ACT-IAC Emerging Technology Community of Interest

Response to NIST RFI on Artificial Intelligence

References
NIST RFI on Artificial Intelligence Standards
https://www.federalregister.gov/documents/2019/05/01/2019-08818/artificial-intelligence-standards

NIST Release
https://www.nist.gov/news-events/news/2019/05/nist-requests-information-artificial-intelligence-technical-standards-and

Table of Contents

Introduction

Technical Standards

Data and Infrastructure

AI Test and Evaluation Standards Framework

Sub-Components

Robotic Process Automation (RPA):

Intelligent Automation (IA):

Cognitive Computing:

Machine Learning (ML):

Speech to Text and Text to Speech:

Keyword Extraction, Entity Extraction, and Relationship Extraction:

Sentiment Analysis, Tone Analysis, and Personality Insights:

Natural Language Processing:

Sensors and Autonomous Mobility:

Consumer Choice Recommenders:

Behavioral Analysis:

Modeling:

Structure:

Supervised Training

Biometrics

Business Standards

Mission

Administrative Services - Mission Support

Biometrics, Biographic and Identity

Medical Safety

Contracting Standards

Bias

Ethics

Elimination of Bias

Need for Governance in AI Development:

Current Use Cases

Use Case 1: HHS Reimagine Program: Spend Analysis Proof of Concept

Use Case 2: GSA Prediction of Regulatory Compliance: Solicitation Review Tool (SRT)[7]

Use Case 3: Leveraging Commercial Video Analytics to Find Killers Faster

Use Case 4: Chatbot Improves Customer Service

Use Case 5: Machine Learning to Improve Reliability Centered Maintenance (RCM) Analysis

Use Case 6: Recommender Engine in Education Outcomes

People

Culture

Workforce

Introduction

In the beginning of 2018, the sub-committee on Information technology of the House Committee on Oversight and Government Reform held a series of hearings on Artificial Intelligence (AI) to understand America's position in the world, opportunities AI will create and the impact AI will have on US workforce. Several recommendations were proposed from the Subcommittee meeting chief among them is a call for increased engagement by the congress and administration. The United States have traditionally led the world in development of new and emerging technologies through large investments in Research & Development (R&D) and the hope is that, this trend will follow within AI as well.

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. According to Stanford's Human Centered AI Institute's AI Index 2017 Report, 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.

This response attempts to provide NIST with information on the current state of Artificial Intelligence, opportunities and challenges and provides a few use cases to bring to light current federal involvement in AI. Additionally, we would like to offer a few recommend that the creation of standards in AI should be an ongoing process as the field is evolving rapidly and new developments are uncovered regularly.

Technical Standards

Data and Infrastructure

- AI requires a lot of characterized, well-organized data. The necessity for well cleansed data that is structured sufficient to be of operational value is intensive.
- Disaggregated and disconnected data is often 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, and load (ETL).
- Given this importance of fact-based decisions, it is critical for agencies to establish strategies around collecting data, correlating information through sound governance, characterizing knowledge by sufficiently defining ontologies.
- Once accomplished they can provide the means to capture data, engineer, architect and maintain it sufficiently to the necessary ingredients of well formulated data and highly structured information to feed the algorithms formulated to provide AI capabilities.
- 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 create data architectures to establish and sustain them.
- AI and Machine Learning (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.
- In order for AI systems to perform better within a certain domain, it needs to be trained by SME's using domain data. Agencies have to make sure proper training is provided to the SME's on all factors mentioned in this paper, in order to create a robust, accurate and unbiased AI models. Additionally, continuous training of the model and learning of the model will need to be performed over a period of time to create well performing models.
- From a technology perspective, agencies should lean more toward using open source technologies and frameworks. This will provide algorithmic transparency the ability for agencies to explain AI decisions made by the model. Black box solutions will not have this capability and might make it prohibitively difficult to evaluate the decisions made by the AI systems.

- When evaluating AI systems, accessibility should be taken into consideration as an evaluation standard. Additionally, security risks will need to be considered as well.

AI Test and Evaluation Standards Framework

This section attempts to begin the process of identifying critical elements, and defining a standard which includes a set of activities suitable for testing and documenting Artificial Intelligence (AI) solutions. One of the main goals for developing a “test” standard is to provide a rubric that will enable AI developers to create innovative solutions that are “more secure, usable, interoperable and reliable.” Indeed, developing a standard is critical to ensure public trust of rapidly evolving AI solutions- so that the world can benefit from all that the field has to offer. We’ve identified the following areas as critical for testing and evaluating AI solutions.

AI SOLUTION SECURITY

1. SAFETY ASSURANCE
   The AI solution developer should identify critical portions of the solution. In this instance, critical portions are identified as portions who failures could lead to hazardous system state. A hazardous systems state is defined as something that could results in unintended death, injury, loss of property or environmental harm. The AI developer shall develop a safety assurance strategy, including both tests and analysis to ensure that the requirements, design, implementation, and operating procedures for the identified AI solution minimizes for eliminates the potential for hazardous conditions. The AI developer shall record the strategy in the AI development plan, implement the strategy and produce evidence, as part of the AI requirement products that the safety assurance strategy has been carried out.

2. SECURITY ASSURANCE
   The AI solution developer should identify all portions of the product, whose failure could lead to a breach of system security. If this weakness exists in the solution, the AI developer shall develop security assurance strategy to ensure that the requirements, design, implementation and operating procedures for the identified AI Solution are minimized or eliminate the potential for breaches of system security assurance strategy has been carried out.

3. PRIVACY ASSURANCE
   The AI solution developer should identify as privacy critical portions whose failure could lead to a breach of system privacy. If this exists in a solution, the AI developer shall develop privacy assurance strategies to assure that the requirements, design, implementation, and operating procedures for the identified software minimize or eliminate the potential for breaches of system privacy. The AI developer shall record the strategy in the software development plan, implementing the strategy, and

produce evidence as part of required software products, that the privacy assurance strategy has been carried out.

4. ASSURANCE OF OTHER CRITICAL ELEMENTS / REQUIREMENTS
   If the AI solution relies on the solution to satisfy other requirements deemed critical by the contract or by system specifications, the developer shall identify those portions whose failure could lead to violation of those critical requirements; develop a strategy to assure that the requirements, design, implementation, and operating procedures for identifying software minimize or eliminate the potential for such violations' record the strategy in the AI solution development plan, implement the strategy and assure that the strategy has been carried out.

AI SOLUTION TEST PLANNING

5. AI SOLUTION TEST PLAN
   The AI developer will develop and record plans for conducting qualification testing. This plan will include an AI solution test plan.

6. AI SOLUTION INSTALLATION PLAN
   The AI developer will develop and record plans for performing software installation and training at user sites as specified. This planning will include all applicable items in the plan.

7. AI SOLUTION TRANSITION PLANNING
   The AI developer will identify all resources that will be needed by the support organization.

AI SOLUTION UNIT TESTING PREP

8. AI SOLUTION IMPLEMENTATION AND UNIT TESTING
   The developer will perform software implementation and unit testing in accordance with the following requirements.

9. AI SOLUTION IMPLEMENTATION
   The AI developer shall develop and record software corresponding to each software unit in the design. This activity will include, as applicable, coding computer instructions and data definitions, building databases, populating databases and other data files with data values, and other activities needed to implement the design. For deliverable solutions, the developer shall obtain acquirer approval to use any programming language.

10. PREPARING FOR AI SOLUTION UNIT TESTING
    The developer shall establish test cases in terms of inputs, expected results, and the evaluation criteria. Test procedures, and data for testing the software corresponding to each AI solution unit. The test cases shall cover all aspects of the unit's details

design. The AI developer shall record this information in the appropriate AI development files.

11. PERFORMING AI SOLUTION UNIT TESTING
The developer will test the AI solution corresponding to each solution unit. The testing will be in accordance with unit test cases and procedures.

12. ANALYSIS AND RECORDING AI SOLUTION UNIT INTEGRATION AND TEST RESULTS
The AI developer will analyze the results of unit integration and testing and shall record the test and analysis results in appropriate solution development files.

13. TESTING ON THE TARGET SYSTEM / SOLUTION
The qualification testing will include testing on the target system or an alternative system approved by the acquirer.

14. PREPARING FOR AI SOLUTION QUALIFICATION TESTING
The AI solution developer will define and record the test preparations, test cases, and test procedures to be used for qualification testing and the traceability between the test cases and the requirements. The results shall include all applicable items in solutions. The developer shall prepare the test data needed to carry out the test cases and provide the acquirer.

15. DRY RUN QUALIFICATION TESTING FOR AI SOLUTION
Qualification testing shall be witnessed by the acquirer and the test cases and procedures to ensure they are complete and accurate and that the solution is ready for witnessed testing. The AI developer shall record the results of this activity in appropriate software development files and shall update the test cases and procedures as appropriate.

16. PERFORMING SOLUTION QUALIFICATION TESTING
The AI solution developer should participate in system qualification testing. This participation will be in accordance with the system test cases and procedures.

17. REVISION AND RETESTING FOR THE AI SOLUTION
The developer shall make necessary revisions to the AI solution, provide the acquirer advance notice of retesting, participate in all necessary retesting, and update the solution development file and other products as needed, based upon the results of the solution qualification testing.

18. ANALYZE AND RECORD THE AI SOLUTION TEST RESULTS
The developer will participate in analyzing and recording the results o system qualification testing. For AI solutions, the results will include all applicable items in the test report.

AI SOLUTION TESTING

19. AI SOLUTION PERFORMING UNIT INTEGRATION TESTING

The AI solution developer will establish tests cases in terms of inputs, expected results and evaluation criteria. Additionally, the AI solution developer will identify test procedures, and test data for conducting unit integration and testing. The test cases shall cover all aspects of the architectural design. The AI solution developer shall record this information in the appropriate development files.

20. AI SOLUTION REVISION AND RETESTING

The AI solution developer shall make all necessary revisions to the solution, perform all necessary retesting and update the development files and other software products as needed, based on the results of unit integration and testing.

21. AI SOLUTION ANALYZING AND RECORDING UNIT INTEGRATION AND TEST RESULTS

The AI solution developer will analyze the results of the unite integration and testing and shall record the test and analysis results in appropriate solution files.

AI SOLUTION QUALIFICATION TESTING

22. INDEPENDENCE IN AI SYSTEM MALFUNCTION TESTING

The person(s) responsible for qualification testing of the developed AI solution shall not be the person who performed detailed design or implementation of that solution. This does not preclude persons who performed detailed design or implementation of the of the AI solution from contribution to the process, for example, by contributing test cases that rely on knowledge of the AI's internal implementation.

23. TESTING ON THE TARGET COMPUTER SOLUTION

The AI solution qualification testing shall include testing on the target computer system or an alternative system approved by the acquirer.

24. PREPING FOR IA QUALIFICATION TESTING

The AI developer will define and record the test preparations, test cases, and test procedures to be used for the AI qualification testing and the traceability between the test cases and the requirements. The developer will prepare the test data needed to carry out the test cases and provide the acquirer advanced notice of time and location for qualification testing.

25. DRY RUN AI SOLUTION QUALIFICATION TESTING

The qualification testing is to be witnessed by the acquirer, the developer will dry run the test cases and procedures to ensure that they are complete and accurate and that the solution is ready for witnessed testing. The developer will record the results of this activity in appropriate solution development files and shall update the test cases and procedures as appropriate.

26. PERFORMING AI SOLUTION QUALIFICATION TESTING

The AI solution developer will perform qualification testing for each item. The test shall be in accordance with the test cases and procedures.

27. AI SOLUTION REVISION AND RETESTING
The AI developer will make the necessary revisions to the solution, provide the acquirer advance notice of retesting, conduct all necessary retesting, and update the solution development files and other products as needed based on the qualification testing.

28. ANALYZING AND RECORDING AI SOLUTION QUALIFICATION TEST RESULTS
The AI developer will analyze and record the results of the qualification testing. The results will include all applicable items in the test report.

Sub-Components

Robotic Process Automation (RPA):

A technology platform enabling robots to interact with applications. RPA software can be programmed to do basic tasks across applications just as human workers do; such as cutting and pasting information between application and entry boxes. 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 to include invoice entry, human resources personnel action entry, and data verification to name a few. RPA can bridge the automation gap serving as the on-ramp to building the automated superhighway, by permitting the subject matter expert (SME) for a given process to be trained to integrate an RPA bot, and therefore democratizing the integration of process automation at a basic level.

Intelligent Automation (IA):

Uses AI techniques to improve business process automation by making cognitive decisions and taking actions. IA is an evolution of RPA whereas it automates processes that are relatively simple and static. Through IA discovery the opportunity to apply intelligent business rules affords the means 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 for comparative analysis to validate and make assessments in an effort to autonomously take an iterative approach to improve operations through a prognosis in which to learn and evolve without being explicitly programmed. Machine learning is the application of AI techniques using statistical methods that enable machines to optimize correlations through the validation of additional data as it is introduced into the model. Machine learning evolves through the resulting feedback loops impeded in the system via the connectionist theory of human cognition through pattern recognition and computational learning theory. Approaches include neural networks and deep learning.

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.

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[1].
  - 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.

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.

Natural Language Processing:

Are advancing to provide fluid interfaces for interactive voice recognition acting as personal assistants in today's 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:

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:

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:

Less well known than consumer choice recommenders, these systems seek to predict future behaviors of persons based on AI algorithms that correlate adaptations of 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.

Modeling:

Given the many approaches available to manage and monitor autonomous systems, it is important to have standards to ensure consistent results that provide contextual understanding. Thus the following frameworks upon which to quantify the observations that assure similar perspectives are applied to ensure results that are are reliable and consistent across the enterprise:

Descriptive modeling: A mathematical process that describes and de-conflicts the definitions and references to data to define the facts that inform the knowledge base that evolve and contextualize understanding. Through this de-confliction process data quality is evaluated and validated thus improving the common references it.

Predictive modeling: A statistical technique to prescribe the present state informed by past activities to elevates situational awareness. These models offer a quantifiable insight as to the probability of current outcomes. Thus the more historical data available afford greater insights into the probability of things occurring in the present.

Prescriptive modeling: Informs the cost/benefit analysis to ascertain the available options given the current circumstances. The result is a mechanism to consistently holistically weight the operational benefits to be derived from actions taken measured against results to be achieved across the organization.

Potentiality modeling: Assessing the probability of potential outcomes given past influences and how they affect the environment. Through pattern correlations and temporal trends, future activities can be ascertained by correlating past performance's influences and their potential impact upon the present state. Thus the advantage by Winston Churchill of "The farther back you can look, the farther forward you are likely to see" holds true.

Deterministic Model: A deterministic model assumes certainty in all aspects. The output of the model is fully determined by the parameter values as related to current conditions. Deterministic models are based on known rules so that given input values will consistently 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 in a variety of different conditions. 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 and the interdependencies that influence the environment to effect the outcome as conditions and circumstances evolve.

Structure:

The most important stand is that of the identifiers and descriptors of data. As the maturation of data in the Enterprise Information Model (EIM) is continually updates and thus changes the data. Thus, we must appreciate that relative to the dependencies within the computation equation relations rarely if ever change.

Therefore, consistency in the METAtag is important to ensure shared data is fit for purpose and operationally relevant. META should identify the WHO, WHAT, WHERE as data is registered into an enterprise data dictionary (EDD). This allows the formulation of the Information Asset Catalog (IAC) which informs the community WHAT data is available, WHO is the Subject Matter Expert (SME), and HOW to get access to it. Additionally the EIM puts a Date Time Stamp (DTS) on it to ensure WHEN the data is pulled, the most current (version control) is made available during aggregation process.

We must stop trying to standardize data for this precludes its evolution and creates a suboptomized outcome as a result of antiquated or outdated inputs. Therefore the META is critical to provide minimum linkage to correlate disparate data sets. The tags serve as a descriptor for the purposes of validation, deconfliction, and maturation during the enrichment process.

Standards for the META must be invoked and processes for deconfliction of the tags is imperative if we are able to consistently create referential information that can be shared and collaborated upon across the information environment.

Supervised Training

- 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.
- 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.[2]

Biometrics

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.

Business Standards

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.

These insights can assist organizations to appreciate the dimension and complexities associated with selecting and applying the right capabilities in sufficient capacity to effectively leverage the plethora of options to capitalize on AI opportunities.

- 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.
- 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.

Mission

- It can be very difficult to discern which AI technologies and use cases will bring value to a mission.
- Taking time to outline agency aspirations, assess existing value chains, and explore opportunities for where AI could solve specific mission problems can help prioritize where to start.
- 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.
- In the end, it is through an elevated awareness of how to apply AI which will ultimately assure the appropriate application of this capability 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.

Administrative Services - Mission Support

The Federal government alone spends more than **$25 billion / year** of taxpayer money on common mission-support services such as processing hiring transactions, managing Federal finances, travel, and payroll (reference: M-19-16, Centralized Mission Support Capabilities for the Federal Government).

As the Federal government continues its progressive journey to achieve greater efficiency across common mission-support functions such as, but not limited to:

- Financial Management
- Human Capital
- Procurement
- Travel
- Grants Management
- Real Property Management
- Records Management

The opportunity to leverage this work as a force multiplier for automation towards outcomes of improved timeliness, accuracy, and cost savings should not be lost.

Federal, state, and local governments as well private sector organizations rely substantially on these common supportive functions to promote and drive effectiveness of mission. When automation opportunities in this space are properly implemented in the support lifecycle, these translates into lower spend and increased performance for mission.

An AI Standard should reference as a pointer the substantial progress and cross-governmental investment already underway within each of these mission support functional areas. Specifically, as an AI is implemented to automate financial management, procurement, human capital and other functional areas these AI implementations and standards should reference work that has been achieved by the cross-functional Federal Business Standards Council to develop a **Federal Integrated Business Framework (FIBF)** which establishes government-wide standards for end to end processes, use cases, data items, and performance.

The FIBF (https://ussm.gsa.gov/fibf) is a model that enables the Federal government to better coordinate and document common business needs across agencies, focusing on outcomes, data, and cross-functional end-to-end business processes. It is the essential first step towards standards that will drive economies of scale and leverage the government's buying power. The FIBF includes five components:

- **Federal Business Lifecycles, functional areas, functions, and activities** serve as the basis for a common understanding of what services agencies need and solutions that should be offered.
- **Business Capabilities** are the outcome-based business needs mapped to Federal government authoritative references, forms, and data standards.
- **Business Use Cases** are a set of agency "stories" that document the key activities, inputs, outputs, and other LOB intersections to describe how the Federal government operates.
- **Standard Data Elements** identify the minimum data fields required to support the inputs and outputs noted in the use cases and capabilities.
- **Performance Metrics** define how the government measures successful delivery of outcomes based on timeliness, efficiency, and accuracy targets.

More information about the government-wide mission support lifecycle standards expressed in the FIBF can be found at https://ussm.gsa.gov/fibf

This model is being used to develop government-wide standards for mission support based on auth that will be used to work with industry and agencies in the implementation and delivery of mission support

Biometrics, Biographic and Identity

- 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.”[5]
- 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.

Medical Safety

- 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.
- 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.

Contracting Standards

- 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.

Bias

- Is not inherently bad, it can be a critical factor in a decision process, and positive where transparent
- Transparency is essential so that a bias can be known and "judged", and risk determinations made.

Ethics

- 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. .
- 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[6], should be mirrored in businesses and working groups of all sizes. Three different areas areas must be considered in this area which is garnering a lot of interest.
- Safeguarding the personal nature of data, autonomous systems parametric frameworks, cultural and cognitive bias are areas worth exploring in consideration of the ethical dilemma faced consciously and unconsciously when developing autonomous self-governed systems with the capacity to evolve.
- When deciding how to build, assess and implement an ML-based system, the implementer of a system has to balance several competing interests and be willing to defend its choice, in public and possibly in the courts.
- Federal agencies will have to choose whether to prioritize accuracy, false positive rate, false negative rate, disparate treatment (affirmative action or discrimination of a class of people) or disparate impact (a criterion negatively affects a class of people). To avoid controversy and potential liability, experts in ethics and the law

should be involved as soon as possible rather than after the system is affecting people's lives. Two concrete examples of these choices are airport screenings and recidivism risk scores.

- Airport screeners accept a high false positive rate (people who receive additional scrutiny but are not dangerous) in order to get a low false negative rate (letting dangerous people and weapons onto flights). This choice, however, can have a disparate impact on transgender persons because people who wear a lot of objects under their clothes will get picked out for additional scrutiny every time they fly.
- As another example, some state courts use a recidivism risk score algorithm, which calculates how likely a convicted criminal is to commit more crimes, during sentencing or parole hearings. This algorithm does not include race, so there's no disparate treatment. But the algorithm does include whether a person has received public assistance, which improves accuracy but leads to disparate impact against minorities.

Elimination of Bias

- Many AI systems[3] 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 systems underlying algorithm.

Need for Governance in AI Development:

AI development needs to be monitored and governed with diligence. The risks with lax regulatory governance and oversight will lead to biased, socially harmful systems that will pose significant threat to federal systems. AI policy needs to be put in place that will drive more transparency and explainability for solutions that are being built. This will prevent bias and risk that can stem from black box solutions and will foster trust within the public.

Current Use Cases

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. Example include the following:

Use Case 1: 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.

Use Case 2: GSA Prediction of Regulatory Compliance: Solicitation Review Tool (SRT)[7]

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 in 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.

Use Case 3: Leveraging Commercial Video Analytics to Find Killers Faster

The Federal Bureau of Intelligence is eying video and photo analysis capability to save time for its agents, money for its budget, and ultimately American lives. The Deputy Assistant Director for the FBI's Counterterrorism Division, cited the FBI's work after the 2017 shooting in Las Vegas[8]:

"*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 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 had involved in conspiracy, the time saved by use of artificial intelligence would have prevented co-conspirators from escaping or killing again. Artificial intelligence may help law enforcement respond to the next major attach to more quickly differentiate between a lone actor or conspirators in the 2013 Boston Marathon bombers who killed a police officer during their search for them.

Use Case 4: 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.[9] 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."[10] By constantly reviewing Emma's responses, USCIS is able to continually improve the chatbot to give USCIS customers the best possible experience.

Use Case 5: 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.

Use Case 6: Recommender Engine in Education Outcomes

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).

People

Culture

- 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 being transparent with employees about the technology and involving them in the design of how they will ultimately collaborate with a machine.

Workforce

- 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*).
- 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.
- 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.

- 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[11].
- 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[12].

Credits

- ACT-IAC Artificial Intelligence Sub-Group of Emerging Technoilogy Community of Interest
- Nevin Taylor, General Services Administration
- Joyce Hunter, Vulcan Enterprises LLC
- Sandy Barsky, General Services Administration
- Janelle Billingslea, Department of Health and Human Services
- Mallesh Murugesan, Abeyon
- Robert Wuhrman, General Services Administration

References

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

[2] 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

[3] 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

[4] 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/

[5] 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/

[6] IEEE. *Ethically Aligned Design V2.* https://ethicsinaction.ieee.org/

[7] 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

[8] 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/

[9] 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

[10] 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/

[11] 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/

[12] Chenok, Dan. *How can AI help government improve?* November 2018. GCN. https://gcn.com/Articles/2018/11/14/AI-improve-government.aspx?Page=2