive, and contextual.

Machine Learning (ML): A technique within the field of artificial intelligence which uses algorithms that automatically improve with experience or "learn" without being explicitly programmed. Machine learning is the application of AI techniques using statistical methods that enable machines to improve correlations as more data is used in the model, and for models to change over time as new data is correlated. Machine learning evolved from the connectionist theory of human cognition, pattern recognition, and computational learning theory. Approaches include neural networks and deep learning.

Neural Networks/Deep Learning: A class of machine learning algorithms that use a cascade of multiple layers of nonlinear processing units for feature extraction and transformation. Figure 1 represents how raw data is processed and categorized into outputs. Each successive layer uses the output from the previous layer as input. Deep Learning can use supervised (e.g. classification) and/or unsupervised (e.g. pattern analysis) methods. When designed appropriately, these systems can "learn" multiple levels of representations that correspond to different levels of abstraction; which may form a hierarchy of concepts familiar to human users.

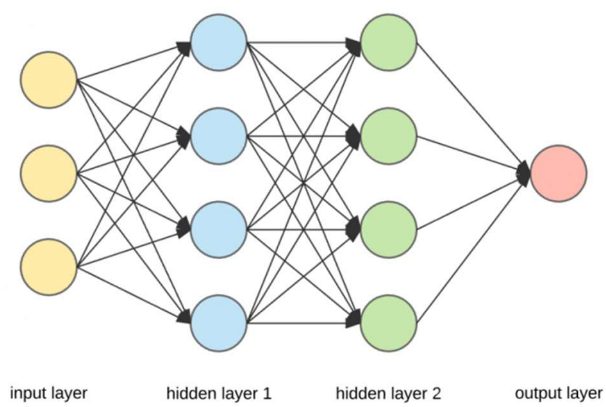

Figure 1: A neural network (in this case, with two hidden layers) identifies hidden relationships between data to create final outcomes

Deterministic Model: A deterministic model assumes certainty in all aspects. The output of the model is fully determined by the parameter values and the initial conditions. Deterministic models are based on known rules so that given input values will always produce the same result. Examples include timetables, pricing structures, linear programming model, the economic order quantity model, maps, accounting, etc.

Probabilistic (Stochastic) Model: Stochastic models possess some inherent randomness. The same set of parameter values and initial conditions will lead to an ensemble of different outputs. Stochastic models accept that there is a probability distribution associated with the inputs so that the same input can yield different output values. Because stochastic models utilize probability density functions in one form or another, they need to be well based in statistical theory. By contrast, deterministic models do not depend on statistics. Examples include class grades on an exam (normal distribution) and expected weather forecast for a day.

Supervised Machine Learning: In supervised machine learning, a model is given data that is labeled in an organized fashion by a subject matter expert. In figure 2, which shows an example of supervised learning, one might provide pictures of cats with the label "cat." Once enough structured and labeled data is provided, the AI model built can recognize and respond to patterns in data without explicit instructions. The output and the accuracy of supervised learning algorithms are easy to measure. Supervised learning is a common method of machine learning today.

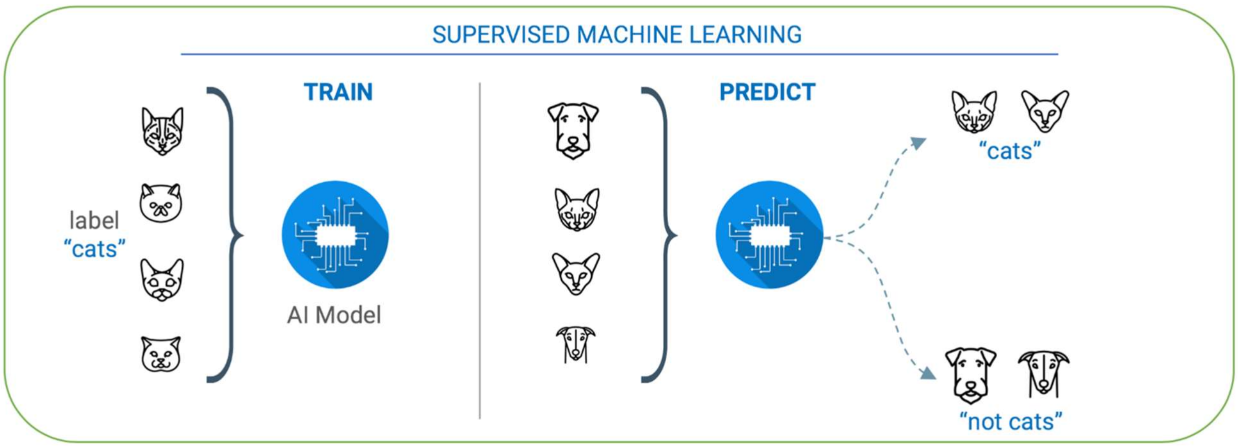

Figure 2: Supervised learning using subject matter experts to "teach" the correct choices[4]

Unsupervised Machine Learning: A model is given data that are not labeled in an organized fashion. For example, one might provide pictures of animals (cats, dogs, etc.) without any labels. The model organizes items into logical groupings based on patterns recognized by the model. This method is used to identify underlying patterns from previously unstructured data. The expectation is not to derive the right output but to explore datasets and draw inferences. Unsupervised learning is often applied to cases where the researcher or user does not already know the significant structures or relationships among the data. Figure 3 depicts this approach to learning.

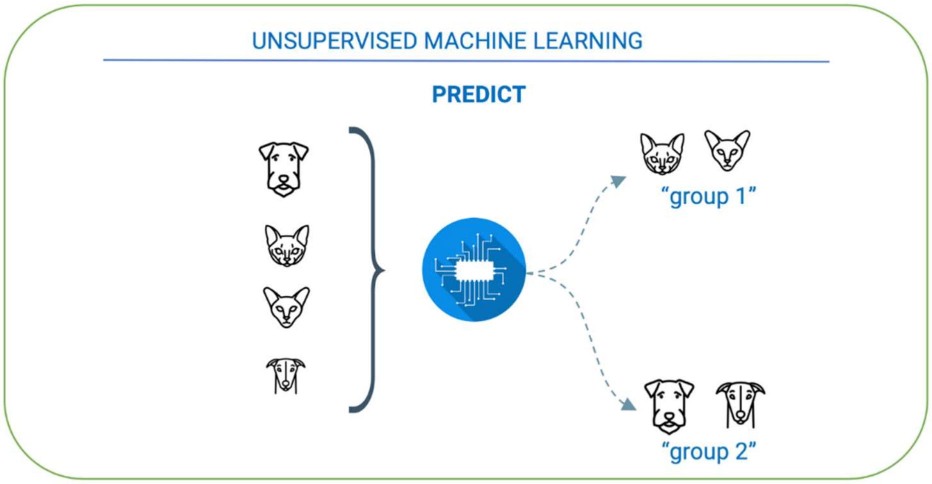

Figure 3 Unsupervised learning enables the machine to identify data with like associations[5]

The following terms describe the common types of algorithms used in AI/ML:

Bayesian Networks: Bayesian models apply Bayes' theorem[6] of conditional probability to provide predictions. Generally designed using probabilistic graphing techniques, these models represent a set of variables and their conditional dependencies. For example, a Bayesian network could represent the probabilistic relationships between diseases and symptoms. Given symptoms, the network can be used to compute the probabilities of the presence of various diseases. Bayesian networks that model sequences of variables (e.g. speech signals or protein sequences) are called dynamic Bayesian networks.

Clustering: Cluster analysis or clustering is the task of grouping a set of objects in such a way that objects in the same group (called a cluster) are more similar (in some sense) to each other than to those in other groups.

Decision Tree Learning: Uses a decision tree to go from observations about an item (represented in the branches) to conclusions about the item's target value (represented in the leaves). Models where the target variable can take a discrete set of values are called classification trees. In these structures, leaves represent class labels (predictions, dependent variables) and branches represent conjunctions of features (known information, independent variables) that lead to those class labels

Reinforcement Learning[7]: Reinforcement learning systems "learn" what output to generate by maximizing a numerical reward signal. The system is not given decision rules but instead must discover which outputs (from a predetermined set) yield the most reward by trying a variety of them. In the most interesting and challenging cases, outputs may affect not only the immediate reward signal but also the next input set, and through that, all subsequent reward signals. These two characteristics—trial-and-error search and delayed reward—are the two most important distinguishing features of reinforcement learning systems.

Representation Learning: Representation learning aims at discovering better representations of inputs by learning transformations of data that disentangle factors of variation in data while retaining most of the information. Features are extracted from unlabeled data by training a neural network on a secondary, supervised learning task.

Artificial Intelligence/Machine Learning Technologies

Commercial applications of AI/ML technologies include the following well known commercial services.[8] Many common online services use advanced machine learning "recommender" software with extensive behavioral algorithms that predict our likes and dislikes based on the recorded preferences of persons with similar demographic backgrounds.

Platforms such as Facebook and Amazon combine AI technologies to deliver customer experiences designed to benefit the platforms. For example, Amazon's Alexa platform combines a consumer recommendation engine, natural language processing, audio sensor technology, and advanced speaker design to provide consumers with an attentive, informed, personal assistant. These "recommenders" improve their recommendations based on analysis of ever larger datasets.

Such systems use algorithms tuned to specific types of data as identified by many of the following popular services powered by artificial intelligence being used today:

Natural Language Processing: *Apple's Siri, Amazon's Alexa, Google Home, Microsoft's Cortana*
Many are familiar with voice recognition personal assistants common on smartphones. Through speech recognition and semantic analysis technology, these systems understand the context of many common spoken words and derive results tailored to a user's past behavior. These assistants help us find information, give directions, add events to our calendars, help us send messages, and so on. These systems use machine learning technology to get smarter and better able to predict and understand our natural language questions and requests.

Sensors and Autonomous Mobility: *Tesla, Uber*
Self-driving cars are now a reality. By combining the multiple environmental sensors (forward-looking and side–looking lidar/radar and multiple real-time video image analysis), GPS location-based mapping, advanced servo and accelerometer controls, and other information, vehicle manufacturers and transportation service providers provide autonomously navigating vehicles based on deterministic and probabilistic predictive algorithms which provide the self-driving features.

Consumer Choice Recommenders: *Netflix, Amazon, Spotify, Pandora*
These systems collect the choices (likes and dislikes of millions of users) and provide customers recommendations for books, music, videos, products, and movies based on the choices of micro-segmented "like" users. The more a customer uses these services, the more likely a recommended

product will meet the needs of a user. These deterministic algorithms are refined more and more each year as these firms acquire more data about the choices and demographics of their users. These recommendation engines analyze billions of records to suggest choices that you might like based on your previous reactions and those of other customers closely matched to your preferences and your demographics profile.

Behavioral Analysis: *Cognito, Care, Predictum*
Less well known than consumer choice recommenders, these systems seek to predict future behaviors of persons based on AI algorithms that correlate known behaviors to relevant data collected via publicly available sources. Using AI algorithms trained across many millions of data points, these services seek to accurately predict important personality traits of individuals, e.g likelihood of the individual to exhibit select behaviors, preferences, or judgments.) These services fuse machine learning identified patterns with behavioral science to identify future behaviors or influence future behaviors.

Capabilities

Some of the useful capabilities of AI/ML technologies, especially Natural Language Processing, are:

Sentiment Analysis, Tone Analysis, and Personality Insights: These capabilities are forms of natural language classifying aimed at analyzing unstructured text in the form of social media, emails, customer reviews, and more, to gain insight into feelings. For example, this could help a company understand if customers are responding positively or negatively toward their new product based on a vast amount of social media posts. It could also be used to personalize product recommendations and target advertising, among many more uses.

Speech to Text and Text to Speech: These capabilities convert mediums of text and speech so that both types of data can provide actionable insight. For example, converting speech to text allows one to analyze voice messages to a call center for tone and sentiment by converting it into text in real-time. It can also be used to allow one to speak to a virtual assistant rather than type. And in reverse, Text to Speech can also be used with virtual assistants to provide more accessibility for users. A good example of both together is Apple's Siri, which "listens" to the user speak, converts what it heard into text for display, and then responds with both written text and verbally.

Keyword Extraction, Entity Extraction, and Relationship Extraction: These capabilities can work together to improve search engines and research by identifying terms that contribute to the main point of a document, classifying key elements from text into pre-defined categories, and detecting semantic relationships between entities[9]. For example, this could be used to analyze a large group of accident reports for a particular manufacturer to determine if there is a common cause for the accidents reviewed. Using the unstructured data of accident reports, the model could determine the main point of the report, extract entities such as make, model, part number/batch number, and accident conditions, and then determine if there are any correlations among these factors that might indicate a root cause.

Technologies that Support and Benefit from AI/ML: In many cases, physical robotics and AI/ML technologies can work together. Some examples are listed below.

- Internet of Things (IoT): IoT refers to the combination of hardware devices (such as cellphones, home video surveillance systems, automotive GPS systems) and networks (including cellular data, WAN and LAN, WiFi spots, and process automation and control networks) that together, comprise a vast, decentralized "system of systems." The amount of data ingested and the ability to "act on" the physical world represented by this system far surpasses any individual system ever designed. Industry, government, academia, and the military, among others, are all developing tools, techniques, and technologies that will provide society the potential benefits of being able to tap into these tremendous resources. AI/ML technologies can be applied to this challenge by providing insights to increase efficiency, better leverage its capabilities, and improve our overall ability to manipulate the physical world.

- Blockchain: Representing a fundamental shift in how separate, disparate systems can be made to interact with each other, blockchain technology allows systems, individuals, and companies to share data, collaborate, and execute critical transactions without the need to architect and implement a custom system. Refer to the ACT-IAC Blockchain Primer[10] and the ACT-IAC Blockchain Playbook[11] for more technical information regarding this technology. It should be pointed out that blockchain may benefit AI systems by allowing them to collaborate and transact exchanges at a distance. AI may also benefit blockchain implementations by identifying potential use cases and protecting against cyber threats.

- Sensors: Using sensors that either already exist in a device or that can be added on can tap into a wide range of useful information. For example, voltage sensors in generator batteries could be used to predict device failure using AI/ML technologies, and thus improve device maintenance and reduce replacement costs.

- Specialized computer chips: Graphical Processing Units (GPUs), while originally designed for gaming and other graphics-intensive processes, have proven to be a huge boost to neural networks. GPUs save both time and money compared to a traditional Central Processing Unit (CPU) when training neural networks. For example, one cloud provider tested a single GPU and found it to perform as well as 100 CPUs[12]. Although a GPU costs slightly more than a CPU on an hourly basis, the hundredfold increase in speed results in about a 98% cost reduction.

- TensorFlow[13]: An open source software library for high performance numerical computation, TensorFlow is a machine learning system that operates at large scale and in heterogeneous environments. The flexible architecture allows deployment of computation across a variety of platforms (CPUs, GPUs, TPUs) and from desktops to clusters of servers to mobile and edge devices. Originally developed by researchers and engineers from the Google Brain team within

Google's AI organization, it comes with strong support for machine learning and deep learning and the flexible numerical computation core is used across many other scientific domains.

TensorFlow uses dataflow graphs to represent computation, shared state, and the operations that mutate that state. It maps the nodes of a dataflow graph across many machines in a cluster, and within a machine across multiple computational devices, including multicore CPUs, general purpose GPUs, and custom application-specific integrated circuits (ASICs) known as Tensor Processing Units (TPUs). This architecture gives flexibility to the application developer. Whereas in previous "parameter server" designs, the management of shared state is built into the system, TensorFlow enables developers to experiment with novel optimizations and training algorithms. TensorFlow supports a variety of applications, with a focus on training and inference on deep neural networks.

- Python: Python is the most popular programming language for data science. Python is on open source language available from www.python.org. Data science requires statistical/mathematical operations to be conducted across large multidimensional arrays of data. The open source community supporting Python has developed libraries that enable these computation to be made. As an example[14] libraries such a "scikit-learn" provide developers and data scientists with accessible and open source machine learning libraries.

Limitations and Risks

While the AI/ML capabilities have the potential to completely revolutionize the way humans live, work, learn, and communicate, there are still many limitations and risks that must be taken into consideration in order to effectively utilize these technologies.

A fundamental misunderstanding is that AI/ML technologies train themselves. The reality is that in order for a system to do something as simple as distinguish a cat from a dog, it must undergo supervised deep learning where its neural networks are trained to distinguish one from the other[15].

This supervised training actually requires a great deal of human labor. In the above example, humans would need to go through a large number of images of cats and dogs and label each of them so that the algorithms can better understand them and make predictions.

AI/ML systems also rely on the availability of extensive and accurate data. These models are only as effective as the data that is fed into them. It can take an organization a long time to gather sufficient data for the AI/ML system to be able to draw useful conclusions. Another limitation of AI is lack of transfer learning - the ability to apply insights from one problem to solve a different problem[16].

Rapid development of AI/ML systems lead to improved capabilities but can also increase the cost of implementation. AI/ML systems will frequently be integrated with existing systems to gather data. For an AI/ML system to be effective, it must be supported by the existing technology with which it is being integrated. Ensuring usability and interoperability between all systems and platforms within an organization can require system upgrades and adaptations, which may bear significant costs.

Ethical Considerations and Implications

As an organization considers the use of AI/ML within its administrative and mission operations, government agencies and the private sector should address the ethical use of this technology. In its purest form, AI depends upon computational processes to find hidden positive and/or negative correlations amongst elements within a very large data. In practical application, AI/ML enhanced systems are pointed to add insights to very specific hypothesis around targeted mission areas. To create and foster trust in the use of new technologies, practitioners must understand the ethical resources and standards available for reference during the design, build, and maintenance of AI. The focus on AI ethics by groups like the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems[17], should be mirrored in businesses and working groups of all sizes. Ethics and regulatory issues can differ by mission area as described the following examples:

Medical Safety[18]

This primer is timely as AI systems are augmenting and in some cases replacing human judgement in actual clinics. AI systems could directly harm patients with unsafe systems leading to overspending, injury, or even death. AI systems are able to review and identify "hidden" issues within medical imaging (e.g. Sonographic, X-ray, Computed Tomography [CT], and Magnetic Resolution Imaging [MRI] images). AI systems are beginning to be used to identify unintended issues from the use of pharmaceuticals. Data science practitioners need to recognize that AI applied to medicine is different than AI in almost all other areas. In consumer services, performance is valued and results can be tuned and refined over time. In contrast, the Hippocratic Oath begins with "first, do no harm" and tort law recognizes such "missteps" as malpractice. Figure 4 shows a spectrum of increasing complexity in "cognitive decisions."

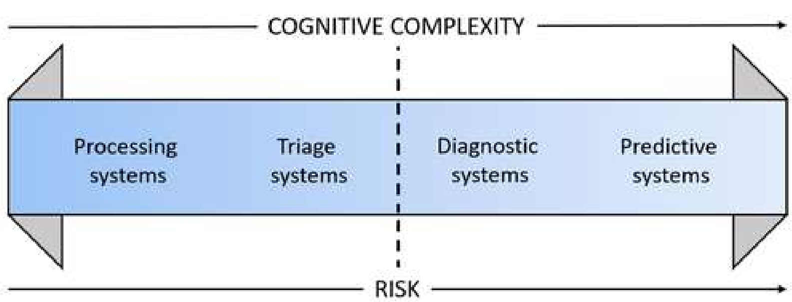

Figure 4: As AI systems move from identifying patterns to diagnosing events and predicting future events the possibly for untraceable error exists[19]

The dotted line is a tipping point where we transition from systems that supply information to human experts and into systems that can make medical decisions independently. Without human review and medical expertise relying on AI/ML driven decision has the potential to cause harm or disaster.

Biometrics, Biographics and Identity[20]
Biometric facial recognition algorithms have struggled with both racial and gender biases, exhibiting higher error rates for both women and non-white subjects. While some products have managed to achieve equitable error rates across the population, many algorithms still struggle with the issue. A study from the Massachusetts Institute of Technology (MIT) earlier this year found significant racial discrepancies in algorithms offered by IBM, Microsoft, and China's Megvii. In a law enforcement context, those error rates would have a serious human cost. Higher false-positives could lead to more police stops and more arrests. "There's a real concern that it could exacerbate the risk of police use of force," Laura Moy of Georgetown Law's Center for Privacy and Technology told The Washington Post. "In a real-time scenario where a police officer is likely armed, the risks associated with potential misidentification are always going to exceed any possible benefits."[21]

As currently envisioned, biometric identity verification is limited to certain areas or needs. Some groups fear that this technology may be expanded to privacy-invasive applications and may be made interoperable with government and law enforcement systems at the state, local, or federal level. Groups such as the American Civil Liberties Union (ACLU) and Georgetown Center for Privacy and Technology believe the effects of such expanded uses may impact free speech and free association rights.

Bias in Data Ingestion/Bias in Algorithm Tuning
Many AI systems[22] rely on machine learning algorithms that are trained with labeled data – the basis of Supervised Learning defined earlier in the document. The opinions, unintended or intended biases of data scientist "trainers" impact the outputs of such trained systems. It has recently been shown that algorithms trained with biased data have resulted in "algorithmic discrimination." MIT research showed that the popular word embedding space, Word2Vec, encodes societal gender biases. The authors used Word2Vec to train an analogy generator that fills in missing words in analogies. The analogy "man is to computer programmer as woman is to [BLANK]" was completed with "homemaker," conforming to the stereotype that programming is associated with men and homemaking with women. The biases in Word2Vec are thus likely to be propagated throughout any system that uses this underlying algorithm.

Replacement of Current Processes, Capabilities, and Workforce
Government agencies are creating far more data now than ever before making it impossible to rely solely on human decision making. With the amount of data increasing, it is important that no data are left unused. Unused data that is invisible to employees may contain valuable information and AI has the capability to capture that value. It is no longer possible for humans to navigate through the data available and cognitive systems will help unlock the data that cannot be described or found by the human eye[23].

With AI becoming integral for unlocking, analyzing, and discovering new data, it is equally important that there are employees capable to fill the jobs that AI creates. Although government employees may fear being displaced by AI, it is more likely that the workforce will need employees with advanced technical skills to take advantage of the new tools. Agencies that invest and prioritize training to improve skills such as critical thinking and data analysis, will help their current and new employees interact with AI systems and ultimately allow the agencies to more effectively and efficiently meet mission requirements[24].

Technical professionals find assisting federal programs that have a large impact on the nation more valuable than assisting social media applications. Government agencies can use this to their advantage and grow a culture of innovative employees within their agencies[25].

Impacts on the Business of Government

AI/ML has the potential to touch literally all aspects of our life – from how we interact with our homes, to work life, to most aspects of government provided services including social services, law enforcement, and national defense. The mission, performance, operation, efficiency, and workforce of every federal agency are likely to be impacted by current and future AI advances.

Mission: The mission of many federal agencies will be impacted by AI as these new tools evolve toward increasingly autonomous behavior and become actors on their own part. These systems are likely to impact the scope and nature of many federal agencies' mission. For instance, the question of whether to grant rights or limited citizenship to AI systems (with legal arguments on both sides of the debate already under way) may redefine the concept of "citizen" with resulting impact on federal agencies' services, role, and scope. Setting the question of AI citizenship aside, as the US population and industry adopt these evolving AI systems, the mission environment of every federal agency will be affected. The military has created entirely new, multi-billion dollar organizations to address these changes, and the civil and legislative spheres will face similar challenges.

Performance: AI systems will improve the quality of information used to support decision making at every level of federal agencies. Mission performance will improve as decision making improves. While agency decision making and mission performance improve in the majority of agencies, some agencies may find their mission performance simultaneously being challenged by an expanding set of demands and threats. Autonomous vehicles, for example, may require some federal standardization for overall safety criteria. The ability of AIs to create art, if applied at the scale of cloud computing, could create an unmanageable flood of copyrightable material. Similar examples of AIs potentially expanding the volume of federal agencies' current mission apply to nearly every civil agency. Finally, it must be assumed that AI will continue to be weaponized by both state and non-state actors, which represents an increasing risk for all federal agencies.

Operations: Through process automation (initially RPA leading to IA), cognitive assistants, real time semantic analysis of all incoming data streams, predictive modeling, and integration with IoT, federal agencies' operations will become faster, better informed, more responsive, and more adaptable.

Efficiency: In addition to applying AI/ML systems to supporting mission performance, they will also be applied to drive operational efficiency. With accurate, timely, and comprehensive operational cost analysis, real-time operational simulation and optimization, and the dramatic increases in IT flexibility enabled by combining AI with cloud computing, federal agencies will be able to spend exactly what is required, not a penny more, to achieve desired operational outcomes.

Workforce: As AI provides the federal workforce with intelligent assistants, real time analysis, and accurate predictions, the workforce will find itself spending more time on policy making and goal setting. Many federal workers will have the opportunity to develop their skillsets in these areas, as well as their human-centric skills – those requiring empathy, creativity, and emotional intelligence that are likely to remain beyond the capability of AI tools. The federal workforce and management are likely to follow the pattern that seems to be playing out in the private sector, where overall staff levels are not being reduced, but staff are being retrained and redeployed into new areas of productivity.

Use of AI/ML is supported by President Trump's Presidential Management Agenda and OMB Memorandum M-18-23 Shifting from Low-Value to High-Value Work[26]. Almost every function of government from federal to state and local can benefit from the deployment of AI.

Government Business Use Cases and Outcomes

According to a recent survey by FedScoop and IBM[27], 37% of federal government IT executives are deriving benefits from artificial intelligence, and another 25% expect to do so in the next 5 years. A common use case is to manage the increasing amount of data, but other uses include video/ imagery analysis, automated decision-making, cybersecurity, and improved customer services. Agencies throughout the government are leveraging existing commercial technology and also building in-house solutions.

Federal organizations ranging from the military to law enforcement to pure civilian departments are incorporating AI/ML in order to best use the overflowing amount of data. In October 2018, the Department of Homeland Security (DHS) formed a Community of Interest (CoI) on Artificial Intelligence and Machine Learning.[28] DHS' CoI is focused on analyzing the Department's "vast datasets" for improved prioritization and decision-making. Both the Census Bureau, under the Commerce Department, and the US Army are leveraging machine learning in their operations to protect against inefficiencies in the former case and enemy actions in the latter.[29]

Realizing These Benefits

While the promise of AI is vast, the challenge of realizing these benefits is not insignificant. Any organization electing to implement an AI solution will first have to decide if AI will offer the expected benefits.

Some specific examples include[30]:

- Citizen help services: Using natural language processing to provide the best possible answer to citizen query.
- Military and veteran health: Use of telemedicine to deliver medical benefits to a soldier or veteran whenever and wherever needed.
- Prevention and recovery of disaster: AI can take necessary steps to learn and prevent any disaster in the IT systems, reducing the operational risk of the systems.

In order to begin to implement systems with AI elements, organizations should recognize that IT and mission organizations are just beginning to implement and understand how to use AI/ML systems.

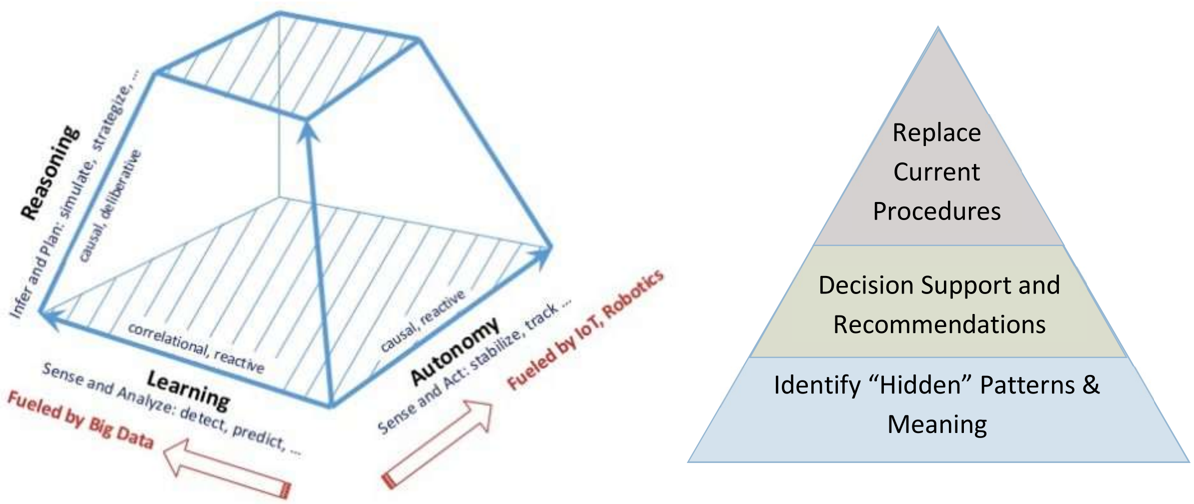

Artificial Intelligence is a capability that is still in its operational infancy

Figure 5: Artificial Intelligence includes a range of techniques and algorithms to assist with finding hidden relationships between data and to assist with real time decision making

Figure 5 depicts a hierarchy of how these system help organizations learn or promote autonomy via physical robots or software bots.

Today's early adopter systems focus on identify hidden patterns and meanings in large amounts of data. As an example, one organization proved the usefulness of a "recommender engine" by applying such an AI algorithm against US Department of Education K-12 educational data collected across every public and private K-12 school in the US. More than 700 data elements were collected for each of the thousands of schools across the US. This analysis identified factors that had positive and negative correlations for educational results and would be useful in defining new school programs and new facilities at the local level. The analysis showed that positive and negative correlations differed by region and grade level (primary school, middle school, high school).

Artificial intelligence applications will assist government organizations to identify new insights from the significant amounts of structured and unstructured data collected by agencies. Agencies such as the General Services Administration (GSA) and the Department of Health and Human Services (HHS) are using AI today to gain insights and to improve operational outcomes. Examples include the following:

HHS Reimagine Program: Spend Analysis Proof of Concept
The HHS Program Support Center (PSC) uploaded 18 months of data from its five enterprise procurement systems (National Institutes of Health, Food and Drug Administration, Indian Health Service, Centers for Medicare and Medicaid Services, and Centers for Disease Control and Prevention) in January 2018. This data represented commodity purchases from 97,400 contracts and more than 1 million pages of text. HHS used natural language processing techniques to translate PDF, MSWord, Excel, handwritten notes, and other data into pure "machine data" that could be analyzed by AI algorithms. Subject matter experts and data scientists worked to ensure that the data was in the proper format to enable the computer algorithms to understand the context of the data, and what the more important contract data elements are (vendor, period of performance, pricing, units of measure, terms and conditions). The figure below shows the results and lessons learned.

97,400 Contracts Reviewed = ~1,000,000 Pages
HHS Wide Contract Data for 18 Months (FY '16 - '17)

PRODUCT/SERVICE	HIGH	LOW	% DIFFERENCE
VMware vSphere	$ 1,028.00	$ 571.00	80%
Case Copy Paper	$ 59.41	$ 27.00	102%
AED Defibrillator	$ 2,627.00	$ 946.00	178%
RN Services	$ 96.20	$ 42.00	129%
Adobe Connect (500 Seats)	$ 18,149.00	$ 4,309.00	322%

HHS Lessons Learned

- One must teach the AI tool to deal with unstructured data
- Humans have to be involved
- Requires an interdisciplinary approach – no one person or organization can be successful without collaboration
- Many view AI technologies as a threat to their job/way of life
- New and emerging technologies offer significant opportunity to remove the mundane and boring tasks, freeing people up to focus on the more critical aspects of their jobs
- There has never been a better time in federal government to embrace technology and modernize the way we do business

Figure 6: A Department of Health and Human Services AI project identified commodity procurement anomalies not recognized by acquisitions staff across the agency[31]

GSA Prediction of Regulatory Compliance: Solicitation Review Tool (SRT)[32]
The SRT AI platform uses natural language processing, text mining, and machine learning algorithms to automatically predict whether federal solicitations posted on fbo.gov are compliant with Section 508 of the Rehabilitation Act and alert responsible parties of non-compliance so that corrective actions could be taken. Through independent review, the predictions have a 95% accuracy rate. This innovation substantially alleviates the human resources needed to identify, audit, and enforce compliance.

The SRT platform is innovative because it helps GSA focus the limited resources available on the non-compliant solicitations identified and alert contracting staff to make the changes for compliance. The SRT tool is slated to go into production on cloud.gov in spring 2018. Future plans for the SRT AI platform include a scope expansion to predict whether solicitations contain other federal regulatory requirements such as cybersecurity or sustainability.

Leveraging Commercial Video Analytics to Find Killers Faster

The Federal Bureau of Intelligence is eying Amazon Web Services' video and photo analysis capability, called Rekognition, to save time for its agents, money for its budget, and ultimately, American lives. Christine Halvorsen, Deputy Assistant Director for the FBI's Counterterrorism Division, cited the FBI's work after the 2017 shooting in Las Vegas[33]:

> "*We had agents and analysts, eight per shift, working 24/7 for three weeks going through the video footage of everywhere Stephen Paddock was the month leading up to him coming and doing the shooting… If we had loaded that up into the cloud, the estimate is it would've taken us a day using Amazon Rekognition to recognize where he was in the videos. That's all we were trying to do: narrow down where in the videos he was and who he was meeting with to make sure there wasn't anybody else part of the conspiracy*."

If the shooter, Stephen Paddock, had been involved in conspiracy, the time saved by using artificial intelligence would have prevented co-conspirators from escaping or killing again. Artificial intelligence may help law enforcement respond to the next major attack to more quickly differentiate between a lone actor like Paddock or conspirators like the Tsarnaev brothers, the 2013 Boston Marathon bombers who killed a police officer during the search for them.

Chatbot Improves Customer Service

The US Customs and Immigration Services (USCIS) offers a text-based virtual assistant named Emma, after the poet Emma Lazarus, to answer questions about the services that USCIS offers.[34] Emma answers questions in both English and Spanish, handling about 6 million conversations every year. Emma's knowledge base is a combination of machine learning performed on human interactions and human analysis of the automatically generated responses. USCIS refers to this as "the best of both worlds… the powerful backend analytics of the Emma platform coupled with the subject matter experts who make the final decision to ensure the accuracy of her responses."[35] By constantly reviewing Emma's responses, USCIS is able to continually improve the chatbot to give USCIS customers the best possible experience.

Machine Learning to Improve Reliability Centered Maintenance (RCM) Analysis

The Military Sealift Command (MSC), part of the US Navy, has several maintenance and repair items documented in numerous electronic formats. A large portion of the domain data MSC possesses is unstructured text. MSC wants to have full visibility into this data to move from a preventive, condition monitoring-based maintenance approach to a proactive, reliability-based maintenance approach, while decreasing inefficiencies and cost.

MSC sought to use natural language processing to derive this information from unstructured documents. The solution applied NLP technology to extract entities, relations, and other machine information from unstructured repair documents. AI specialists reviewed and annotated sample documents and used the sample dataset to build a machine learning model in which unstructured data could be analyzed automatically, and entities (e.g. relationships between equipment data and important data) could be identified. This "training" of the machine learning model was done in an iterative manner, with feature extraction results verified and validated along the way. This effort enabled MSC to have complete visibility into a large set of unstructured data to make efficient and educated decisions regarding its ships and maintenance.

Transformation Recommendations

The decision to use AI/ML is not just technical. It is a decision that, if planned for and responsibly adopted, can be transformational to an agency's stakeholders, workforce, and long-term mission trajectory. Agencies planning to embark on the AI journey should consider the following factors while developing implementation strategies or else run the risk of slow or impeded adoption.

Demystify AI/ML for Executives. The marketing hype is in full force and it is incumbent upon agency leaders to learn the key terminology, relevant use cases, and how AI differs from previous analytic tools. Understanding the skills needed, development approaches, and characteristics of a successful AI implementation will help organizations navigate the landscape of available AI solutions, machine learning frameworks, and the large ecosystem of AI startups to determine what is real, and just as important, what is not.

Prove Mission Value Incrementally. It can be very difficult to discern which AI technologies and use cases will bring value to a mission. Take time to outline agency aspirations, assess existing value chains, and explore opportunities for where AI could solve specific mission problems. Additionally, agencies should look to not create major disruptions to their mission. Focus instead on narrowly scoped process improvements that are deployed into production environments to get executives, users, and stakeholders motivated to move onto the next AI opportunity.

Ensure Humans Remain in Control. It is important to remember that AI is both an emerging and evolving technology that comes with both opportunity and risk. To that end, agencies will need to evaluate existing laws and regulations, as well as internal governance, risk, and control policies and procedures to ensure appropriate safeguards are implemented to mitigate both intended and unintended consequences. In addition, agencies should develop a clear articulation of how much risk the organization is willing to tolerate when it comes to AI so that they can set and communicate clear parameters and business rules around topics like bias, discrimination, ethics, safety, security, privacy, explicability, accountability, compliance, and incident response.

Get Data and Infrastructure in Order. AI requires a lot of characterized and well-organized data. This data intensiveness is a stumbling block for many organizations looking to capture the value of AI. Even those with huge data assets find data is unusually hard to extract, structure, label, and organize. Given this importance, it is critical for agencies to establish strategies around data governance, defining ontologies, data capture, data engineering and maintenance, etc. Equally important is the design and implementation of an infrastructure that can support the significant computational and storage requirements associated with AI and high-performance computing capabilities. Agencies will need to plan for these types of capital and operational expenditures as they design their infrastructure and data architectures.

Invest in Talent. AI talent is scarce and the battle for hiring and retaining experts can be fierce. Agencies will need to explore and implement strategies that balance between partnering with research and academic organizations, contracting for talent, and investing in building AI talent in-house. Most agencies will not have large pools of in-house technical AI experts but can upskill those employees who have technical competencies through online and in-person technical courses. Those employees who are not technical should be trained to have a basic digital literacy (i.e., be familiar with what AI is, how it works, and how to interact with it).

Contract Differently. Given the state of the current AI technology and the availability of AI talent, contracting AI expertise or consultants can be an option for agencies that have strategies in place, use cases identified, and project goals well defined. In these scenarios, it is important that agency acquisition officials and contracting officers allow bidders to demonstrate their capabilities instead of telling you. Contracting labor-based AI support services could be done in a phased approach where a firm-fixed price Request for Proposal (RFP) is issued asking vendors to produce an initial operating capability based on the agency's data, and then switch the model to allow for more development and enterprise scale. It might also be in the government's best interest to consider buying AI solutions at cost rather than labor at level of effort. In this scenario, contracts would need to be set up to buy a mix of labor and AI solutions/software in either a subscription or consumption-based buying model.

Do Not Let Culture Constrain the Pace of Adoption. Implementing AI is really no different than other technologies except that current hyperbole is creating additional anxiety and risk across an organization. To mitigate the concern, agencies should design a change management program that communicates both a common goal and the importance of AI in empowering the workforce, not replacing them, is essential to overcoming internal and external resistance to the adoption of AI. Equally important is transparency with employees about the technology and involving them in the design of how they will ultimately work with a machine.

Conclusion

Artificial intelligence and machine learning is one on many new technologies entering IT operations. The technology can improve insights from the petabytes worth of structured and unstructured data retained by agencies. AI early adopters include many areas where analysts – without automation – must sift through tremendous amounts of data to find issue requiring attention. Figure 7 depicts the many diverse applications of AI techniques. Cybersecurity operations use algorithms that automate the search for network anomalies and patterns of behavior which may indicate a possible external cyber-attack or an insider threat. Autonomous vehicles identify vehicle threats based on identification of nearby objects by sensors. With cloud-based tools, an agency can easily stand up an AI pilot to prove the utility of this technology against a critical business problem.

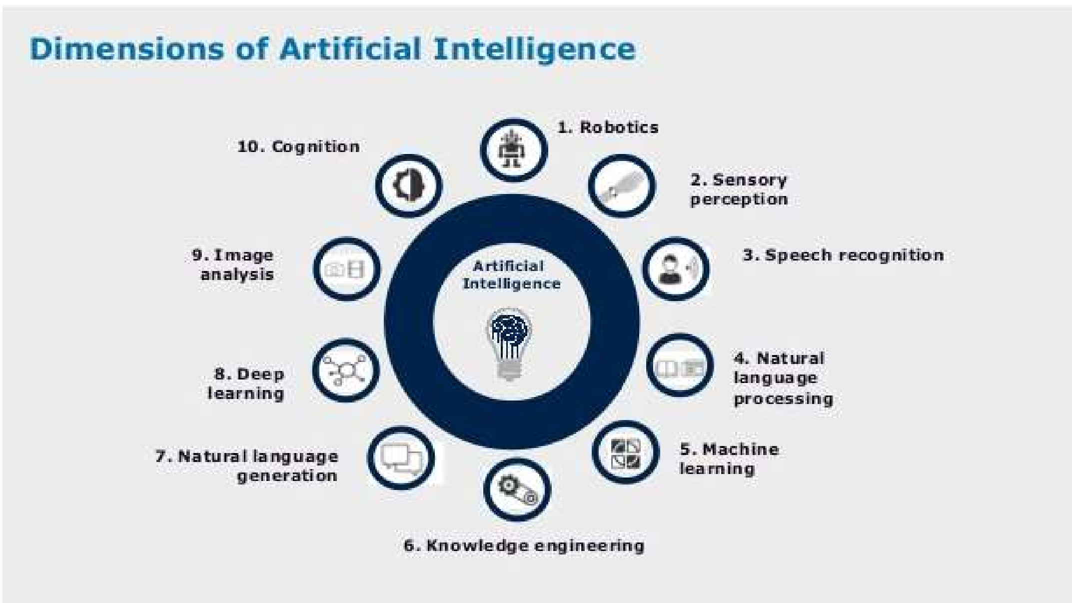

Figure 7: Artificial Intelligence techniques can assist agencies across many diverse business areas where large amounts of data makes it difficult for analysts to identify trends and new needs

These insights can assist organizations to improve citizen services, reduce waste and fraud, improve acquisition outcomes, and identify potential security threats. Agencies must carefully plan their AI/ML transformational journey and adoption to better realize the technology's potential to meet their mission objectives.

Authors

This white paper was written by a consortium of government and industry. The organizational affiliations of these contributors are included for information purposes only. The views expressed in this document do not necessarily represent the official views of the individuals and organizations that participated in its development.

The ACT-IAC Artificial Intelligence Working Group wishes to thank the many contributors who provided a tremendous amount of time and good humor to bring this primer to completion.

Sandy Barsky	General Services Administration
Luwan Bokure	IBM Corporation
Michael Bruce	Leidos, Inc.
Josh Elliot	Booz Allen Hamilton
Ken Farber	The Ambit Group
Sukumar Iyer	Brillient LLC
Venkat Kodumudi	CGI, Inc.
Aaron Margolis	Harper Paige LLC
Mallesh Murugesan	Abeyon Corporation
Jack Tatlow	IBM Corporation
Sabrina Wallace	IBM Corporation
Mimi Whitehouse	Accenture Federal Services
Karen Buckley	Fusion Applied

References

[1] Artificial Intelligence Index. http://aiindex.org/

[2] The White House. *Artificial Intelligence for the American People*. May 10, 2018. https://www.whitehouse.gov/briefings-statements/artificial-intelligence-american-people/

[3] ACT-IAC. *Intelligent Automation Primer: An overview of the journey to intelligent automation*. February 2019. https://www.actiac.org/act-iac-white-paper-intelligent-automation-primer

[4] Murgugesan, Mallesh. How do Machines Learn? November 15, 2018. https://abeyon.com/how-do-machines-learn/

[5] ibid

[6] Bayes Theorem. In Wikipedia. From https://en.wikipedia.org/wiki/Bayes%27_theorem

[7] Sutton, Richard and Barto, Andrew. *Reinforcement Learning: An Introduction Second edition*. January 1, 2018. A Bradford Book The MIT Press Cambridge, Massachusetts

[8] Adam, RL. *10 Powerful Examples of Artificial Intelligence In Use Today*. Jan 10, 2017. Forbes.

[9] Expert System. *Entity extraction: How does it work?* May 16, 2016. https://www.expertsystem.com/entity-extraction-work/

[10] ACT-IAC. *Blockchain Primer: Enabling Blockchain Innovation in the U.S. Federal Government*, October 2017. https://www.actiac.org/act-iac-white-paper-enabling-blockchain-innovation-us-federal-government
[11] ACT-IAC. *Blockchain Playbook for the U.S. Federal Government.* April 2018. https://www.actiac.org/act-iac-white-paper-blockchain-playbook-us-federal-government
[12] Microsoft Azure. *GPUs vs CPUs for deployment of deep learning models.* https://azure.microsoft.com/en-us/blog/gpus-vs-cpus-for-deployment-of-deep-learning-models/
[13] *TensorFlow: A System for Large-Scale Machine Learning*. 12th USENIX Symposium on Operating Systems Design and Implementation, https://www.usenix.org/conference/osdi16/technical-sessions/presentation/abadi
[14] Raschka, Sebatian. *Using Python for Machine Learning*. Python Machine Learning, Packt Publishing Ltd. September 23, 2015.
[15] Pontin, Jason. *Reedy, Brittle, Opaque, and Shallow: The Downsides to Deep Learning*. February 2, 2018. Wired. https://www.wired.com/story/greedy-brittle-opaque-and-shallow-the-downsides-to-deep-learning/
[16] McKinsey Podcast. *The real-world potential and limitations of artificial intelligence.* April 2018. https://www.mckinsey.com/featured-insights/artificial-intelligence/the-real-world-potential-and-limitations-of-artificial-intelligence
[17] IEEE. *Ethically Aligned Design V2.* https://ethicsinaction.ieee.org/
[18] Oakden-Rayner, Luke. *Medical AI Safety: We have a problem.* July 11, 2018. https://lukeoakdenrayner.wordpress.com/2018/07/11/medical-ai-safety-we-have-a-problem/
[19] Oaken-Rayner, Luke. Medical Safety AI. July 11, 2018. https://lukeoakdenrayner.wordpress.com/2018/07/11/medical-ai-safety-we-have-a-problem
[20] Vincent, J, and Brandom, R. *Axon launches AI ethics board to study the dangers of facial recognition*. April 26, 2018. The Verge. https://www.theverge.com/2018/4/26/17285034/axon-ai-ethics-board-facial-recognition-racial-bias
[21] Harwell, D. *Facial recognition may be coming to a police body camera near you*. April 26, 2018. Washington Post. https://www.washingtonpost.com/news/the-switch/wp/2018/04/26/facial-recognition-may-be-coming-to-a-police-body-camera-near-you/
[22] Buolamwini, Joy and Gebru, Timnit. *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*. Proceedings of Machine Learning Research 81:1{15, Conference on Fairness, Accountability, and Transparency, 2018
[23] Sahota, P. *How artificial intelligence can help unlock data*. November 15, 2018. https://www.ibm.com/blogs/cloud-computing/2018/11/15/artificial-intelligence-unlock-data/
[24] Chenok, Dan. *How can AI help government improve?* November 2018. GCN. https://gcn.com/Articles/2018/11/14/AI-improve-government.aspx?Page=2
[25] Chenok, Dan. *How can AI help government improve?* November 2018. GCN. https://gcn.com/Articles/2018/11/14/AI-improve-government.aspx?Page=2
[26] OMB Memo M-18-23 *Shifting from Low-Value-work to High-Value work*. August 2018. https://www.whitehouse.gov/wp-content/uploads/2018/08/M-18-23.pdf

[27] Fedscoop. *Federal agencies are achieving demonstrable value in AI, survey shows*. October 24, 2018. https://www.fedscoop.com/new-survey-shows-federal-agencies-already-achieving-demonstrable-value-ai/
[28] DHS. *Snapshot: S&T Founds DHS-Wide Artificial Intelligence and Machine Learning Community of Interest.* October 9, 2018. https://www.dhs.gov/science-and-technology/news/2018/10/09/snapshot-st-founds-dhs-wide-ai-and-ml-community-interest
[29] Thomas, Steff. *Women of Washington: Data science and machine learning in government*. August 30, 2017. Federal News Network. https://federalnewsnetwork.com/federal-insights/2017/08/women-of-washington-data-science-and-machine-learning-in-government/slide/1/
[30] Forbes Technology Council. *14 Ways AI Will Benefit or Harm Society*. March 1, 2018. https://www.forbes.com/sites/forbestechcouncil/2018/03/01/14-ways-ai-will-benefit-or-harm-society/#6b108fc64ef0
[31] Bruce, Michael and Pinto, Earl. Artificial Intelligence & Machine Learning: Commercially Proven Methods – Recommenders and HHS Acquisitions Results. Presentation July 2018. ACT-IAC – HHS Emerging Technology "Reimaging – Buy Smarter" Day
[32] 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. *Game Changers: Artificial Intelligence Part II; Artificial Intelligence and the Federal Government*. March 7, 2018 https://www.gsa.gov/about-us/newsroom/congressional-testimony/game-changers-artificial-intelligence-part-ii-artificial-intelligence-and-the-federal-government
[33] Mitchell, Billy. *The FBI is looking to the cloud to stop the next terrorist attac*k. November 28, 2018. Fedscoop. https://www.fedscoop.com/fbi-data-aws-christine-halvorsen/
[34] USCIS. *Emma: Friendly Presence and Innovative USCIS Resource - Available 24/7*. September 1, 2016. https://www.uscis.gov/archive/blog/2016/09/emma-friendly-presence-and-innovative
[35] Federal News Network Podcast. *Vashon Citizen: USCIS' new virtual assistant Emma gets service award.* May 31, 2018. https://federalnewsnetwork.com/federal-drive/2018/05/vashon-citizen-uscis-new-virtual-assistant-emma-gets-service-award/