

SYM-AM-24-070



EXCERPT FROM THE
PROCEEDINGS
OF THE
TWENTY-FIRST ANNUAL
ACQUISITION RESEARCH SYMPOSIUM

**Acquisition Research:
Creating Synergy for Informed Change**

May 8–9, 2024

Published: April 30, 2024

Approved for public release; distribution is unlimited.

Prepared for the Naval Postgraduate School, Monterey, CA 93943.

Disclaimer: The views represented in this report are those of the author and do not reflect the official policy position of the Navy, the Department of Defense, or the federal government.



The research presented in this report was supported by the Acquisition Research Program at the Naval Postgraduate School.

To request defense acquisition research, to become a research sponsor, or to print additional copies of reports, please contact any of the staff listed on the Acquisition Research Program website (www.acquisitionresearch.net).



ACQUISITION RESEARCH PROGRAM
DEPARTMENT OF DEFENSE MANAGEMENT
NAVAL POSTGRADUATE SCHOOL

Enhancing Acquisition Outcomes through Leveraging of Artificial Intelligence

Christopher R. Barlow—is a seasoned Senior Acquisition Analyst with over 6 years of experience who has demonstrated expertise in audit risk mitigation, acquisition data analytics, process development and optimization, as well as demonstrated experience with independent assessments, challenge-based acquisitions (ChBA), artificial intelligence applications, and acquisition innovation. His diverse experience spans multiple federal agencies. Barlow holds a bachelor's degree in business administration from the University of Pikeville and is a Certified Business Intelligence Professional from Villanova University. [cbarlow@mitre.org]

Kevin M. Forbes—is a Principal Software Engineer at the MITRE Corporation, with over 20 years of software development experience. His expertise spans a wide range of technologies, including full stack software engineering, agile/DevOps, mobile apps, robotics, cloud computing and artificial intelligence. Forbes is deeply passionate about leveraging technology to address critical and urgent problems to have mission impact. His current work centers on the application of Generative AI to tackle challenges in government acquisition and software development life cycle. Forbes earned his bachelor's degree in computer science from Georgetown University and a master's degree in systems engineers from Johns Hopkins University. [forbes@mitre.org]

Rachel T. Giachinta—an early careerist passionate about applying cutting edge acquisition strategies to efforts that enhance U.S. access to, and competitive utility of, emerging technologies, Rachel Giachinta holds a bachelor's degree from the University of Virginia in global Middle Eastern and South Asian studies and a minor in global sustainability, laying the foundation for her interdisciplinary and critical thinking skills that she exercises daily at MITRE. With a diverse support portfolio across DoD and civilian agencies, Giachinta has supported both high-profile acquisitions and innovative pilot programs, approaching both as a mission driven problem solver. [rgiachinta@mitre.org]

Jay Kim—with an 8-year career in government and corporate contract acquisitions, Jay Kim has cultivated a rich array of experiences. His roles, including Contracting Officer Representative (COR), Contract Specialist, Contract Lead, and Acquisition Decision Analyst, have allowed him to refine his skills within the government sector and at the MITRE Corporation. Kim has made significant contributions to various aspects of government acquisitions, including acquisition planning and strategies, market and technical research, cost analysis, and contract negotiations and awards. He holds a master's degree in operations management, a bachelor's degree in political science, and an associate degree in accounting. [jaykim@mitre.org]

Zachary G. Levenson—is a Senior Contract Analyst at the MITRE Corporation who supports Intelligence Community, Defense, and Federal sponsors throughout their acquisition life cycles. His interests include security and counterintelligence risk associated with the acquisition and contracting practices of the federal government. He has 5 years of experience in acquisition and contracting and is a former Subcontract Administrator. He holds an BA in political science from West Virginia University and a Certificate in Contract Management from Villanova University. [zlevenson@mitre.org]

Ryan M. Novak—acquisition Outcome Lead from the MITRE Corporation and AI Lead for Acquisition Innovation Center, has earned two MBAs, an MS in strategic purchasing, and an MS in project management, and is DAWIA III Contracting certified. He has authored the Challenge-Based Acquisition Handbook, taught the approach to 500+ staff at MITRE, Industry, and the Government, and helped numerous agencies employ innovative acquisition approaches. He is a former USAF Contracting Officer with 28 years of leadership experience in the areas of innovation, contracting, acquisition, negotiation, strategic purchasing, consulting, and project and program management. He is a published author, speaker, and practitioner on innovative acquisition approaches and solutions. He helped achieve strategic program and acquisition successes for the DoD, Civilian Agencies, and the IC across the board. Novak led the start-up on iCAMS, has held an unlimited CO warrant, and completed both the Executive Strategy Program at Oxford and the Accelerated Management Program from Yale University. [rnovak@mitre.org]



Justin G. Raines—is a Principal Decision Analyst at the MITRE Corporation, where he assists government teams to develop effective acquisition strategies. His work focuses on developing innovative capability and industry communication in order to improve product delivery and meet national priorities. Raines has extensive experience in nontraditional procurement strategies, challenge-based acquisition, market research, and strategic planning. He is currently an Acquisition Program Manager in the USAF Reserves within the Kessel Run program office; he holds a bachelor's from the USAF Academy and a Master of Public Policy from the University of Maryland. [jraines@mitre.org]

Stephen W. Roe—managing Director at MITRE for Veterans Affairs programs since 2020, has significantly contributed to advancing the VA's strategic initiatives by generating cross-federal collaboration opportunities and developing several agile teams across VA's field offices. With a MITRE tenure starting in 2004, his efforts have enhanced system strategies and business transformations within the VA and beyond. Previously, as Portfolio Manager and National Background Investigation Services (NBIS) Project Leader, he led a 45-member team on a \$15 million budget DevSecOps project. Roe also co-authored MITRE's Challenge-Based Acquisitions (ChBA) Handbook. He holds two systems engineering degrees from Virginia Tech, Johns Hopkins University, and an MBA from the University of Virginia's Darden School of Business, and is the Pride Council Business Resource Group Co-Chair. [sroe@mitre.org]

Abstract

The extraordinary advancement of Artificial intelligence (AI) technology emerges at a critical juncture in which the Federal acquisition workforce is ill-equipped to meet the sky rocketing demand for products and services, alike. AI poses the opportunity to overcome data-intensive, laborious tasks and expedite the speed in which acquisition professionals operate; potential benefits may increase efficiency, enhance transparency, and reduce workload. While the use of AI across the Federal Government differs between agencies, the significance and scrutiny of Government Acquisition makes implementing AI across the acquisition process uniquely challenging. This paper will explore the current state of AI; who (i.e., which agencies) and how AI currently supports the acquisition process across the Federal Government. Next, the future state of AI and anticipated applications for the acquisition community will be discussed...think the future, think the next generation of Acquisition! This will be developed through strategic exploration across thought leaders, academic research, and working within our own AI model for acquisition. Next, we will discuss how the risks of this new technology -- new tools and novel concepts -- introduce both procedural, ethical, and operational risks that must be taken into consideration. Finally, we will offer a set of recommendations on how best to implement AI in the acquisition process as well as a list of best practices to maximize utility, mitigate risks, and ensure the acquisition workforce is well positioned to embrace the benefits and efficiencies of integrating AI capabilities.

Introduction

Artificial intelligence (AI) has stimulated widely divergent popular opinions and is a topic at the forefront of thought leaders' attention. As the world grapples with understanding AI's potential impacts, especially in augmenting or eliminating existing processes, these technologies continue to gain prominence. So-called "generative AI" has captured attention due to its ability to replicate human talents such as producing text, audio, and imagery. AI takes only seconds to perform activities that can take humans years to master. Beyond its "Wow" factor, generative AI has the potential to streamline labor-intensive and tedious tasks.

This paper explores the opportunity that judicious use of AI provides in reducing or eliminating tedious, labor-intensive acquisition tasks that often divert attention from strategic planning and execution. While the acquisition profession has already embraced technology, the Federal Acquisition process remains notoriously slow in acquiring and ultimately delivering goods, services, and/or capabilities to those who need them.



This paper does not propose a future in which AI completely replaces human involvement in acquisition activities. Instead, it investigates how AI can serve as a supportive tool that augments the workload of basic, routine, standardized actions and, thus, enables Federal Acquisition to operate more efficiently and rapidly. It envisions a balanced integration of AI in acquisition by leveraging AI's potential to enhance human productivity and decision-making, thereby increasing efficiency without disregarding human oversight and expertise.

Acquisition often views itself as more of an art than a science, valuing the unique perspectives that individuals contribute to the decision-making process. Within this framework, it is understood that certain parts of the acquisition process should continue to benefit from individuals' expertise. However, some portions of the process involve extensive manual handling of data, which can distract acquisition experts from concentrating on overarching strategic objectives. This paper envisions a world in which the art of acquisition is blended with scientific components of AI, enabling acquisition professionals to focus, prioritize, and work smarter, not harder.

Unraveling how the Federal Acquisition process uses AI today reveals that some acquisition professionals acknowledge its value and have adopted and implemented the technology, despite the perception of AI's relative immaturity. To accomplish this, the paper begins by examining the current state of AI in the Federal Acquisition community, including insight from early adoptions and lessons learned. The discussion then shifts to the future, as envisioned by technologists and forward-leaning acquisition professionals. The discussion identifies tangible points where AI can enhance and expedite the acquisition lifecycle.

Finally, the paper contains recommendations on what acquisition professionals should consider or do today. Whether they fear or overestimate and potentially even sensationalize the impact of AI, acquisition professionals must rid themselves of common misconceptions around AI, such as the belief that AI will have an immediate and overwhelming impact on their work environment. Like it or not, the transition to AI has already begun. Organizations leading this shift are proceeding cautiously with development, testing, and integration. Many of the personnel within these flagship organizations may still be unaware of AI's inevitable integration into their workflows. This indicates that new adopters must integrate AI into their operations gradually, building employee trust and confidence in AI security and the validity of its outputs, while maintaining the "trust but verify" model. This process will require strategic planning and careful communication. As Executive Order (EO) 14110 states, "Artificial Intelligence must be safe and secure. Meeting this goal requires robust, reliable, repeatable, and standardized evaluations of AI systems, as well as policies, institutions, and as appropriate, other mechanisms to test, understand, and mitigate risks from these systems before they are put to use" (The White House, 2023c).

Understand the Nomenclature

Before exploring the impact of generative AI on the Federal Acquisition lifecycle, one must first understand what AI *is* and *is not*. The extraordinary speed at which AI has emerged and its wide-ranging applications have led to a proliferation of definitions, potentially hindering practical, consistent, and regulated implementation. Acquisition professionals should consider this section as a reference source on the essential attributes of AI!

According to the American National Standards Institute (ANSI) International Committee for Information Technology Standards (INCITS), AI is "a branch of computer science dedicated to creating systems that perform task associated with human intelligence, like reasoning, learning, and self-improvement." Generative AI—the focus of this paper—is a specific branch of AI, akin to how engineering encompasses various specialties such as mechanical, electrical, etc. It is a type of artificial intelligence that can learn from and mimic large amounts of data to



create content such as text, images, music, videos, code, and more, based on inputs or prompts (Harvard University Information Technology, n.d.).

The overall discipline of AI has subfields:

- **Machine learning (ML)** is the capability of a machine to imitate intelligent human behavior, enabling computers to learn without explicitly being programmed (MIT Sloan, 2021).
- **Natural language processing (NLP)** combines computational linguistics—rule-based modeling of human language—with statistical, machine learning, and deep learning models to give computers the ability to understand text and spoken words in much the same way human beings can (IBM, n.d.).
- **Large language models (LLMs)** “are neural network models designed to process sequential data (e.g., can be trained by giving it access to a large corpus of text (such as Wikipedia, digitized books, or portions of the Internet) and using that input text to learn to predict the next word in a sequence, given what has come before” (Brynjolfsson et al., 2023).

Before making or implementing policy, decision-makers must first understand this basic taxonomy; it underpins a comprehensive understanding of how AI works, and it informs the ways that AI can be effectively and safely harnessed.

Current State: AI in Federal Acquisition in the Year 2024

The legislative and executive branches of Government have recently focused significant attention on AI policies. The 117th Congress, for example, introduced 75 AI-focused bills, of which six (6) were enacted (Congressional Research Service, 2023). From initiatives such as the National Artificial Intelligence Initiative (NAII) Act of 2020, which aims to codify and support American AI investment and advancement, to EO 14110, which has the primary objective of regulating that advancement amidst rapid expansion, official guidance has attempted to unify the fragmented approach to developing and deploying AI technology.

By design, legislation and executive branch policy documents often provide conceptual rather than practical guidance on how government organizations should leverage and develop AI for their specific missions. Examples include the National AI R&D Strategic Plan (The White House, 2023b), National AI Research Resource Roadmap (The White House, 2023a), and the Office of Management and Budget’s Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (Young, 2023), although other AI-based policies abound. This absence of detailed direction creates barriers to action. Identifying these gaps can aid acquisition organizations in developing future guidance and approaches for integrating AI that enhances the Federal Acquisition process.

Across the Federal Government, various working groups, governance bodies, and agency initiatives provide disjointed, local guiding principles, scope, frameworks, and tools. In order to implement the AI-driven acquisition enhancements of tomorrow, today’s acquisition professional must transcend these silos and understand insights and contributions from today’s AI pioneers.

Fundamentally, AI is no longer a theoretical technology, poised to shape our future defense initiatives; rather, it is already here, influencing the acquisition process in a host of ways. Acquisition leaders must accept this and securely embrace the power that AI brings. While the future of what AI can do remains relatively opaque, its immediate impact on the acquisition process can already be imagined. Interviews by this authorship team with acquisition thought leaders and AI technologists revealed potential “quick win” areas for AI in Federal



Acquisition. To demonstrate the power of AI, the authors of this paper cross-referenced the collected feedback with a ChatGPT-generated list of “how the Federal Government should use AI/ML in Procurement” (DeITek, 2023). Table 1 presents the ChatGPT-suggested areas, with green boxes highlighting those also identified by interviewed subject matter experts (SMEs).

Table 1. AI-Identified Areas for AI Application in Federal Acquisition

Identified only by SMEs	Identified by Both SMEs and ChatGPT	Identified only by Chat GPT
Requirements Development	Supplier Performance Monitoring	Market Intelligence
Drafting Contract Language	Vendor Selection	Predictive Maintenance
Commercial Analysis Software for Drafting Clauses in Compliance with Plain Writing Principles	Contract Compliance Monitoring	Demand Forecasting
Drafting Statutes and Regulations	Spend Analysis	Risk Assessment
Developing Evaluation Criteria	Contract Management	Fraud Detection
	Workflow Automation	

Table 1 does not suggest that AI should be the arbiter of how it is applied, as shown by its omission from the possible application areas identified by SMEs. Nevertheless, the green areas in Table 1 indicate that AI can aid in brainstorming ideas that human SMEs can refine and expand upon.

Industry has also recognized contract management and consistency as key areas in which AI can augment human skills. The *Harvard Business Review* (Rich, 2020) reports that AI can reduce human error in contract management and compliance monitoring by identifying and extracting key data points. Feeding historical contracts into AI systems could help establish performance standards or evaluation criteria. However, human oversight is essential to verify the relevance of AI-generated wording and to capture a program’s contractual nuances.

Some of these techniques are indeed being applied to the Federal Acquisition process today. Table 2 shows a sampling of the current use of AI tools in Federal Acquisition, serving as proofs of concept for the potential applications detailed above. This should inspire, not limit, ideas for broader AI adoption. While these tools provide foundational capabilities, organizations should customize tools to meet specific Federal Acquisition needs, drawing on lessons from these early use cases to improve future iterations.

Table 2. Current State of “Who’s Using What” AI in Federal Acquisition

Organization	Tool
Department of Agriculture (USDA)	<i>Acquisition Approval Request Compliance Tool.</i> Utilizes the text in the procurement header and line descriptions within USDA’s Integrated Acquisition System (IAS) to determine the likelihood that an award is IT-related and therefore might require an After-Action Report (AAR). The model uses the text characteristics for awards that have an AAR number entered into IAS and then calculates the probability that those



Organization	Tool
	procurements that did not have an AAR Number entered in IAS are in fact IT related (USDA, n.d.).
Health and Human Services	<i>Federal IT Acquisition Reform Act (FITARA) Tool.</i> Automates the identification of National Institute of Allergy and Infectious Disease (NIAID) IT-related contracts (HHS, 2023).
Treasury Department	<i>Digital Accountability and Transparency Act (DATA) Bot.</i> Automates the verification process of reports coming from the Internal Revenue Service (IRS) Federal Procurement Data System (FPDS) with the information in contract documents. Leverages NLP to extract unstructured information from contract documents and uses F1 scores to measure performance of validation models for each specific data element (U.S. Department of Treasury, 2024).
Department of Labor (DoL)	<i>Intranet Website Chatbot Assistant.</i> Uses a conversational chatbot on DoL intranet websites to help answer common procurement questions as well as questions about specific contracts (DoL, n.d.).
Department of State – With 37 reported AI use cases, two serve the acquisition process.	<i>Federal Procurement Data System (FPDS) Auto-Populate Bot.</i> Automates the data entry in FPDS to reduce the burden on the procurement staff and drive improved compliance on DATA reporting. It is used to update ~300 FPDS awards per week.
	<i>Production Service Code Automation ML Model.</i> Scans unstructured, user-entered procurement data (e.g., requisition title, line descriptions) to automatically detect the commodity and service types being purchased and thus enhance procurement categorization (Department of State, 2023).
General Services Administration (GSA)– With 12 reported AI use cases, five pertain specifically to the acquisition process (AI Inventory — Tech at GSA, n.d.).	<i>Solicitation Review Tool (SRT).</i> Pulls in SAM.gov data about solicitations, then compiles the data into a database for use by ML algorithms. An NLP model determines if a solicitation contains compliance language. If a solicitation does not include compliance language, it is marked as non-compliant. Each agency is asked to review its data and validate the SRT predictions. The GSA also conducts random manual reviews monthly.
	<i>Acquisition Analytics.</i> Takes detailed data on transactions and classifies each transaction within the Government-wide category management taxonomy.
	<i>Category Taxonomy Refinement.</i> Uses NLP to extract tokens from product descriptions more accurately to shape intended markets for Product Service Codes (PSCs).
	<i>Contract Acquisition Lifecycle Intelligence (CALI).</i> Streamlines the evaluation of vendor proposals against the solicitation requirements to support the Source Selection process using an automated ML evaluation tool. CALI is currently being trained with sample data from the End User License Agreements (EULAs) under the Multiple Award Schedule (MAS) program.
	<i>Chatbot for Federal Acquisition Community.</i> Enables the GSA FAS NCSC (National Customer Support Center) to streamline the customer experience process and automate documentation of answers to commonly asked questions through public-facing knowledge articles. The end goal is to reduce staffing requirements for NCSC's live chat programs and allow NCSC resources to be dedicated to other



Organization	Tool
	proactive customer service initiatives. Customers still have the option to connect to a live agent if they choose by requesting an agent.
DoD’s Chief Digital and AI Office (CDAO)	<i>Acqbot</i> . Like ChatGPT, generates text to accelerate authorship of acquisition artifacts (e.g., problem statement; Heckman & Heckman, 2023).
Department of the Army	<i>Determination of Responsibility Assistant (DORA)</i> . Pulls information from SAM.gov and the Federal Awardee Performance and Integrity Information Systems (FAPIIS) to deliver relevant information to Contracting Officers (COs; i.e., as specified by inputting a vendor’s DUNS number), rather than having COs access each system separately. Within minutes, the DORA bot sends COs a summary document of the vendor’s responsibility status and results, including screenshots of what the vendor’s file contains (Kanowitz, 2023).

Executive Order 13960, “Promoting the Use of Trustworthy AI in the Federal Government” (Federal Register, 2020) plays a key role in making AI usage data accessible by requiring Federal agencies to document AI use cases; a webpage has been stood up to communicate these to the public to ensure transparency (The White House, n.d.). However, SMEs suggest that these documented use cases may not capture all AI explorations being pursued by Federal Acquisition professionals. For example, the MITRE Corporation is working on an LLM that could function as an interactive Federal Acquisition Regulation (FAR) bot and one that can quickly assess market research. Wolf Stake has also developed prototypes to assist contracting efforts. Government may have been pursuing similar efforts, but no documented cases exist (Wolf Stake, n.d.). Despite the availability of the Federal inventories, the Federal Acquisition community has not catalogued, or shared lessons learned to aid broader AI adoption, integration, and/or training.

A growing trend in government is the development and use of AI chatbots that function like help desks but also assume a small degree of AI discretion (i.e., complete full tasks). Agencies have used these chatbots to perform simple actions such as data summarization and key word searches through records. The DoD’s Defense Information Systems Agency (DISA) has considered a ‘digital concierge’ LLM to assist employees “in all aspects of their job.” DISA estimated that about 80% of the data analyses currently performed by defensive cyber analysts could be automated, freeing up human resources for complex tasks so that “their brains can be applied to those really high-end problems” (Gill, 2023). This concept could also benefit acquisition professionals, automating tasks such as market research.

Future State: 20 Acquisition AI Use Cases for the Year 2035

The Acquisition Management Landscape

Government acquisition managers are intimately familiar with the so-called ‘Iron Triangle’ of cost, schedule, and performance. Successful acquisition management requires balancing the three. Moreover, acquisition managers are deeply aware that perturbations to one element of that triangle have impacts on the other two. For acquisition managers, applying AI fundamentally distorts the dynamics of the Iron Triangle. It becomes possible to perform their tasks faster or less expensively without compromising the other elements. While AI can never fully replace the talent of an acquisition management team, it can be used across the acquisition lifecycle to augment their efforts and improve outcomes.



20 AI Use Cases for Acquisition

Consider a thought experiment. The year is 2035. Use of AI has become ubiquitous. It permeates the daily operations of the acquisition management team. Imagine, then, the work of the Integrated Product Team going through the end-to-end acquisition and contracting lifecycle to procure a new system or service. How does the team use AI?

The following depicts the outcome of such a thought experiment. While a subsequent section of this paper describes those threats and drawbacks in detail, it is important to note for the purposes of this thought experiment that acquisition programs can manage those risks successfully. The potential applications are nearly limitless; however, those identified here may be likely candidates for implementation in the near-term.

1

Building a Request for Information (RFI)

Detailed Description: As the acquisition team begins to fully understand and address the agency’s need, it recognizes the need to engage industry immediately. Knowing that more detailed, face-to-face market research will occur later, the team wants to gauge initial industry interest and “big picture” suggestions. Applying AI, the team rapidly creates an RFI that draws upon historical agency documentation, the acquisition team’s existing products (e.g., briefings to leadership), and notes from several brainstorming sessions. A second AI-enhanced tool reviews the AI-generated RFI to ensure the RFI does not contain controlled unclassified and proprietary information, compartmentalizing different tasks across different AI tools. The acquisition office uses different AI systems to perform different tasks, demonstrating the limitations of a single AI system to accurately execute all tasks. The acquisition team then reviews and enhances the RFI, enabling rapid release to industry.

Impact on Acquisition Outcomes: A process that used to take 6 to 8 weeks is shortened to 10 days, accelerating schedule and enabling the program to perform other, more detailed vendor engagements without adversely impacting schedule.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

2

Industry Constructing Responses to Government

Detailed Description: Each prospective vendor’s AI system examines the Government’s front-end portal on a daily basis, generating a ‘score’ for each newly posted solicitation or request; past performance criteria may influence the scores. This AI tool informs vendors’ capture managers about opportunities, reduces the amount of labor required to find opportunities, and allows industry to better invest scarce business development resources in the most potentially productive lines of effort. After identifying a highly scored RFI (i.e., one that might lead to contract award), vendors use AI applications to respond to the Government’s request. By leveraging existing capability statements, project information, resumes, and other corporate data, the AI system can rapidly produce a high-quality response that includes several key innovations the company has produced for the Government to consider.

Impact on Acquisition Outcomes: Industry only needs to dedicate small amounts of time, resources, and effort to produce high-impact and meaningful RFI responses, including ones that give game-changing technical and strategy recommendations. This has two positive impacts. First, vendors can reduce their overhead rates, thus increasing competitive pricing for future opportunities. Second, vendors become more willing to engage in one-on-one and in-depth market research activities with the Government, as responding to the RFI did not consume their scarce resources.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------



3

Building Market Research Assessments

Detailed Description: The Government acquisition office uses AI tools to examine all industry responses to the RFI. The AI system can search through all the elements of each response, determine overall capabilities of different companies, and provide insight into the quality of the response, as envisioned in the RFI. The AI system’s reports allow the acquisition office to more rapidly assess RFI responses, attach meaning to the different recommendations from industry, and efficiently gain acquisition insights. Specifically, the AI assessments enable the Government to prioritize its actions prior to formal solicitation and weigh different industry recommendations.

Impact on Acquisition Outcomes: Often, Government acquisition offices do not make full use of industry input (e.g., RFI responses) during the market research phase – they do not read the responses in detail and sometimes overlook specific recommendations. The AI-assisted reports on each response focus on critical information, potentially avoiding future risks and accelerating acquisition delivery timelines.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

4

Performing Open-Source Market Intelligence

Detailed Description: The acquisition office uses an AI tool to ask pointed, capability-driven questions. Specifically, the program seeks to understand what potential industry partners that are not part of the typical Government contracting ecosystem may be qualified to bid for a particular contract. The AI produces a summary report for the acquisition office, which uses it to quickly understand the state of practice in industry, including pricing models. The acquisition office then publishes an open RFI on SAM.gov. Once approved by the acquisition office, the AI sends an email to each qualified company’s cognizant business office recommending that the company consider responding to the open RFI.

Impact on Acquisition Outcomes: Acquisition offices always welcome better information about leading-edge capabilities available from industry. Increasing participation by all segments of industry in the pre-solicitation process expands competition to ultimately drive down cost, improve schedule, and enhance vendor performance.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

5

Assessing Risk of Existing Efforts

Detailed Description: Ongoing contracts demonstrate varying degrees of successes and failures across a multitude of cost, schedule, and technical performance metrics. Monthly program management reports, irregular reports and briefings to agency leadership, corrective action reports to contractors, cost and financial status reports delivered under contracts, user feedback on deployed capabilities, and, in extreme cases, legal records contain evidence of such outcomes. The acquisition office tasks a custom AI tool to examine all of these data sources in detail and identify (1) potential areas of risk or failed performance that may be relevant to the current acquisition/contract and (2) potential metrics, program management approaches, and insights that enabled successes in similarly scoped programs. The acquisition office uses the AI-generated report to add and modify contract requirements and internal program management processes.

Impact on Acquisition Outcomes: Learning from the successes and failures of other programs is critical for a successful acquisition. Acquisition offices often write ‘lessons learned’ that frequently become no more than ‘lessons documented’ – acquisition offices simply do not have the time or resources to fully understand and internalize the experiences of other programs. AI can reduce the need for such resources and the effort involved in reviewing the historical record, allowing current acquisitions to connect to valuable lessons from the past. This



increases the likelihood that the acquisition will not repeat past mistakes and will repeat successful approaches.

Primary Improvements:

Cost	Schedule	Performance
------	----------	-------------

Identifying Similar Programs

6

Detailed Description: The program office tasks its AI tool to comb SAM.gov, Government spending data, and contract repositories from across the Federal Government. The tool makes data readily available and informs the acquisition office of other agencies and departments that are acquiring similar (or the same) capabilities. Communication across organizational boundaries and stovepipes can be extremely challenging for Government agencies. The nature of Government operations makes it difficult for agencies to gather current or complete information on the activities of all other departments and agencies. The AI tool can bridge the communication gap when given access to data generated by various agencies and provide detailed insight into similar acquisitions by other organizations.

Impact on Acquisition Outcomes: Using the AI-generated information, acquisition offices can access information about acquisitions conducted by other agencies, learn from these various sources, and potentially leverage their existing contract vehicles. They can use this information primarily to enhance affordability and improve technical performance by leveraging existing capabilities.

Primary Improvements:

Cost	Schedule	Performance
------	----------	-------------

Developing Analyses of Alternatives

7

Detailed Description: Using market research, open source, and internal agency data, the AI system can perform detailed course-of-action analyses for the proposed acquisition. Specifically, it develops informed acquisition strategy alternatives, including contracting options, and delivers technical alternatives for the acquisition team to consider. By pairing these different options with associated risks and opportunities, the AI system gives the acquisition office in-depth insight into different approaches that it can leverage. Moreover, the AI system may identify alternatives that acquisition staff may otherwise have overlooked.

Impact on Acquisition Outcomes: Acquisition offices should make decisions based on the greatest possible amount of information. They can use AI to organize available information, better structuring an acquisition manager's decision space. While AI should not make independent recommendations, it can inform available options from the risk and opportunity perspectives, positively impacting all elements of program success.

Primary Improvements:

Cost	Schedule	Performance
------	----------	-------------

Building Management Plans

8

Detailed Description: Procedurally, acquisition offices expend many resources on building, refining, and reviewing documentation. While the direct utility of many documents may vary by program, the need for such documentation is rarely in question. The law directs programs to develop, and follow, plans that describe their processes and procedures. Often, documentation takes the form of a Systems Engineering Plan, a Project Management Plan, and/or a Test and Evaluation Plan. The program office uses AI to produce these required plans. Starting with past approved plans, augmenting them with the team's copious notes as to how the acquisition office will manage the program, and ingesting recordings of team meetings, the AI system generates partially complete plans. It fills in the areas that it can from the source material and highlights missing material. The acquisition team can use this material as a starting point for developing the required documents, reducing the time and complexity of such undertakings by freeing the team to spend more time on developing new concepts and less time on writing.



Impact on Acquisition Outcomes: High-skilled acquisition labor is best suited to innovating, analyzing, and synthesizing information. Writing routine program documentation can often distract from the acquisition team’s most valuable work. Using AI to help produce some of this documentation frees the acquisition team to focus on more important tasks.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

Creating Work Statements and CDRLs

9

Detailed Description: Many components of contractor requirements, such as the work statement (e.g., Performance Work Statement or Statement of Work) and Contract Data Requirements Lists (CDRLs) include basic, repeated, or high-level requirements that can be derived from existing sources. Additionally, those documents often incorporate requirements, or parts of requirements, from prior contracts. The AI system produces first drafts of these critical contract documents for the program office to consider. In the course of doing this, the AI can determine ‘best of breed’ requirements from different sources of materials and create a superior document. This allows the program office to focus its time on innovations, key requirements, and other differences from status quo.

Impact on Acquisition Outcomes: This AI implementation primarily affects the schedule-to-release time for generating solicitation documentation, enabling the Government to have a better set of documents earlier. Additionally, the quality of the documents is likely to improve, as they will incorporate advances and recommendations across the breadth of requirements.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

Estimating Cost

10

Detailed Description: Using an extensive database of market research, open source data, agency contract data, work statements, and financial data, the AI system builds a comprehensive cost estimate for the acquisition. Each calculated element includes a tailored confidence level and highlights areas with insufficient data to accurately calculate an estimate. The program office uses this AI-generated estimate to start its costing and make well-informed trade-off decisions.

Impact on Acquisition Outcomes: The AI-informed cost estimate more accurately reflects reality and enables the program office to procure the right solution at the right price.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

Building Contract Considerations and Recommendations

11

Detailed Description: Using an extensive library of case law, protest responses, and previous contracts, the AI tool analyzes the Government’s draft solicitation and model contracts. It seeks out potential vulnerabilities, protest grounds, and areas that require clarification(s). Furthermore, the AI tool makes recommendations on how to improve the solicitation and model contract. The CO then uses this report to bolster the solicitation and improve the overall quality of the Government’s contract.

Impact on Acquisition Outcomes: While the improved contract does not directly affect the schedule, cost, or performance of the acquisition, implementing this AI use case does reduce the likelihood of a sustainable protest and increases the clarity of the solicitation. Ultimately, this may reduce the need for re-work and protest-imposed pauses as well as improving the quality of responses.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

Tailoring Training for New Team Members

Detailed Description: The acquisition team brings new technical SMEs and functional



teammates on board throughout the pre-solicitation period. To rapidly bring these new team members up to speed, the AI system produces a custom, high-quality training video for each team member, including the information most relevant to each person’s area of expertise. The system asks each team member to fill out a survey regarding his or her experience with the acquisition, agency, technology, and other factors. It compiles programmatic documentation, slides, presentation videos, and other agency information into a succinct 6-hour training video. This video addresses key elements of the acquisition in question. For members who are new to the agency, the video includes detailed background on the agency and its mission. For those with expertise in a specific domain (e.g., software engineering), the video describes relevant capabilities sought and summaries from market research.

Impact on Acquisition Outcomes: The AI-generated training produces high-quality materials that enable new team members to attain full effectiveness. It does so without burdening existing staff members and allows new staff members to rapidly begin contributing to programmatic outcomes.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

Compiling a Comprehensive Bidders Library

Detailed Description: Offerors require detailed information on existing capabilities, contracts, needs and requirements, priorities and challenges to provide the best quality and risk-informed proposals. Information is key to reducing uncertainty and de-risking contract performance. In general, the more information a potential vendor can apply to its proposal, the better that proposal and the subsequent capability will be. The acquisition office tasks its AI system to curate a library of information for the contractor community. The system includes all relevant information, filters out all sensitive information, flags any conflicting information for program office action, and neatly indexes the content for offeror consumption. The AI system organizes the bidders’ library in a logical flow, allowing vendors to more rapidly understand the information that the library contains.

Impact on Acquisition Outcomes: Creating a comprehensive and impactful bidders library can be challenging for program offices. Moreover, organizing that information into a logical index consumes a considerable amount of time. By using AI, the acquisition office can improve the quality of proposals and the product, service, or /capability to be delivered, positively affecting cost, schedule, and performance during execution.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

Checking Proposals for Compliance

Detailed Description: The AI system compares each received proposal to the Section L or the instructions to offerors that accompanied the solicitation. In this analysis, the AI system scores each proposal to determine if any proposal elements are missing, if the elements contain insufficient information, or if pieces of the proposal fail to meet Government requirements. The AI system generates a report that the contracting office can then use to verify if the proposal(s) in question fails to meet compliance requirements.

Impact on Acquisition Outcomes: This use of AI helps the source selection team perform its tasks faster, reducing the need to analyze deficient proposals and reducing the number of clarifications that the agency must provide to industry.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

Industry Generating Proposal Content

Detailed Description: After receiving the solicitation, offerors use a variety of AI tools to build the basis of their proposals. Specifically, the AI system creates the basis and baseline of

13

14

15



the proposal by comparing the solicitation requirements to other, similar proposals that the company has submitted previously. It also relies on a library of information about each vendor, including capability statements, RFI responses, and other business development and technical implementation material. While the data does not provide enough information for the AI system to formulate a comprehensive proposal, it does provide a starting point. Once the proposal managers and technical team have finalized the proposal, they use a separate AI tool to compare the proposal to the Offeror Instructions and Evaluation Criteria in the Request for Proposals (RFPs). The tool provides recommendations on how to improve the clarity and effectiveness of the proposal.

Impact on Acquisition Outcomes: Commercial companies' use of AI to build proposals enhances the quality and readability of vendor proposals. Furthermore, it allows vendors to focus their limited business development resources on presenting the best approach to solving the most complex technical problems. Ultimately, this ensures that the Government obtains the best innovations and technical solution within the limited time available for bidders to submit their proposals.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

Conducting Initial Proposal Scoring

Detailed Description: Given evaluation criteria and a set of proposals, the acquisition office can use AI to perform initial proposal analysis. Specifically, an AI system can compare proposals to the solicitation requirements and evaluation criteria. The resulting report can identify risk hot spots to which the evaluation team should pay especially close attention and identify areas of each proposal that may have failed to meet requirements. The acquisition office can also use the AI system to inform relevant exchanges with industry, identifying inconsistencies within the proposal or illogical statements.

Impact on Acquisition Outcomes: Source selection is a time-intensive process, often requiring months to explore and evaluate comprehensive offeror-provided information, compare it to evaluation criteria, and determine the Government's best option. Using AI to perform some initial proposal scoring can accelerate this process and make it less laborious. While AI will never replace the role of a source selection evaluation team or technical evaluation team, an AI-based system can help streamline the process and augment such efforts.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

16

Building Source Selection Technical Reports

Detailed Description: The source selection's technical evaluation team has reviewed every proposal. The team has created individual notes on the proposals using templates that focus on risks and opportunities. For consensus events, the team members have created consolidated notes and recorded their conversations. The AI system uses all these sources of information to compile a comprehensive first draft of the technical evaluation report. It compares the source data to the evaluation criteria in the RFP and the content of the proposal. It accurately cites all solicitation and proposal references. It highlights areas of the evaluation that do not make sense or content that is potentially inconsistent with the solicitation or irrelevant. The evaluation team reviews this draft document, adjusts it appropriately, and finalizes it to state the agency's official position. The AI system does not create new evaluation material; it simply condenses the existing evaluation data into a readable technical evaluation report.

Impact on Acquisition Outcomes: Source selection teams spend a tremendous amount of time organizing thoughts, finding references, and documenting their findings. Use of AI accelerates this process and allows the teams to spend more time focusing on the proposal content rather than on how to document it.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

17



18

Compiling Past Performance Data

Detailed Description: The source selection team needs past performance data to evaluate each vendor. For each offeror that submits a proposal, the AI system uses the Contractor Performance Assessment Reporting System (CPARS) and other contract performance data to form comprehensive past performance reports on each company. Furthermore, the AI system automatically reaches out via email to program managers across Government who might potentially have input on the past performance rating. Those program managers then submit their input to the AI database, which uses the inputs to enhance the past performance findings. The AI system compiles a comprehensive past performance report for the source selection evaluation team to consider as it forms its ratings.

Impact on Acquisition Outcomes: AI can consolidate much of the work needed to support an evaluation of past performance. Specifically, the AI system compiles a past performance report without human intervention, enabling the Government to make better use of its scarce acquisition resources and focus only on generating ratings from the gathered information.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

19

Accelerating Transition-In

Detailed Description: Once the contract has been awarded, the awardee uses an AI-based tool as a primary resource to begin work under the contract. Specifically, the contractor relies on the AI tool to help identify the right staff to bring onto the team – both current employees of the company and possible new hires. The AI system uses contract and corporate information to draft initial reports, such as program management plans, risk management plans, monthly reports, and other necessary documents. The system drafts initial kick-off slides for the contractor’s program team to consider. In each case, the AI system seeks to reduce the labor and risk required to complete the project.

Impact on Acquisition Outcomes: Using AI in this way allows contractors to perform their work more efficiently and at a reduced cost. It frees the contractor’s team to focus less on administrative work and more on higher-risk elements of contract performance. This reduces the burden imposed by routine programmatic work and allows the contractor to address higher-risk technical aspects of the work more thoroughly.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------

20

Detecting Performance Assessment Hot-Spots

Detailed Description: The Government program office’s AI system uses programmatic data, including contractor program reviews, contract data requirements list (CDRLs), and meeting notes to create a comprehensive program risk register. This register provides the Government acquisition team with a detailed view of the program’s cost, schedule, and technical performance variables. Moreover, the AI system helps prioritize the higher risk items and draws attention to those areas that require Government action. While the AI system does not replace the need for diligent programmatic oversight, it enhances the Government team’s ability to rapidly and thoroughly process the vast quantities of information generated by the program and its contractors.

Impact on Acquisition Outcomes: It is difficult for acquisition teams to rapidly process the vast quantities of information generated by any given program. Judicious use of AI helps acquisition managers to gain timely insight into the risks and opportunities that permeate their program. This helps the Government team to manage performance while simultaneously enabling the team to focus on the highest priority items.

Primary Improvements:	Cost	Schedule	Performance
-----------------------	------	----------	-------------



While it is important to note that each of these use cases is notional, the examples all demonstrate “the art of the possible” that AI may soon deliver to the acquisition workforce. AI can complement the detailed work that acquisition teams must carry out, augmenting human intelligence, labor, and process flows, and synergistically improving the overall quality and timeliness of the acquisition process. The structure to implement this AI-driven future is already here. Acquisition leaders must learn to embrace how AI will change their functions, the functions of their teams, and the ways in which they build and buy tomorrow’s key technologies.

Before these ideas can become reality within the field, acquisition offices must be open to this new world, be avid learners and early adopters, and ultimately, and most importantly, address risks and challenges. While the opportunities for AI-based improvements are nearly limitless, practitioners must recognize, and mitigate, the associated threats and risks.

AI Threat Landscape – A Manageable Risk Today and Tomorrow

It is clear that AI can transform the processes used by the acquisition community. AI-enhanced applications offer many opportunities to make the acquisition workforce more efficient, decrease individual workloads, and decrease Procurement Action/Administrative Lead Time (PALT). However, while it can bring huge benefits, AI also poses additional risks.

By analogy, the internet offers tremendous advantages to the acquisition process; however, like AI, threat vectors abound across the internet. Universal reliance on the internet has made contractors and the Government increasingly susceptible to attack from any location and by any adversary around the world. Instead of eschewing the internet, though, acquisition leaders mitigate and manage these risks through training, policies, and procedures. This same approach applies to the use of AI in acquisition. As occurred when the internet and computers became tools to assist in acquisition and contract management, use of this new tool in making decisions within the acquisition and contracting process carries similar, if not exponential, national security, financial, and legal implications.

Procedural Risks – If an acquisition office plans to use AI in a source selection or acquisition decision-making, the process must be able to withstand significant legal scrutiny. Any decision that the AI system supports must be substantiated by evidence of the reasoning underlying the decision. Furthermore, as protest decisions are released, AI systems must adapt to incorporate the most recent case law/guidance established by the GAO or the U.S. Courts to supersede any previous guidance. With this understanding, acquisition offices should not use an AI system as the sole decision-maker, but rather as a decision support tool to enhance the knowledge of the human decision-maker. The AI-based system should be considered as a verification tool or an advisor. Acquisition programs must consider additional risks and must proceed cautiously as the Federal Government adopts strategies for implementing AI systems in support of their acquisition processes.

Ethical Dilemmas – With the prevalence of socio-economic programs that ensure small and disadvantaged businesses receive adequate and direct support in the U.S. Government contracting system, training AI models to interpret and weigh the importance of the business type in making an informed and accurate decision presents a unique challenge. For this reason, keeping a “human in the loop” is essential to the introduction and operation of AI in acquisition. All decisions must be fully transparent and traceable. Further, the use of proprietary information to train the model may create significant privacy concerns. For an AI-assisted acquisition, the models must be thoroughly reviewed, vetted, and quarantined to prevent unintended influence.

AI Bias – AI bias refers to the presence of systematic errors in the outputs generated by ML algorithms which can arise from various sources, including the data used to train the algorithms, the design of the algorithms themselves, or the objectives set by the developers (Manyika, 2022). The introduction of biased information into the corpus (the book of knowledge)



or into the training data can undermine the benefits of AI-based technologies, leading to outcomes that compromise the integrity and fairness of a procurement effort. Discrimination caused by AI bias can lead to favoritism towards certain solutions based on irrelevant factors, ultimately reducing competition and increasing costs to the Government. Furthermore, AI bias can cause inaccurate predictions or evaluations, resulting in suboptimal decision-making within the acquisition process. Programs should thoroughly vet both the algorithms and training for bias. Furthermore, following the completion of the task or project, the model should be wiped clean of all data or destroyed. This will prevent unintentional exposure of trade secrets, source selection information, proprietary information, and private information.

Malicious Intent – Knowingly or unknowingly manipulating data to influence an acquisition would be extremely effective and dangerous. Influence operations conducted by adversaries consist of hundreds of thousands of social media profiles with similar hashtags, phrases, or reshared topics. Thus, it would be easy for an adversary, competitor, or interested party to skew research, and possibly an associated AI model or tool, in one direction by corrupting the data being ingested by removing or adding information. With the AI models running on algorithms that generate inferences based upon the data to which they have been exposed, the risk of so-called ‘hallucinations’ arises. Hallucinations in data can be extremely dangerous because they are not based on facts and therefore are simply untrue and inaccurate. Ensuring transparency within the decision-making process means that as training, analysis, and decisions are published and become accessible, the attack landscape becomes larger. More access by nefarious actors creates a larger attack surface and better understanding of how to manipulate, obfuscate, or obtain unauthorized access.

Deep Fakes – AI and deep learning techniques have enabled the creation of exceptionally realistic manipulated digital content, often in the form of images, videos, or audio. This phenomenon, known as “deep fakes,” signifies the application of advanced AI techniques to producing counterfeit content (Taha et al., 2022). Candidate vendors can use deep fakes to fabricate credentials, such as documents, images, or videos, falsely showcasing a team’s expertise, experience, or qualifications, and thus making their proposal appear more attractive to the Government. Vendors competing for the contract, or third parties, could use synthetic media to conduct a virtual presentation, answer questions, and perform a variety of demonstrations. As a result, Government agencies must exercise increased vigilance in verifying the authenticity of information presented in proposals. They should conduct thorough background checks on contractors’ credentials and experiences, strictly verify financial statements, and utilize software or tools to detect deep fake images, videos, or audio.

AI Poisoning – AI poisoning is a type of cyberattack conducted by injecting malicious code or misleading data into the AI training code or dataset. Using AI systems to facilitate Federal Acquisition processes without mitigation control for AI poisoning could potentially lead to biased decision-making, manipulated activities, delays in procurement, and mistrust in systems. To address the risks of AI poisoning in the Federal Acquisition process, it is essential that programs implement a robust process for data validation and verification, following the principles of “Trust but Verify,” and monitor malicious activities in Government AI systems.

Cyber-Enabled Espionage – As AI is integrated into the acquisition process through the use of code, software, and the internet, the attack surface to which these vectors can be exploited increases through cyber-enabled espionage. Obtaining source selection information may give a competitor an illegal advantage or an investor a lead, but, more significantly, adversaries constantly seek to undermine U.S. security systems to obtain designs, intellectual property, or trade secrets at no cost after a company has spent millions of dollars on research and development (Federal Bureau of Investigation, 2022). Writing a spear phishing attack and coding a malicious email may have previously taken one person 3 hours; AI can ingest a



person's biography to create and send an email with malicious code within seconds. This increases the risk that Federal Acquisition staff will encounter exploitation of proprietary, source selection, and national security information.

Mitigating the Threats, Risks, and Biases in AI

While the threats detailed in the previous section may trigger a very reasonable “extreme caution signal” for the risk-averse, acquisition programs can, and should, employ various strategies to mitigate these risks. The National Institute of Standards and Technology (NIST) is developing Federal AI acquisition testing and evaluation capabilities and guidance for other Federal agencies to reference NIST's Artificial Intelligence Risk Management Framework. Within its AI Risk Management Framework (AI RMF 1.0), NIST identified the “characteristics of trustworthy AI systems” as: Valid and Reliable, Safe, Secure and Resilient, Accountable and Transparent, Explainable and Interpretable, Privacy Enhanced, and Fair – with Harmful Bias Managed (NIST, 2024).

Ensuring that each program considers these characteristics as criteria that determine acceptability of any AI system used in Federal Acquisition will be paramount in ensuring the success of the system(s) and optimizing Federal Acquisition processes with AI. The list below summarizes some of the best practices for inspecting and approving AI-based systems. These guidelines provide a framework that helps Federal agencies to ensure their AI systems are not only efficient and effective, but also ethical and in line with necessary Federal regulations.

- **Clearly Defined Requirements:** Use non-ambiguous, precise and succinct language when defining system requirements to limit the risk of misinterpretation and misunderstanding. By clearly defining requirements, programs will ensure that potential vendors have a fundamental understanding of the intent and outcomes the system should strive to achieve.
- **Vendor Assessment:** Evaluate the credibility and previous performance of AI application vendors, especially their performance history with AI deployments. Acquisition professionals can assess the vendors' approach from an ethical and compliance perspective based on current standards, regulations, and best practices.
- **Transparency and Comprehensibility:** Ensure the AI system's decision-making processes are transparent and explainable. This is essential to understanding how the AI system arrives at conclusions and ensuring accountability.
- **Ethical Considerations:** Incorporate ethical guidelines in the acquisition process. This includes considerations pertaining to privacy, data security, fairness, and avoidance of bias.
- **Testing and Validation:** Implement rigorous testing and validation protocols to test the AI system in controlled environments and validate its performance against predefined criteria.
- **Third-Party Audits:** Assess the AI system's compliance with standards and absence of biases through independent audits or certifications from reputable third parties. This will also act as an independent check to mitigate any unintentional biases in an organization's culture or processes.

Federal Acquisition professionals may find themselves leveraging AI to augment or expedite acquisition processes for the procurement of an AI system. In such cases, both sides of the acquisition will present threats and risks, which demand careful attention and oversight of the systems, their logic processes, and their auditing outcomes to ensure that the mitigation strategies are effective against the threats previously described as well as unknown threats yet to emerge.



Recommendations

First and foremost, leaders in the acquisition field must accept the inevitability of AI permeating every facet of the acquisition lifecycle. Like every technical revolution that has preceded it, this technology cannot be wished away. Defense policy makers and acquisition professionals who ignore the promise of AI squander an incredible opportunity while inevitably harming the national security of the United States. Rather than adopting a wait-and-see philosophy or adopting a philosophy of ignoring AI, acquisition leaders must embrace AI in a safe and responsible way. Doing so will open the door to a new acquisition revolution, enabling acquisition teams to procure higher quality capabilities and solutions at lower cost and at a much greater speed.

The natural question is, “what next?”

Recommended Next Steps

Do not categorically deny the use of AI in Acquisition

Acquisition leaders must understand the potential that AI has in enhancing the outcomes of the acquisition process. Unrestrained fear of this technology is undeniably a formula for deterioration in defense and the efficient means of executing the Agency’s missions. Rather, as technologists, acquisition leaders must embrace their role on the front lines of technology adoption and implementation.

Across the acquisition lifecycle, acquisition leaders must identify opportunities to implement AI, use AI to enhance their organizational outcomes, and mitigate the risks that come along with this technology. Education, communication, and prototyping are paramount. Educate acquisition teams, understand the risks, leverage institutional systems for risk management and security, and above all, relentlessly pursue solutions that responsibly harness this power.

Establish a cross-functional team of experts for implementation and oversight, and ensure AI transformational leadership at various levels

Gaining momentum for organization-wide adoption and implementation of AI systems with a top-down strategic focus on executing discrete and shared mission goals and objectives requires an AI Governance Board composed of the various organizational stakeholders. This board should focus on implementing AI at the speed of relevance. AI adoption also requires a dedicated group of professionals who maintain current awareness of AI development, deployment, and maintenance trends and best practices to advise the Governance Board when actionable insights become necessary.

AI Governance

Regardless of an organization’s size or mission, any governance approach must address organization-wide guidance for establishing an approach for acquiring and implementing AI systems, while respecting the nuanced needs of various components within the organization.

Patrick T. Blitgen (2024), PhD, author of *AI for Defense and Intelligence*, offers four strategies that organizations can leverage to begin AI governance but must be tailored to their need and intended use of AI.

Establish clear policies: Policies should be in place to guide the ethical use of AI, the handling of data, and the lifecycle management of AI models. These policies should be regularly reviewed and updated to keep pace with technological advancement.



Enhance transparency: Organizations should strive for transparency in their AI systems, including clear documentation of data sources, model architectures, training procedures, and decision-making processes.

Implement robust oversight mechanisms: This includes establishing dedicated governance bodies or committees, conducting regular audits and reviews, and implementing mechanisms for reporting and addressing issues.

Provide training and education: All stakeholders, from decision makers to end-users, should be educated about the principles and practices of AI governance. This will help ensure that everyone understands their roles and responsibilities and can make informed decisions about using AI.

Task Force

For an AI Governance Board to be effective, it must focus on taking inventory of the organization's current AI capabilities and potential areas of application to responsibly and strategically monitor how to mature the implementation approach. However, with summits on AI advancement occurring almost continuously – at the time of publication (January 2024), summits have spanned the global level (World Economic Forum) and individual Federal entity levels such as the Federal Communications Commission – the Governance Board cannot extract relevant and impactful updates from the flood of daily developments.

The DoD's (2023) AI Hierarchy of Needs (pictured below) serves as a frame of reference to distinguish between the AI Governance Board's and its Task Force's areas of responsibility. The Governance Board would be responsible for setting the foundation of quality data, subsequently building the pyramid upward, and then coordinating internal efforts by the various components charged with continuously implementing enabling actions. The Task Force would be responsible for maintaining situational awareness across the public, private, academic, and international community of external factors that could affect different levels of the pyramid. This Task Force would not function merely as a watch dog but would (1) brief the Governance Board on impact areas and (2) suggest courses of action. This would require an interdisciplinary group of AI SMEs with policy, technical, and behavioral backgrounds.

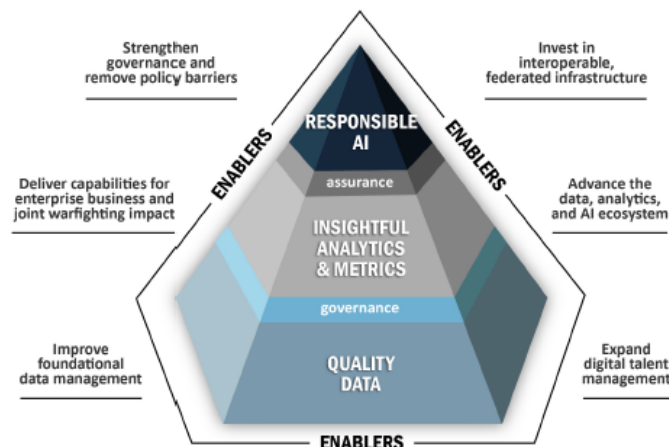


Figure 1.

Establish a holistic AI acquisition framework

Government agencies must create a comprehensive framework encompassing the complete lifecycle of AI systems from inception to implementation and ongoing maintenance. This applies to both commercial off-the-shelf (e.g., IBM Watson, ChatGPT) and in-house-



developed AI products and services (MITRE, 2023) and is in line with the National Security Commission on Artificial Intelligence's (NSCAI's) final report (Schmidt et al., n.d.). The framework should incorporate guidelines for

- A. AI system requirements definition: Clearly define performance, safety, security, and ethical requirements for AI systems.
- B. AI system design and development: Establish best practices for designing and developing AI systems that meet the defined requirements.
- C. AI system testing and evaluation: Develop rigorous testing and evaluation methodologies to ensure AI systems perform as intended and meet safety, security, and ethical standards, in line with the NSCAI's suggestions.
- D. AI system deployment and monitoring: Provide guidance on deploying AI systems in operational environments and monitoring their performance, safety, and security.
- E. AI system maintenance: Establish processes for maintaining and updating AI systems to ensure their continued safety, security, and trustworthiness.
- F. AI maturity model: Build a model of AI integration and maturation within the organization according to its performance goals, mission objectives, and organizational capacity.

Promote transparency and accountability both internally and with external partners

Government organizations can achieve transparency in AI system development and deployment by providing clear documentation of system capabilities and design, objectives, limitations, and potential risks. This must go beyond technical documentation and explain in plain English where in the process AI is being leveraged so that humans either using the AI's output or working with it (i.e., in the loop) can justify and explain the final product/decision. The Government should prioritize safety and ethics to ensure that AI development follows Federal ethical guidelines and safety principles to minimize potential harm to the public and promote safe and effective AI systems that safeguard civil rights. This will enable Government stakeholders to make informed decisions about AI system acquisition and deployment, while mitigating the likelihood that risks will materialize (OSTP, 2023).

Foster collaboration and information sharing

Organizations should consider beginning with small, specialized AI systems to ease the transition into broader AI use and conduct outreach to other Government agencies, industry partners, and academia leading the charge on AI in acquisition efforts. They should use internal and external lessons learned to shape and institutionalize best practices. This aims at enhancing the Federal Acquisition process and promoting the development and adoption of best practices for AI system safety, security, and trustworthiness (OSTP, 2023).

Similarly, Government organizations should leverage proven acquisition techniques to spur a competitive atmosphere and accelerate innovation and acquisition. For example, challenge-based events are known to garner participation from all parts of the private sector (large businesses, small businesses, and academia). High-profile problem spaces, such as the Defense Advanced Research Project Agency's (DARPA's) Grand Challenge, accelerated advancement in autonomous vehicle technologies and incentivized involvement through a \$1 million prize, the winner being a Stanford University team. CDAO is currently undertaking a similar effort, dubbed an AI Bias Bounty, to increase the speed at which the Government acquires leading-edge technology (DoD, n.d.). CDAO could consider elevating this challenge framework to the intergovernmental level by showcasing AI tools that may offer cross-cutting benefits and then sponsoring private sector challenges to meet shared problems in the Federal arena and thus accelerate government-wide access to trusted solutions.



Advocate policy and regulatory updates

Government agencies should regularly evaluate and update their AI systems to ensure they continue to meet the evolving needs of the Federal Acquisition process. By adopting AI-driven solutions, the Government can significantly improve the efficiency, effectiveness, and security of its acquisition processes, ultimately delivering better value for taxpayers and enhancing the overall quality of public services (Hamilton, 2022).

Acquisition leaders are the first line of defense to understanding the gaps and needs in this space. As situations identify various policy and regulatory shortcomings, acquisition leaders must use their platform to communicate and educate their findings.

Understand the Drivers of AI Success

Pilot projects and ideation sessions are critical first steps in AI adoption to identify impactful use areas and obtain organizational user buy-in. After identifying tasks for AI integration, organizations should draw on lessons from current AI adoptions, anticipating challenges and seizing opportunities.

PREREQUISITE: Prioritize data hygiene across datasets. Well-maintained data is crucial in developing effective AI tools and preventing errors or hallucinations.

- **Data Management in AI:** Effective data management is pivotal in AI implementation, serving as the foundation for accurate and reliable AI outputs. High-quality data ensures the AI systems are trained on relevant and comprehensive information to deliver desired outputs, reducing the risk of biases or errors. This has particular importance in environments where AI influences decisions that have significant consequences, such as developing requirements and determining vendor viability. As AI continues to evolve, the role of robust data management becomes even more critical, ensuring AI tools perform not only optimally but also ethically and responsibly.
- **Organizational Change Perspective on AI Implementation:** AI implementation represents a significant shift in organizational dynamics, necessitating a thoughtful change management approach. It requires balancing technological advancements with human factors, ensuring that employees are adequately prepared and supported. This transition involves not just the adoption of new technologies but also a cultural shift towards embracing digital transformation. Organizations successful in AI integration often prioritize continuous learning, adaptability, and employee engagement. Organizations must recognize that the human element is as crucial as the technology and that this approach helps to ensure a smooth and effective integration of AI into existing workflows.
- **Contrasting Culture and AI Adoption:** Organizational cultures resistant to change or lacking technological literacy face greater challenges in integrating AI into their processes. These dynamics highlight the necessity of building a culture that not only embraces technological advancements but also actively prepares for them through training, awareness, and leadership support. The success of AI adoption hinges not just on the technology itself, but on how well it aligns with and is nurtured by the organization's cultural ethos.
- **Establish a maturity model:** Immediately establish a maturity model, overseen by the organization's appropriate governance board, such as the Department of Energy's AI and Technology Office (AITO) and the VA's NAII. This model should guide the transition of AI tools from development to implementation and be aligned with workforce and mission needs. Creating such a model also presents an opportunity to distinguish between AI-generated automation and general software automation, avoiding potential misunderstandings.



- **Identify the tasks that AI could absorb or assist in:** Determine the types of tasks best suited to AI integration, focusing initially on “low-hanging fruit” tasks with low complexity/risk. This approach would allow documentation of organization-specific lessons before progressing to more complex AI-compatible tasks. For example, SMEs suggested using AI to automate the drafting of iterative reports, enabling human experts to focus on authoring initial reports and reviewing subsequent versions.

Challenges: Be aware of these, but they can be overcome by the guidance provided!

- **Acquisition professionals largely lack technical backgrounds:** Since many acquisition professionals are unfamiliar with technology, they require training in AI usage. Developing or hiring in-house AI expertise takes time, so users primarily need to learn how to use AI tools responsibly, confidently, and effectively. Industry adopters of AI who struggled to promote end user adoption, specifically of chatbots, reported, “most of us are using the technology in a suboptimal way, largely because the tech companies gave us poor directions” (“We’re Using AI Chatbots Wrong. Here’s How to Direct Them,” 2023). Users need a clear and direct interface (i.e., prompts) to explain how to pose the right questions to elicit accurately sourced responses.
- **Misconceptions create pushback:** Programs should directly address misconceptions as they arise. It is crucial to emphasize that humans will remain integral in decision-making. AI will enhance, not replace, human judgement. Security is paramount. All AI integrations will occur in secure, Government-approved environments, with necessary security protocols and application programming interfaces (APIs) in place. The bottom line is that programs will not integrate AI into any of their decision-making processes without the appropriate security assurances.
- **Decisions must be traceable:** At critical points, programs require comprehensive documentation to trace how AI analysis informs decisions. For instance, AI could identify the source(s) of its findings, such as the FAR (and the specific section and sub-section) or vendor documents (and specific page and paragraph numbers).

Develop training and education programs for an ‘AI Ready Workforce’

Prior to preparing the Federal workforce to use AI, the organization must understand its employees’ experience and how changes will affect them. This will help to inform communication and education campaigns. *“Part of the challenge is that AI is evolving so quickly that frameworks, tools, and guidance will need to be continuously updated and improved as we learn more”* (AI COE, n.d.).

Organizations must foster a cultural shift towards embracing curiosity and encouraging staff members to ask questions about AI early, often, and repeatedly. The GSA’s AI Guide for Government offers these organizational suggestions to foster responsible and trustworthy AI (AI COE, n.d.):

- **Focus on the root problem.** *Why is AI being considered as a solution? Is it the best option to solve this problem?*
- **Be accountable to the users.** *Establish clear roles and responsibilities and ensure the outcomes of the systems are justifiable to the users who interact with it.*
- **Define and avoid harm.** *Evaluate what possible harms could be and how bias might cause disparate, negative impacts to create mitigation strategies to reduce that possibility.*
- **Monitor the outcomes.** *Are there regular management reviews of changes? Are the systems auditable so that the drivers of incorrect or inequitable outcomes can be identified and fixed?*



Once an organization has established this culture of continuous learning, it must turn its attention to the individuals who comprise its workforce.

First and foremost, acquisition programs must “avoid centralizing AI practitioners and leaders in one unit. AI talent must be accountable to the business needs and therefore should exist across the organization” (AI COE, n.d.). This talent will constitute the bedrock of knowledge to support upskilling of their peers. The AI Guide for Government also recommends that this AI talent “be involved in further talent recruitment, certification, training, and career path development for AI jobs and roles” (AI COE, n.d.).

Continuous Learning and Upskilling

The existing workforce must understand the AI system’s decision-making process so that they can trust its output and provide “explainability” and traceability as they use it to inform their next steps. Developing and retaining AI talent among new staff members will require (AI COE, n.d.):

- Incentives for skill development
- Formal education opportunities
- Optional training, conferences, and exchanges with industry and academia.

The workforce will be learning about a technology that is itself still rapidly evolving; therefore, these support resources should be designed to encourage personnel to provide feedback on ineffective or irrelevant processes.

Adopt a Security-Forward Mindset

To ensure the Federal Acquisition process can fully utilize, operationalize, and conceptualize the use of AI, acquisition offices must ensure they understand the severe consequences of not utilizing AI/ML properly. They must consider procedural risks, ethical dilemmas, malicious intent, deep fakes, AI poisoning, AI bias, and cyber-enabled espionage when adopting this novel technology, which has the potential to change the future of acquisition. With technology, change, and tax dollars come increased scrutiny, oversight, and importance. If acquisition offices do not understand, consider, and address the issues above when utilizing a radically new capability, the public and Congress may have reduced confidence in the Government’s ability to fairly and properly manage taxpayer funds and deliver essential Government functions.

Conclusion

The defense acquisition system rarely gets an opportunity to implement the solutions that it acquires for the warfighter to improve its own ends. AI offers a compelling case for acquisition leaders seeking to enhance their impact on the defense ecosystem. While there are many different use cases for AI systems in the acquisition process workflow, the underlying assumption across the acquisition enterprise must be that these technologies are ubiquitous and essential to mission needs/objectives. As leaders on the front lines of technology adoption, acquisition teams must drive a risk-informed policy of seeking, understanding, and employing these vital technologies. Ultimately, the successful adoption of AI is dependent upon seeing the promise of AI and safely bringing it into the acquisition lifecycle.

References

AI COE. (n.d.). *AI guide for government*. <https://coe.gsa.gov/coe/ai-guide-for-government/print-all/index.html>

AI inventory — tech at GSA. (n.d.).

Blitgen, P. T. (2024). *AI for defense and intelligence*.



- Brynjolfsson, E., Li, D., & Raymond, L. (2023). *Generative AI at work*. <https://doi.org/10.3386/w31161>
- Congressional Research Service. (2023, August 4). *Artificial intelligence: Overview, recent advances, and considerations for the 118th Congress* (CRS Report No. R47644). <https://crsreports.congress.gov/product/pdf/R/R47644>
- DelTek. (2023, August). *Federal artificial intelligence landscape, 2024* [Slide show]. GovWin.
- Department of Labor. (n.d.). *Artificial intelligence use case inventory*.
- Department of State. (2023). *AI inventory – United States*.
- DoD. (n.d.). *CDAO launches first DOD AI bias bounty focused on unknown risks in LLM*. <https://www.defense.gov/News/Releases/Release/Article/3659519/cdao-launches-first-dod-ai-bias-bounty-focused-on-unknown-risks-in-llms/>
- DoD. (2023, June 27). *2023 data, analytics, and artificial intelligence adoption strategy*. https://media.defense.gov/2023/Nov/02/2003333300/-1/-1/1/DOD_DATA_ANALYTICS_AI_ADOPTION_STRATEGY.PDF
- Federal Bureau of Investigation. (2022, July 21). *Combating economic espionage and trade secret theft*. <https://www.fbi.gov/news/testimony/combating-economic-espionage-and-trade-secret-theft>
- Federal Register. (2020, December 8). *Promoting the use of trustworthy artificial intelligence in the federal government*. <https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government>
- Gill, J. (2023, November 7). *DISA wants to use AI as a 'digital concierge' for its workforce*. Breaking Defense. https://breakingdefense.com/2023/11/disa-wants-to-use-ai-as-a-digital-concierge-for-its-workforce/?utm_campaign=BD%20Daily&utm_medium=email&hsmi=281636445&hsenc=p2ANqtz-8NEr73yhYBheB8kT5ouCfqiPUcGVcKOSvLrqqFHUf-c0Q-ZsaquzJsNs0gEOKYEB2eZL9rRljaieEe8xAwKQ7RmNAMXQ&utm_content=281636445&utm_source=hs_email
- Hamilton, T. (2022). *Artificial intelligence acquisition guidebook*. In *Department of the Air Force / MIT Artificial Intelligence Accelerator*. https://aia.mit.edu/wp-content/uploads/2022/02/AI-Acquisition-Guidebook_CAO-14-Feb-2022.pdf
- Harvard University Information Technology. (n.d.). *Generative artificial intelligence (AI)*. <https://huit.harvard.edu/ai>
- Heckman & Heckman. (2023).
- HHS. (2023). *2023 public inventory of AI use cases – HHS.gov*
- IBM. (n.d.). *What is natural language processing?* <https://www.ibm.com/topics/natural-language-processing>
- Kanowitz. (2023).
- Manyika, J. (2022, November 17). *What do we do about the biases in AI?* *Harvard Business Review*. <https://hbr.org/2019/10/what-do-we-do-about-the-biases-in-ai>
- MIT Sloan. (2021, April 21). *Machine learning, explained*. <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>
- MITRE. (2023). *A sensible regulatory framework for AI security*. https://www.mitre.org/sites/default/files/2023-06/PR-23-1943-A-Sensible-Regulatory-Framework-For-AI-Security_0.pdf
- Newman, D. (2024, January 25). *AI rising: 7 technology leaders provide a view from Davos*. *Forbes*. <https://www.forbes.com/sites/danielnewman/2024/01/24/ai-rising-7-technology-leaders-provide-a-view-from-davos/amp/>



- NIST. (2024, January 5). *AI risk management framework*. <https://www.nist.gov/itl/ai-risk-management-framework>
- OSTP. (2023, November 22). *Blueprint for an AI bill of rights*. The White House. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
- Rich, B. (2020, October 27). How AI is changing contracts. *Harvard Business Review*. <https://hbr.org/2018/02/how-ai-is-changing-contracts>
- Schmidt, E., Work, R., Catz, S., Horvitz, E., Chien, S., Jassy, A. ... McFarland, K. (n.d.). *Final report: The National Security Commission on Artificial Intelligence*. <https://cybercemetery.unt.edu/nscai/20211005220330/https://www.nscai.gov/>
- Taha, M., Khudhair, W. M., Khudhur, A. M., Mahmood, O. A., Hammadi, Y. I., Al-Husseinawi, R. S. A., & Aziz, A. (2022). Emerging threat of deep fake: How to identify and prevent it. *International Conference on Future Networks & Distributed Systems*, 6. <https://doi.org/10.1145/3584202.3584300>
- U.S. Department of Treasury. (2024). *Artificial intelligence (AI) use cases*
- USDA. (n.d.). *2023 public inventory of AI use cases - USDA.csv*.
- We're using AI chatbots wrong. Here's how to direct them. (2023, July 20). *The New York Times*. https://www.nytimes.com/2023/07/20/technology/personaltech/ai-chatgpt-bing-directions.html?name=styln-artificial-intelligence@ion=TOP_BANNER&block=storyline_menu_recirc&action=click&pgtype=Article&variant=undefined
- The White House. (n.d.). *Government use of AI*. <https://ai.gov/ai-use-cases/>
- The White House. (2023a, January 24). *National artificial intelligence research resource task force releases final report*. <https://www.whitehouse.gov/ostp/news-updates/2023/01/24/national-artificial-intelligence-research-resource-task-force-releases-final-report/>
- The White House. (2023b, May 23). *Fact sheet: Biden-Harris administration takes new steps to advance responsible artificial intelligence research, development, and deployment*. <https://www.whitehouse.gov/briefing-room/statements-releases/2023/05/23/fact-sheet-biden-harris-administration-takes-new-steps-to-advance-responsible-artificial-intelligence-research-development-and-deployment/>
- The White House. (2023c, October 30). *Executive order on the safe, secure, and trustworthy development and use of artificial intelligence*. <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>
- Wolf Stake. (n.d.). *GPTs for use*. <https://www.wolfstake.org/home/gpts-for-use>
- Young, S. D. (2023). *Proposed memorandum for the heads of executive departments and agencies*. Executive Office of the President. <https://www.whitehouse.gov/wp-content/uploads/2023/11/AI-in-Government-Memo-draft-for-public-review.pdf>





ACQUISITION RESEARCH PROGRAM
DEPARTMENT OF DEFENSE MANAGEMENT
NAVAL POSTGRADUATE SCHOOL
555 DYER ROAD, INGERSOLL HALL
MONTEREY, CA 93943

WWW.ACQUISITIONRESEARCH.NET