ARTIFICIAL
INTELLIGENCE
ACCELERATOR

Advance • Educate • Deliver

DEPARTMENT OF THE AIR FORCE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Artificial Intelligence Acquisition Guidebook

CAO: 14 FEB 2022

300 TECHNOLOGY SQUARE | CAMBRIDGE MA 02139

Foreword

The Department of the Air Force's acquisition professionals must adapt to how Artificial Intelligence/Machine Learning (AI/ML) will continue to impact our defense ecosystem. These technologies will transform the way we develop and procure capability. From our data, to our algorithms, to our warfighters – we must responsibly address implications of the technology, including the way we acquire, develop, implement, and sustain.

This guidebook is a starting point, an opportunity for our acquisition professionals to garner a valuable viewpoint as they deal with a technology that will impact every program office. The topics and lessons covered herein are not exhaustive and instead offered as a catalyst for our acquisition professionals aiming to leverage AI/ML with a healthy and more well-rounded perspective. The AI Accelerator is available and ready to help empower you and your organization to reach an interconnected AI/ML future. Let us make this technology real for our Airmen & Guardians…together!

TUCKER R. HAMILTON, Colonel, USAF
Director, DAF / MIT Artificial Intelligence Accelerator

About the Artificial Intelligence Accelerator

The Department of the Air Force (DAF) signed a cooperative agreement with the Massachusetts Institute of Technology (MIT) to jointly create an Artificial Intelligence Accelerator (AIA) hosted at MIT. The partnership leverages the combined expertise and resources of MIT and the Department of the Air Force. At the AI Accelerator we conduct fundamental research to enable rapid prototyping, risk reduction, and address the ethical considerations of AI to advance both the DAF and society in general. A multidisciplinary team of embedded officers and enlisted Airmen and Guardians join MIT faculty, researchers, and students to tackle some of the most difficult challenges facing our nation, ranging from the technical to the humanitarian. Additionally, the Phantom Fellowship is a specialized AIA program that brings in a group of AI-driven acquisition officers and enlisted members for four months to learn how to manage AI programs and develop policies, processes, and lessons learned to expand that knowledge to the greater community.

In January of 2020, the AI Accelerator launched interdisciplinary projects, involving researchers from MIT Campus, MIT Lincoln Laboratory, and the Department of the Air Force. The projects, which encompass a total of 15 research workstreams, advance AI research in a broad range of areas, including for example weather modeling and visualization, optimization of training schedules, and autonomy for augmenting and amplifying human decision-making.

The AI Accelerator's historic partnership embraces a robust and ongoing dialogue with America's leading academics and technologists to help accomplish an AI-driven future that reflects our nation's values.

To contact the DAF-MIT AI Accelerator and view additional information, please see the AIA website here.

Table of Contents

I. Purpose

Advancements in useable AI throughout the academia and commercial sectors have outpaced the Department of Defense's (DoD) ability to acquire these newest technologies at the speed of relevance.

The purpose of this guidebook is to provide a basic understanding of the AI acquisition lifecycle with respect to data, finance, contracting, and legal considerations. Its aim is to provide deliberate and digestible information to help entry-level to senior program managers navigate the various stages of the AI lifecycle. Additionally, this guidebook provides resources and points of contact that may be helpful for your local projects. While this guidebook includes references to the controlling statutory and policy provisions, it is not a formal policy document. Organizations should consult with legal counsel before entering any agreement to ensure proper adherence to laws and policies.

It is important to note that this AI guidebook is operating under the assumption that the reader has a basic technical understanding of AI, Machine Learning (ML), and data analysis. Prior to working on an AI-related project, the program team should have access to adequate AI education to make effective decisions related to cost, schedule, and performance for the project or program. If you do not have a basic technical understanding of the aforementioned subjects, then please review the information provided in Appendix A - AI Lifecycle and the education opportunities in the Resources section prior to delving into this material.

This guidebook is organized to first preface the state of AI and then delve into some of the considerations for program managers attempting to contract AI/ML capabilities to include contract strategy, testing AI solutions, and ultimately delivering an effective capability to the end-user.

II. Authorities Driving AI

Within the last three years, there has been a marked increase in the proliferation of AI value from both the White House and Congress. This is reflected in key strategy documents to include the National Defense Strategy (NDS) and National Defense Authorization Acts (NDAA). This section is to provide the reader an overview of the current state of national AI policies and initiatives to better support their own project activities. Below is a non-exhaustive list and description of key documents to review before pursuing AI in your organization.

National Guidance (most recent to older):

- National Security Commission on Artificial Intelligence, Mar 2021: *Final Report*
    - Presents recommendations to advance the development of AI and associated technologies that comprehensively address the national security and defense needs of the United States.

- Executive Order 13960, Dec 2020: *Promoting the Use of Trustworthy AI in the Federal Government*
    - Encourages the adoption of AI across the US with nine principles to follow when developing AI. 1) Lawful and respectful of our Nation's values, 2) Purposeful and performance-driven, 3) Accurate, reliable, and effective, 4) Safe, secure, and resilient, 5) Understandable, 6) Responsible and traceable, 7) Regularly monitored, 8) Transparent, and 9) Accountable. This EO does not apply to "AI used in defense or national security systems (as defined in 44 U.S.C. 3552(b)(6) or as determined by the agency)," but they shall adhere to the DoD principles.

- National AI R&D Strategic Plan, Jun 2019:

- Defines the priority areas for Federal investments in AI Research and Development (R&D). Leading AI researchers and research administrators from across the Federal Government developed this plan, along with input from the broader civil society, including from many of America's leading academic research institutions, nonprofit organizations, and private sector technology companies. It lays out eight strategic priorities. 1) Make long-term investments in AI research, 2) Develop effective methods for human-AI collaboration, 3) Understand and address the ethical, legal, and societal implications of AI, 4) Ensure the safety and security of AI systems, 5) Develop shared public datasets and environments for AI training and testing, 6) Measure and evaluate AI technologies through standards and benchmarks, 7) Better understand the national AI R&D workforce needs, and 8) Expand public-private partnerships to accelerate advances in AI.

- Executive Order 13859, Feb 2019: *Maintaining Leadership in Artificial Intelligence*
  - Outlines the basic policy and principles for researching and investing in AI. It enumerates high-level objectives that executive departments and agencies must strive for in this highly technical field, while also stipulating fundamental requirements for protecting the safety, security, privacy, and confidentiality of those involved or affected by the work accomplished. Finally, it establishes several educational grants and programs aimed at increasing the American workforce capability of developing AI.

DoD Guidance (most recent to older):

- Interim National Security Strategic Guidance, 2021:
  - Highlights the need to modernize our military, recognizing that there is a revolution in technology and AI is a technology that must be invested in to retain our scientific and technological edge.

- DoD Ethical Principles for AI, 2020:
  - Defines the five DoD principles that will apply to both combat and non-combat functions and assist the U.S. military in upholding legal, ethical and policy commitments in the field of AI. The DoD's AI ethical principles encompass five major areas...Responsible, Equitable, Traceable, Reliable, and Governable.

- USAF AI Annex to the 2018 DoD AI Strategy, 2019:
  - Acts as the administrative catalyst for the Air Force to ensure AI is a priority when developing budget proposals and planning for the use of funds in future years. It is a framework that implements the Air Force Science and Technology 2030 Strategy and details the principles surrounding the topics of information technology, data, algorithms, people, and partnerships. It highlights five focus areas which connect to the four DoD AI Strategy focus areas. Focus Area 1: Drive down technological barriers to entry. Focus Area 2: Recognize and treat data as a strategic asset. Focus Area 3: Democratize access to artificial intelligence solutions. Focus Area 4: Recruit, develop, upskill, and cultivate our workforce. Focus Area 5: Increase transparency and cooperation with international, government, industry, and academic partners.

- National Defense Strategy, 2018:
  - Describes artificial intelligence within the general topic of rapid technological advancements and how the security environment is affected by these advanced autonomous systems of which the DoD will invest broadly.

- DoD AI Strategy, 2018:

- Recognizes the importance AI will have on future operations citing that it is poised to...impact every corner of the Department. The DoD established the Joint Artificial Intelligence Center (JAIC) to lead the four DoD AI focus areas of: 1) delivering AI-enabled capabilities that address key missions; 2) partnering with leading private sector technology companies, academia, and global allies and partners; 3) cultivating a leading AI workforce; and 4) leading in military ethics and AI safety.

The following National Defense Authorization Acts (NDAA) highlight Congress' increasing awareness of the importance of AI. In these Acts, they have directed DoD agencies to develop, procure, and deliver AI-enabled organizations and programs.

- FY2022 NDAA, P.L. 117-81, Sections 214, 226-228, 232, 247, 803, 833, 1118, 1251, 1509, 1511, 1531
  - Section 214: Adding a new statute (10 U.S.C. § 4010) that requires USD(R&E) to establish a program with the objectives to increase the number of researchers at universities capable of performing science and engineering research responsive to the needs of the DoD.
  - Section 226: Directs DoD to assess potential applications of AI and digital technology to DoD's platforms, processes, and operations, and to establish performance objectives and metrics within 180 days of enactment. These metrics will cover skill gaps within the DoD and AI modernization activities. The skill gaps metrics will identify the qualifications required for military personnel needed for both management and specialist tracks.
  - Section 227: Requires the DoD to modify its AI-focused Joint Common Foundation (JCF) program—a cloud-based AI development and experimentation platform managed by DoD's Joint Artificial Intelligence Center (JAIC)—to ensure that DoD "can more easily contract with leading commercial artificial intelligence companies to support the rapid and efficient" development and deployment of AI.
  - Section 228: Requires the Secretary of Defense, within two years, establish executive education activities on emerging technologies for general and flag officers and senior executive-level civilian leaders that are designed specifically to prepare new general and flag officers and senior executive-level civilian leaders on relevant technologies and how these technologies may be applied to military and business activities in the DoD.
  - Section 232: Creates a pilot program to establish "data repositories" to facilitate the development of AI capabilities for the DoD. This authority allows public and private organizations to access the repositories for the purpose of developing improved AI and ML software capabilities that may be procured to satisfy the DoD's requirements and technology development goals.
  - Section 247: Requires the DoD to submit reports and briefings to Congress on DoD's implementation of recommendations made by the National Security Commission on Artificial Intelligence (NSCAI). The NSCAI's March 2021 report to Congress and the President called for legislation and policy changes as part of a comprehensive defense and national security strategy for "winning the artificial intelligence era."
  - Section 803: Permanently authorizes the DoD's use of Commercial Solutions Openings (CSO) to acquire "innovative" commercial products and services. Using a competition involving a general solicitation and peer review of proposals, CSO competitions target small businesses and commercial companies, but without the restrictions applicable to FAR-based procurements. The Senate Armed Services Committee's report explains that the Air Force and the Defense Innovation Unit have successfully used CSOs with commercial firms.
  - Section 833: Establishes a pilot program to develop and implement "unique acquisition mechanisms for emerging technologies" in order to increase the speed of transition into acquisition program or into operational use. This pilot focuses on projects that support National Defense Strategy priorities, to include space-based assets, personnel and quality of life

improvements, and energy generation and storage. In support of the program, the Senate Armed Services Committee states that "more work is required" to improve DoD acquisition.
    - Section 1118: Within 270 days, the Director of the Office of Personnel Management (OPM) will update one or more occupational series covering Federal Government positions in the fields of software development, software engineering, data science, and data management.
    - Section 1251: Within 270 days, the Under Secretary of Defense for Research and Engineering, in coordination with the Director of the Office of Net Assessment, will develop comparative analysis assessments between the US and China in the areas of: directed energy systems, hypersonics, emerging biotechnologies, quantum science, and cyberspace capabilities.
    - Section 1509: Not later than one year after the date of the enactment of the NDAA, the Commander of United States Cyber Command, the Under Secretary of Defense for Policy, and the Under Secretary of Defense for Intelligence and Security, will conduct an assessment, which may include a war-game or tabletop exercise, of the current and emerging offensive and defensive cyber posture of adversaries of the United States and the current operational assumptions and plans of the Armed Forces for offensive cyber operations during potential crises or conflict. The assessment will recommend possible changes to plans and tactics along with developing future cyber targeting strategies to be employed by US Cyber Command.
    - Section 1511: Not later than 180 days after the date of the enactment of the NDAA, the Chief Information Officer and the Director of Cost Assessment and Program Evaluation (CAPE) of the Department of Defense, in consultation with the Principal Cyber Advisor to the Secretary of Defense and the Chief Information Officers of each of the military departments, will conduct a comparative analysis, to be conducted by the Director of the National Security Agency (NSA) and the Director of the Defense Information Systems Agency. Specifically, the cybersecurity tools, applications, and capabilities offered as options on enterprise software agreements for cloud-based productivity and collaboration suites, such as options offered under the Defense Enterprise Office Solution and Enterprise Software Agreement contracts with the DoD will be assessed. AI and ML capabilities associated with the tools are called out explicitly.
    - Section 1531: Develops a working group and plan for the establishment of a modern information technology infrastructure that supports state of the art tools and modern processes to enable effective and efficient development, testing, fielding, and continuous updating of AI capabilities.
- FY2021 NDAA, P.L. 116-283, Sections 231-236, 241, 801, 808, 1751, 5101-5106
    - Section 231: Alters the requirements of the Department of Defense's (DoD) Joint AI Center (JAIC) to include information describing how the JAIC's efforts contribute to the development of standards in AI through its biannual reports. This takes into account developments made in collaboration with agencies inside and outside DoD and the Intelligence Community (IC) and further requires the JAIC to report on the status of active-duty military personnel assigned to it.
    - Section 232: Amends the organizational structure of the JAIC such that the Director will now report to the Deputy Secretary of Defense instead of the DoD Chief Information Officer (CIO). It places additional emphasis on the acquisition and development of mature AI technologies and gives the Secretary of Defense greater latitude to make decisions about JAIC personnel in research, development, and procurement roles.
    - Section 233: Establishes a board of advisors for the JAIC to provide independent strategic advice and technical expertise to the Secretary and the JAIC Director, conduct long-range studies on AI, and assist Pentagon leadership in developing strategic-level guidance on AI-related hardware procurement and supply-chain issues. The board is to be appointed by the Secretary of Defense and composed of experts from academic or private sector organizations outside DoD.
    - Section 234: Consistent with reform efforts to support the National Defense Strategy, this measure directs the Secretary of Defense to identify a set of at least five use cases for existing AI-enabled systems to support improved management of enterprise acquisition, personnel, audit,

or financial management functions. It also directs the Secretary to pilot technology development and prototyping activities that leverage commercially available technologies and systems to demonstrate new AI-enabled capabilities.
    - Section 235: The Secretary of Defense is to assess whether the Department has the ability, requisite resourcing, and sufficient expertise to ensure that any AI technology acquired by the Department is ethically and responsibly developed. In addition to determining how the Department can most effectively implement ethical AI standards in acquisition processes and supply chains, the Secretary is to provide a briefing of the assessment's results to Congress.
    - Section 236: The Secretary of Defense is permitted to establish a steering committee on emerging technology and national security threats to develop a strategy for the organizational change, concept and capability development, and technology investments needed to maintain the technological superiority of the United States military as outlined in the National Defense Strategy. The committee is to provide recommendations to the Secretary of Defense on the implementation of the strategy, the steps that may be taken to address identified threats, as well as any changes to the Defense Planning Guidance. This committee will sunset October 1, 2024; members are to include the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, the Under Secretary for Intelligence and Security, the Under Secretary for Research and Engineering, the Under Secretary for Personnel and Readiness, the Under Secretary for Acquisition and Sustainment, the DoD CIO, and other officials the deemed appropriate.
    - Section 241: Leverage existing civilian software development and software architecture certification programs to implement coding language proficiency and artificial intelligence competency tests within the DoD.
    - Section 801: Requires a report on integration and interoperability risks associated with using AI across multiple domains.
    - Section 804: Gives Director of the JAIC acquisition authority.
    - Section 1751: Requires DoD to develop guidance on direct hiring processes for AI, data science, and software development personnel.
    - Sections 5101-5106: Requires the President to establish the National Artificial Intelligence Initiative. The outcome from these sections is codified on The National Artificial Intelligence Initiative (NAII) website.

- FY2020 NDAA, P.L. 116-92, Section 256, 260, 800, 1735, 5711
    - Section 256: Requires DoD to create a strategy for educating servicemembers in relevant operational fields on matters relating to AI. The JAIC created the strategy and delivered it to Congress as required. It can be found on the JAIC's website – 2020 DoD AI Training and Education Strategy.
    - Section 260: Requires a biannual report to Congress on the JAIC's activities.
    - Section 800: Enacts use of the Software Acquisition Pathway (SWAP) program codified in DoDI 5000.87. SWAP is designed to enable the timely acquisition of custom software capabilities, reducing time and reporting requirements mandated by the Joint Capabilities Integration and Development System (JCIDS) process. Appendix B provides specific guidance as outlined in Section 800.
    - Section 1735: Extends the National Security Commission on AI (NSCAI) created in the FY19 NDAA to end on October 1, 2021.
    - Section 5711: Requires the Director of National Intelligence to brief Congress on the major initiatives of the intelligence community in AI every year through 2026. They are to describe how they coordinated with the DoD's efforts at the JAIC and Project Maven.

- FY2019 NDAA, P.L. 115-232, Section 238 and 1051

- Section 238: Drives the Secretary of Defense to establish a set of activities within the Department of Defense to coordinate the efforts of the Department to develop, mature, and transition AI technologies into operational use. The establishment of the JAIC is one of the byproducts of this law.
- Section 238: Establishes the JAIC and sets the definition of AI as: (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets. (2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action. (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks. (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task. (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.
- Section 1051: Establishes a National Security Commission on AI (NSCAI) to advance the development of artificial intelligence, machine learning, and associated technologies by the United States to comprehensively address the national security and defense needs of the United States.

III. Managing AI Projects

This section provides a management overview of AI/ML projects for the reader. It first introduces AI requirements definitions and then walks the program manager through contract vehicle preparation. Once the appropriate vehicle has been selected, the section provides considerations for project execution, delivery and sustainment.

Artificial intelligence projects can take many forms. AI projects can be just a piece of a larger program, such as the AI copilot project known as ARTUμ which flew on a U-2 supporting the pilot for sensor employment and tactical navigation. In other projects, AI is the crux of a program, such as Project Maven which aims to provide computer vision and like technologies for a number of applications (e.g., detect objects of interest). These two types of AI projects represent different project lifecycles which are illustrated in Figures 1 and 2 below.

Project Type 1. **AI as Part of the Program**

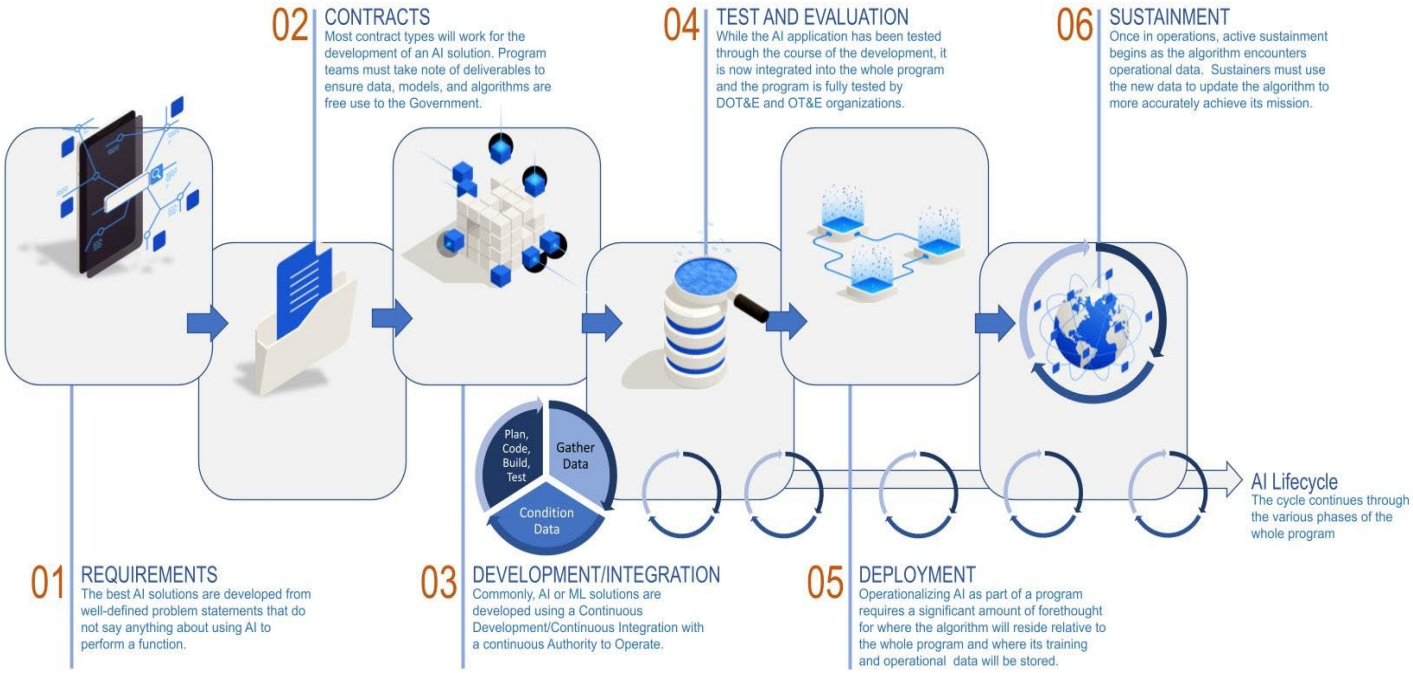

Figure 1: Acquisition lifecycle where AI makes up one component of the program

Project Type 2. **AI is the Program**

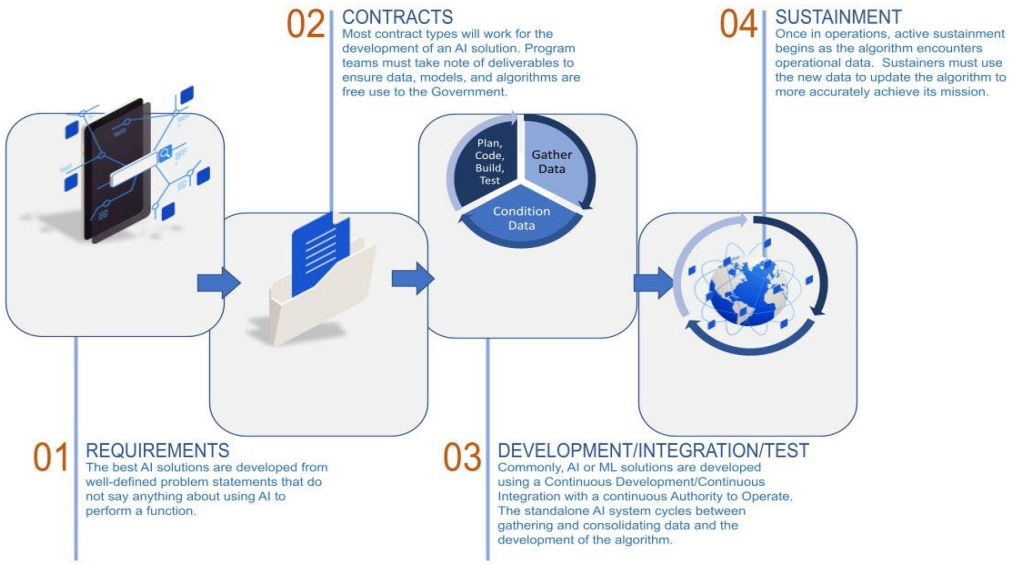

Figure 2: Acquisition lifecycle if the entire program requires AI

There are noticeable similarities between the two figures. In both cases the AI/ML solution will need its scope of work defined normally through a Performance Work Statement (PWS), Statement of Objectives (SOO), or Statement of Work (SOW), depending on the level of requirement specificity. The AI or ML requirements will be placed on a contract (to be further analyzed in later sections of this guidebook); however it is important to note that in Project Type 1, the prime contractor may source the work to a third party, also known as subcontracting.

Another important distinction to make is that while in Project Type 2, the development/integration/test of the AI/ML solution can be accomplished discretely, Project Type 1 will require the AI/ML solution to record its own developmental testing as a sub-system before it is joined with the larger system. Once integrated, the complete system will undergo both developmental and operational testing.

3.1 AI Requirements

The requirements development process is a crucial aspect of DoD acquisition and especially necessary for AI projects. Like any requirement, it is important to start with a robust problem statement that ensures the appropriate problem is being solved with the right solution.

There are many sources that can help a team build problem statements. One resource is the MITRE Innovation Tool Kit (ITK): Problem Framing Canvas, which prepares an integrated product team for a capability acquisition and/or solicitation by helping the user become intrinsically aware of the various base factors causing the capability gap. Appendix C provides additional information on the ITK. Another common practice used to build a problem statement is the "5 Whys" technique, which gets the project team to think about the specific issues and problems.

After using these concepts and ideas from the ITK or another requirements development framework, some requirements may not be fully defined. In this case, there are solicitation tools such as Commercial Solution Offerings (CSO) and Broad Agency Announcements (BAA) that will help the government acquire the relevant capability and refine the problem statement. Once a well-defined problem statement and requirements are created, the program office can move into the contracting phase.

Key Resources for problem framing and/or requirement development:

- MITRE's problem framing toolkit (What, Why, and When to use it)
    - https://itk.mitre.org/toolkit-tools/problem-framing/
- Developing your requirements from a DoD perspective
    - https://acqnotes.com/acqnote/tasks/requirements-development-overview

3.2 Contracts

This section describes a logical, step-by-step process starting at contract strategy preparation through contract execution and sustainment. Prior to starting contract strategy, execution, and administration, we must always conduct market research to develop an accurate depiction of the market environment. Market research can take the form of internet searches, reviewing related publications on the product or service, and reviewing prior vendor assessments from past DoD customers.

To maximize the likelihood of success for an AI project, a program manager must give special consideration to informing the contracting process. As a starting point, program managers can use the following high-level 4-step process to better inform the contracting process.

- Step 1: Contract Strategy Preparation
- Step 2: Explore Existing Contracts
- Step 3: Create a New Contract
- Step 4: Measure Contract Deliverables

Each of these steps is detailed below. Figure 3 provides a visual representation of steps 1 to 3.

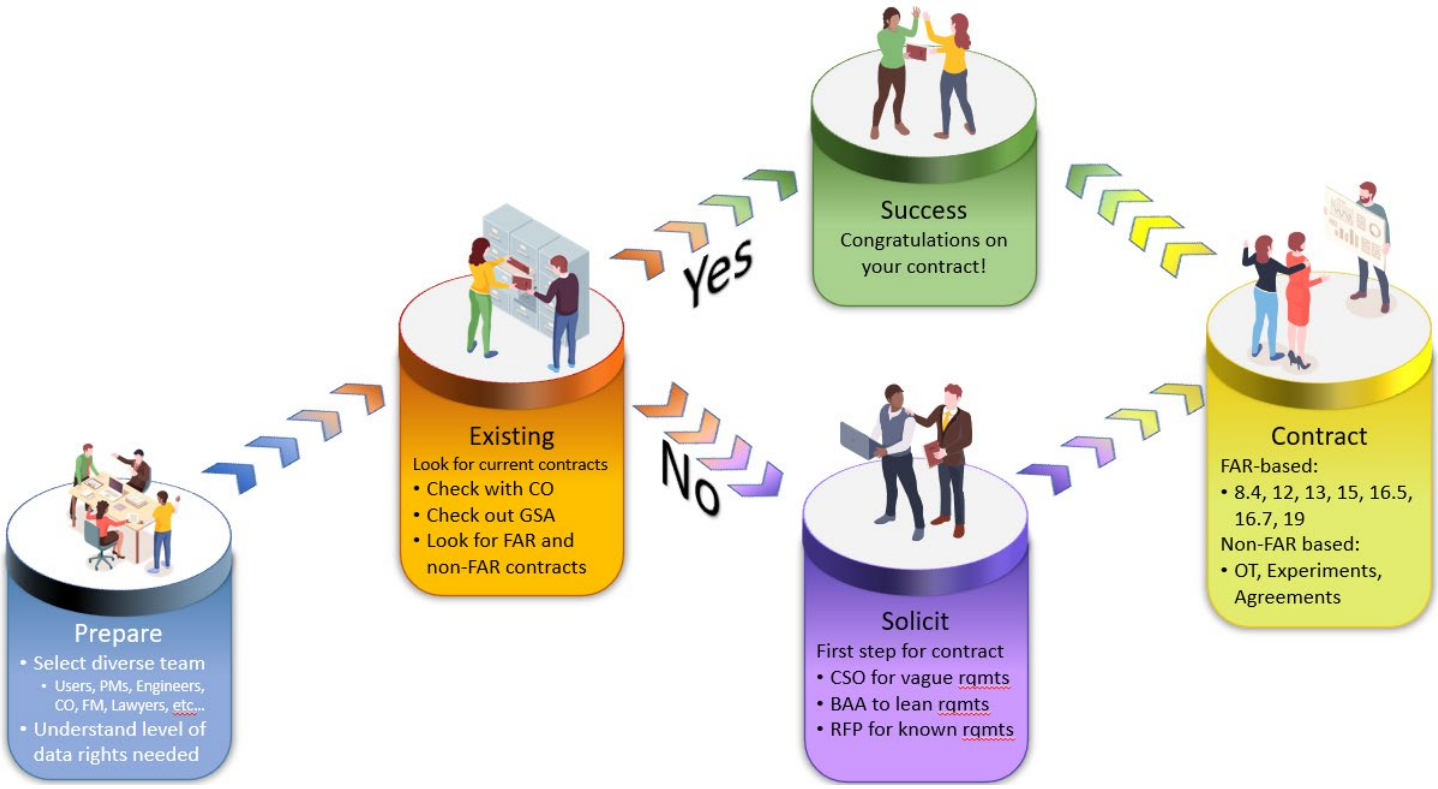

Figure 3: Step-by-step process of contracting process

Step One: Contract Strategy Preparation

Creating a deliberate and thorough contract strategy will positively affect your project's success. Specifically, it is important to begin your contracting strategy preparation with a diverse Integrated Product Team (IPT) including, but not limited to, users, program managers, engineers, data scientists, finance members, contracting officers, and lawyers. Contracting lawyers are privy to relevant contracting procedures, data rights, and intellectual property information not widely available to the general IPT population.

In addition to building a strong and diverse IPT, a key part of the contract strategy preparation is developing your project's data rights' needs because data is essential to AI projects. Data rights refer to the Government's nonexclusive license rights in two categories of valuable intellectual property: "technical data" and "computer software" delivered by contractors under civilian agency and DoD contracts. FAR- and non-FAR-based contracts have vastly different procedures and regulations on data rights. FAR-based contracts are regulated by 10 U.S.C. § 2320-21 and DFARS 227 and DFARS 252.227 for data rights and the Bayh-Dole Act (35 U.S.C. § 200-212) for patents and are described more below. Non-FAR-based contracts are not regulated by most procurement statutes and regulations, thus allowing the government to negotiate beneficial data rights within the contract or agreement. Your contracting officer and lawyer will be able to provide guidance to your team prior to choosing a specific contract type.

AI presents unique considerations regarding data rights and IP. Non-FAR-based contracts, specifically Other Transaction Agreements, offer negotiation-based rather than strictly regulatory-based data rights framework. While this flexibility can attract AI firms that may avoid traditional government data rights/IP frameworks, and meet an agency's needs more precisely, such negotiation is nuanced. Thus, the acquisition team must develop the IP strategy with care. Program managers should lean heavily on subject matter experts.

DFARS 227 outlines patent, copyright, and data regulations for FAR-based contracts, while DFARS 252.227 outlines rights in technical data for non-commercial items. For FAR-based contracts, the standard data rights invoked are (from most government rights to the least):

Unlimited Rights: Allows the government to, in any manner or purpose, use, disclose, reproduce, derive work, distribute, and publicly display the data. Unlimited data rights apply when:

- Data first produced in performance of a contract
- Form, fit, and function data
- Data derived from manuals or instructional and training material for installation, or operation and maintenance
- All other data delivered under contract, other than limited or restricted computer software

Government Purpose Rights (GPR): Provides a middle path between limited and unlimited data rights that allows the contractor to use the IP commercially, while the government maintains usage rights. GPR are used when both government and private funds are leveraged in development efforts.

Limited Rights: Enables the contractor to protect certain limited rights data and restricted computer software from government use. If the contractor has privately developed and funded the technology, then the contractor may limit the government's use of technical data. The government can use the data within the government, but cannot release data outside of the government unless an exception is met. Restricted data rights are a category of limited data rights where the government is only able to run software on one computer at a time and make minimum copies or modifications to the software.

Specifically Negotiated License Rights: Allows the government and contractor to modify the standard license rights to obtain, relinquish, or increase the government's rights in technical data.

The government can negotiate rights within FAR-based contracts; however, deviations are not common. Ultimately, both regulations (DFARS 227 and 252.227) are quite restrictive and make garnering data and IP difficult for AI projects. Non-FAR-based contracts offer the government much greater latitude to mold beneficial data rights and IP, especially in the context of AI.

*TIP: An important consideration during contract preparation is to determine how the project wants the data formatted and transmitted. Data exists in three states: in use, at rest, and in transit, and can be formatted in a multitude of ways. To enable rapid acquisition, development, and operational deployment of an AI capability, all data returned to the government should be AI-ready and, in some cases, even further conditioned to enable model training (click here for an in-depth description of preferred technical data formats).
For more information on data rights, please see the following resources:

- AcqNotes: https://acqnotes.com/acqnote/careerfields/data-rights
- AIDA: https://aida.mitre.org/demystifying-dod/ip/

Step Two: Find an Existing Contract

The quickest way to get on contract is to find one already awarded and captures your specific project's technical and financial scope. The list of FAR and non-FAR contract types described in upcoming Step Three is not exhaustive, but highlights the most practical contract types for an AI project. Additionally, each specific project scenario could potentially call for a different contracting vehicle, which have other considerations that could apply to unique aspects of a program.

Currently, there is not a functioning repository to search for current contracts with AI/ML scope within the DoD. The General Services Administration (GSA) has information on awarded contracts with scope; however, it may not have all the information required to find the right contract for your project. As such, it is crucial to involve a contracting team that can use their experience and contracting tools to conduct research in conjunction with the Program Manager.

If you found an existing contract that meets your needs, skip to Step Four.

Step Three: Create a New Contract

If there are not any available and previously awarded contracts that meet your project's specific scope, then you will have to create a new contract. Your team must first identify what solicitation method they will use to garner proposals. Detailed below are common solicitation methods used for requirements that generate AI solutions—Data Challenges, Defense Commercial Solutions Openings, Broad Agency Announcements, and Request for Proposals. Then your team will use these proposals to award FAR-based or non-FAR-based contracts, which are highlighted after the Solicitations section.

Data Challenges

Background: While not a solicitation method, project teams can use data challenges within solicitation methods to encourage academic, commercial, and government collaboration. The AI community has traditionally leveraged data challenges to engage academia and industry for decades by using available data sets to solve complex problems. The broad collaboration created by data challenges has driven fundamental breakthroughs in data science for many years. Organizations such as DIU and the AIA have begun to use challenges to drive

capability development using unclassified DoD releasable data sets. Program offices can leverage existing challenges or create their own to find solutions to their problem, The links below provide examples of previous data challenges:

- Jan 21: SEVIR Dataset Challenge
- Feb/Mar 22: MIT Data Center Challenge
- 2020-Present: RF Challenge at MIT
- 2020-Present: AI Accelerator Maneuver Identification Challenge
- Nov 22: DIU xView3
- 2004-Present: DARPA Prize Challenges

AI Context: Program offices should include language to leverage data challenges in the solicitation process, which can build a collaborative information environment in which data is shared and mutual efficiencies are realized. In application, a program office can provide vendors with a representative dataset and solicit a model response within weeks. Upon submission, the program office may evaluate the vendor's model to determine its value. Such inclusion will trigger the search for useful data prior to contract award rather than later in the acquisition life cycle. For example, a program office could include language such as:

> "Responders looking to highlight the relevance of their AI capabilities to the Air Force should consider demonstrating their AI innovations on challenge problems described in DAF/MIT AI Accelerator Challenges."

10 U.S.C. § 2374a provides authority for providing monetary or non-monetary prizes for advanced technology achievements such as winning challenges. Prize challenges may be used in conjunction with other authorities to acquire, support, or stimulate research, technology development, or prototype projects carried out under the authorities discussed in the Solicitations section below. Prizes awarded under this authority satisfy competition requirements under the Competition in Contracting Act (CICA).

Solicitations

The following sections of solicitation types are suggested offerings of what is available. This list is not exhaustive but rather highlights solutions typically used.

Defense Commercial Solutions Opening (CSO)

Background: The use of a CSO is permanently authorized by section 803 of the FY2022 NDAA. The purpose of a CSO is to obtain new and innovative capabilities or solutions that can fill requirements, capability gaps, and/or provide technological advances. Also, CSOs can be used to procure commercial products and services that directly meet project/program requirements.

CSOs enable government and industry collaboration on vaguely defined problem sets. Per GSA Manual (GSAM) Part 571, the solicitor shall use the Procurement Innovation Resource Center (PIRC) template for submission. CSOs are a merit-based source selection and solicitation strategy available for the DoD. Each proposal submitted through the CSO process is evaluated by its individual merit rather than on a comparative basis. Additionally, CSOs can use both FAR and non-FAR based contract options allowing for a more versatile contract strategy that hopefully meets the needs of a variety of project and program use cases.

AI Context: CSOs enable rapid solicitation and selection of innovative commercial solutions. Often, the government uses CSOs for IT products and services, commercial technology studies, and commercial

technology maturation. CSOs allow the government to provide general problem descriptions rather than specifically defined requirements to industry, quickly select a partner, and garner a solution. This solicitation method is more flexible than traditional competition methods better aligned with industry practice and encourages AI firms to more creatively solve problems.

Broad Agency Announcement

Background: BAAs are solicitation methods used to obtain proposals for basic and applied research and development used to advance or evaluate innovative technologies. In FAR 35.01, BAAs can be used to award science and technology proposals for: basic research (budget activity 6.1), applied research (budget activity 6.2) advanced technology development (budget activity 6.3), and advanced component development and prototypes (budget activity 6.4). BAAs are generally "open" and proposals from vendors are accepted for a specified period and can be used to award both FAR and non-FAR based contracts.

Two widely used BAA processes are Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR). SBIR is a competitive program that motivates small businesses to spur Federal Research and Development (R&D). STTR is a program to facilitate small business and non-profit R&D. Both processes provide a streamlined avenue for government projects and industry to access additional federal and private funding (click here for additional SBIR opportunities and information).

AI Context: Projects use BAA, SBIR, and STTR for technology maturation, studies, and prototyping. Similar to CSOs, BAAs allow a program to pose problems as an "open topic" to which industry can provide solutions. This allows the government to have a relatively less refined requirement. When the primary purpose of the contract is to advance scientific and technical knowledge and apply that knowledge to the extent necessary to achieve agency national goals, contracting should be approached differently than when the purpose is to procure services or supplies. Like many R&D efforts, contracts for AI-enabled capabilities are directed toward objectives for which the work or methods cannot be precisely described in advance. Thus, problem-focused solicitation methods like CSOs and BAAs or SBIR open topics are generally preferrable to precise performance specifications as typically seen in request for proposals.

Request for Proposals

Background: A Request for Proposal (RFP) is a solicitation method used in negotiated acquisitions to communicate requirements to industry and enable their proposals. In FAR 15.203, RFPs require, at a minimum, that the government communicate its requirement, anticipated contract terms and conditions, and sometimes evaluation criteria depending on the contract strategy. As a note, for commercial acquisition of products and services, which may be needed to supplement the development effort, PMs and Contracting Officers should follow the requirements set forth in FAR Part 12.

AI Context: RFPs require well-defined requirements and/or problem statements, which are not always feasible for problems that may require an AI solution.

Other Transactions

Background: Other transactions can be solicited in a variety of methods, including each of the above methods. Because the Competition in Contracting Act and FAR do not apply to OTs, solicitations can take on flexible formats including demonstrations, hackathons, challenges, pitch days, or more formal processes to include down-selecting and fly-offs. See OT Guide.

AI Context: Given the flexibility of the OT statutes, an agency can engage with industry in a more flexible, commercial-like manner that can result in shorter procurement lead times than traditional contract competition.

Contracts

After a requirement has been solicited to industry, there are a number of contracting options that can be used to engage performers; these include both FAR-based and non-FAR-based contracts. Figure 4 provides an overview of the various contracting mechanisms and their governance structure.

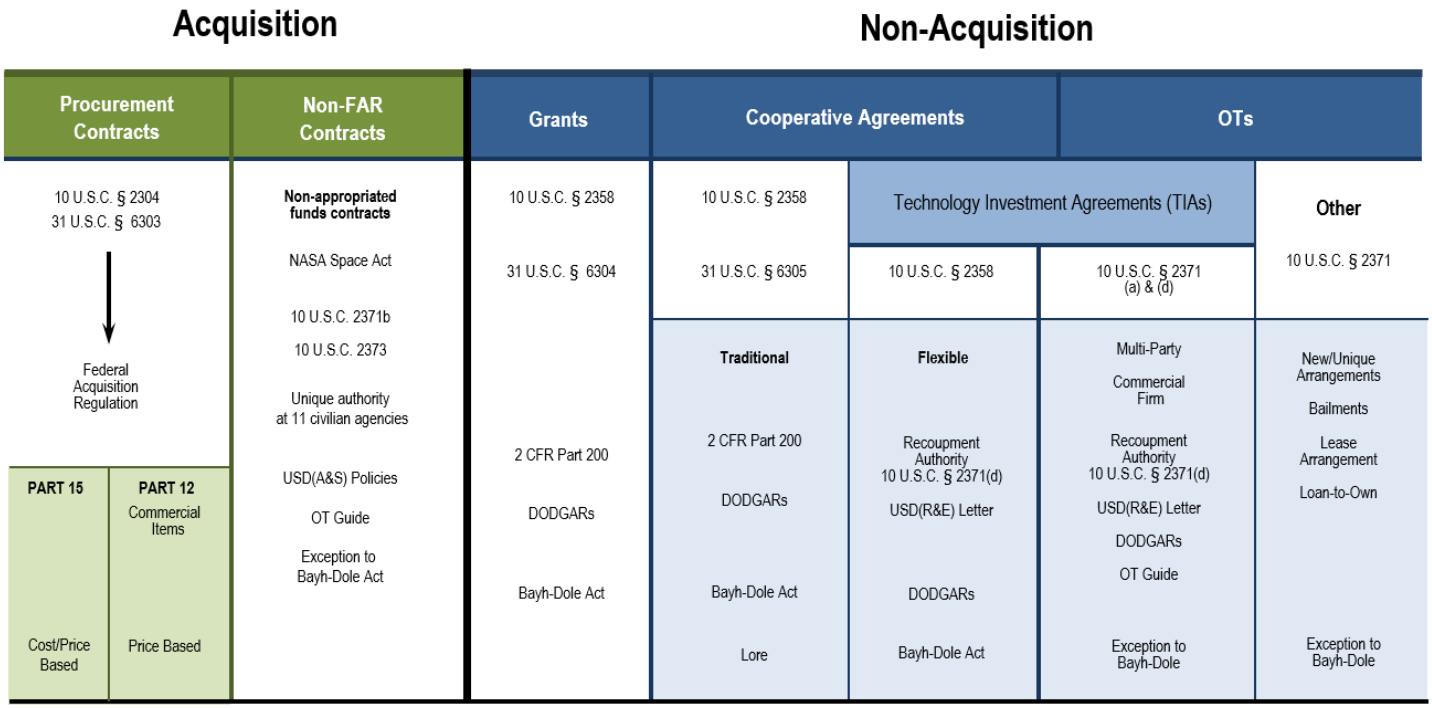

Figure 4: Comparison of acquisition and non-acquisition contracting mechanisms

The following section provides further background on select contracting options; specifically those that are commonly used where AI technology is part of the deliverable. It is broken down into FAR based considerations and non-FAR based options.

FAR Contracts

Federal Acquisition Regulations, are the better-known authorities for DoD contracts and solicitations, with the specific FAR citation listed within the green section of Figure 4. When a Contracting Officer references a "FAR-based contract," they mean a contract under one of these authorities, subject to preset regulations and often pre-determined clauses.

Federal Supply Schedules (FAR 8.4)

Background: Federal Supply Schedules provide the government easy access to simplified contracting processes for procuring commercial supplies and services with volume buying. Schedules contain previously negotiated fixed prices for both products and services. GSA publishes those proposed and negotiated prices on the GSA Schedules e-Library. Under Federal Supply Schedules and GSA, the program manager can execute Task Orders, Delivery Orders, and Schedule Blanket Purchase Agreements (BPA).

AI Context: Use FAR 8.4 for AI services and products that are available commercially and for volume procurement. Also, use FAR 8.4 to contract with any entity that has pre-negotiated prices on GSA for services such as Computer Vision as a Service, Annotation as a Service, licenses, or even tools that could allow the government to create AI/ML applications in house.

Commercial Items (FAR 12)

Background: Use FAR 12's streamlined procedures to procure Commercial and Commercial Off-the-Shelf (COTS) supplies and services (per FAR 2.1). FAR 12 provides streamlined, commercial-like processes when goods and services meet the FAR Commercial and COTS definitions. These procedures can provide an efficient alternative to FAR contracting processes such as FAR 15.

AI Context: Common applications for FAR 12 include products and services for IT, cyber, cloud, software, telecommunication, and various alternative Commercial/COTS technologies. Use FAR 12 for situations where industry and a program office or unit are using AI and are attempting to take that same technology or service and use it for the DoD. Fortunately, dual use technologies, technologies with both commercial and military use cases, are becoming more prevalent.

Simplified Acquisition (FAR 13)

Background: Simplified Acquisition, as it is outlined in FAR Part 13, provides a streamlined process to acquire both supplies and services for research and development, and commercial items as long as those supplies/services do not exceed the simplified acquisition threshold for non-commercial items. FAR 13.5 allows simplified acquisition procedures for supplies and services greater than the simplified acquisition threshold but not greater than $7.5M (This threshold is subject to change so consult with your contracting office before you proceed forward). Simplified acquisition procedures allow contracting officers additional procedural discretion and flexibility for services and supplies within the financial thresholds that eases the administrative costs and burden for both government and industry. FAR 13 can use Simplified Acquisition BPAs, Purchase Orders, and Micro Purchases.

AI Context: Use FAR 13 for repeated orders of both commercial and non-commercial products and services beneath the simplified acquisition threshold. Practical use cases could be engineering services, studies, and other common AI services from a commercial and non-commercial standpoint.

Contracting by Negotiation (FAR 15)

Background: FAR 15 procedures are used for both competitive and non-competitive, open-market acquisitions that exceed the Simplified Acquisition Threshold. Open market services and products are those that are not available from required sources of supply, such as those outlined in FAR 8. FAR 15 delineates rules for negotiating, evaluating and awarding these contracts. FAR 15 requires unique cost and pricing procedures due to the lack of commercial data.

AI Context: Use FAR 15 for technology that is not commercially available or dual use and can be cumbersome and cost-prohibitive for many companies. It generally attracts large national security contractors that provide unique solutions for the government. Often, AI is considered a dual use technology that requires more expeditious contracting; however, certain use cases call for a unique national security capability, which is suited for FAR 15 contracts.

Modular Contracting (FAR 39)

Background: Per 41 U.S.C. § 2308, modular contracting is the favored method for acquiring major software information technology systems. FAR 39 can also be used for non-major systems as well. FAR 39 is designed to reduce program risk and provide performance incentives to the contractor while meeting the government's time-sensitive requirements for rapidly evolving technology. Best practices for Agile and Modular contracting can be accessed in the U.S. Government Accountability Office's Agile Assessment Guide.

AI Context: Modular contracting is an excellent tool for FAR based major AI acquisition. FAR 39 enables rapid contracting procedure that can keep pace with the AI technology evolution.

Indefinite Delivery Indefinite Quantity (FAR 16.5)

Background: Indefinite Delivery Indefinite Quantity (IDIQ) allows customers to order from indefinite delivery contracts originally let from the same agency and from other agencies. IDIQs can take the form of a Single Award IDIQ Contract or a Multiple Award IDIQ Contract. Government-wide Acquisition Contracts (GWACs) and Multi-Agency Contracts (MACs) are examples of available IDIQ contracts.

A GWAC can be a task-order or delivery-order for information technology (IT) established by a single agency and operated by an executive agent designated by the Office or Management and Budget or under a delegation of procurement authority from GSA. A MAC is a task-order or delivery-order contract created by one agency for use by alternative agencies to procure supplies or services in accordance with the Economy Act.

AI Context: Multiple Award IDIQs can be cumbersome to initiate; however, if there is a strong demand signal for repeated acquisitions, then this contracting method can benefit the government. GSA Oasis is an example of a MAC that is available for government stakeholders. IDIQs can offer easy access to products and services at pre-negotiated prices.

Basic Ordering Agreement (FAR 16.7)

Background: A Basic Ordering Agreement (BOA) is not a contract, instead it is a written document containing pre-negotiated contract clauses that could be applicable to future procurements between the parties in the agreement. A BOA includes the description of the product and/or service, method for determining pricing, issuing, and delivery of future orders.

AI Context: BOAs offer a flexible agreement between government and vendors for uncertain product and service requirements. BOAs can allow additional parties to join the agreement. If a BOA has a similar scope and contractor that can fill certain AI requirements, it is a good and efficient tool to get a party on contract for any number of AI use cases including both supplies and services.

Small Business (FAR 19)

Background: There are two categories of a FAR 19 contract: Small Business Set-Aside and Direct 8(a).Set-asides are a way to reserve either the entire acquisition or a portion of the acquisition for a small business. Acquisitions below the Simplified Acquisition Threshold must be set aside for small businesses unless the contracting officer determines an exception applies. Acquisitions over the threshold must set aside work if there are two or more small businesses that can complete it. Direct 8(a) awards allow the government to use sole

source procedures to award to a single contractor within the 8(a) business development program if certain conditions are met.

AI Context: FAR 19 can be used for supplies and services that small and disadvantaged businesses can support. FAR 19 also allows the government to go sole source to certain contracts. For AI, FAR 19 can provide a quick contracting method to acquire certain products and services.

Non-FAR Contracts

"Non-FAR" points to statutory, rather than regulatory authorities. These are listed in the blue shaded section of Figure 4. Non-FAR-based contracts, or statutory agreements such as Other Transactions, may be less familiar but offer unique alternatives versus traditional FAR based contracts. These options may be better suited for a small business or startup because they don't contain standard clauses and can be more freely structured.

Other Transactions

Background (DAU/MITRE): Other Transactions (OT) (click here for guide) enable DoD access to innovative research and development. FAR/DFARS mandated statutes and regulations such as Cost Accounting System (CAS) compliance, data rights procedures, and intellectual property rights limitations are not applicable. The DoD can negotiate relevant terms and conditions that may be more in line with the DoD's interests and commercial business practices. For example, through an OT the government and contractor can negotiate funding arrangements, payment milestones, data rights, and length of agreement.

OTs are designed to decrease barriers to entry for non-traditional defense contractors; however, recently OTs have been made available for traditional contractors who satisfy certain conditions (see in-depth description of traditional versus non-traditional contractors here). As outlined in 10 U.S.C. § 2302, a non-traditional contractor is a business entity that does not have or has not had, within one year, a full Cost Accounting Standards (CAS) covered contract or subcontract. Furthermore, a company qualifies as a non-traditional defense contractor if *any* of the below apply:

- The small business is exempt from CAS requirements
- The company exclusively performs contracts under commercial procedures
- Exclusively perform under firm-fixed-price (FFP) contracts with adequate compensation
- Performed less than $50M in CAS efforts during the previous cost accounting period

The DoD has the authority to use Research OTs (10 U.S.C. § 2371), Prototype OTs (10 U.S.C. § 2371b), and Production OTs (10 U.S.C. § 2371b(f)). Use Research OTs for basic, applied, and advanced research projects for dual-use application. Research OTs do not include the ability to transition to follow-on production contracts or transactions. Consider Prototype OTs for research, development, and prototyping activities. Use a Prototype OT to procure a reasonable number of prototypes for the government to field and test to determine if it would like to procure in quantity. Furthermore, a Prototype OT provides a streamlined effort for a non-competitive follow-on Production OT or FAR contract if the initial Prototype OT was "competed" to the maximum extent possible. OT competition does not have to follow the Competition in Contracting (CICA) Act.

Overall, OTs offer a program manager a streamlined way to acquire AI tools, and, in order to award OTs, a contracting officer must have an additional Agreement Officer authority. OT Consortia offers interesting benefits to program offices seeking partnerships (click here for OT Consortium list). OT Consortia are made up of a government sponsor and contracting office, Consortium Manager, and the Consortium (industry and academia) itself, which is created based on a specific technology focus area (i.e., Space, Weapons of Mass

Destruction, Hypersonics). OT Consortia offers a program increased collaboration and streamlined administrative processes, and a way for programs whose contracting officer does not have the proper authority.

AI Context: OTs offer great opportunities for unique contractors, agile contracting, data rights, and intellectual property. OTs offer flexibility to shape future contracts according to future requirements (FAR or non-FAR based contracts). OTs offer the government flexibility to sidestep cumbersome administrative procedures, such as adding additional vendors, while contracting at the pace of industry. Critically, OTs are inherently flexible and permit robust collaboration and negotiation that is difficult to achieve under the FAR.

Procurement for Experimental Purposes (10 U.S.C. § 2373)

Background: Procurement for Experimental Purposes (often referred to as "2373") allows the government to procure quantities of a product necessary for experimentation, technical evaluation, operational use case assessment, or to maintain operational capability. It allows acquisition in nine (9) areas: ordnance, signal, chemical activity, transportation, energy, medical, spaceflight, aeronautical supplies, and telecommunications. 2373 can be both competitive and noncompetitive and considering it is a non-FAR-based contract, FAR and DFARS are not applicable. However, a Determination and Finding (D&F) identifying certain information including item description, method of test and quantity to be tested is required.

AI Context: For AI, the government has a large amount of flexibility with 2373. AI projects can fall into any one of the nine technology domains which makes 2373 a particularly good contracting option for AI for medical, weapons, signal, or telecommunications. In conjunction with OTs, 2373 enables rapid transition of technologies into fielded systems for test and evaluation. Furthermore, a 2373 contract is not tied to FAR or DFARS data and IP regulations, may or may not be competed, and only requires a D&F to award.

R&D Agreements

Background: There are three types of R&D agreements: a Cooperative Research and Development Agreement (CRADA), Partnership Intermediary Agreement (PIA), and Technology Investment Agreement (TIA). A CRADA, see 15 U.S.C. § 3710a, allows federal labs to enter into agreements with other federal agencies, government, academia, or industry to commercialize intellectual property and/or research derived at the labs. A PIA, see 15 U.S.C. § 3715, can be a contract, agreement, or memorandum of understanding with a non-profit to include academia and/or industry to expedite technology transfer and licensing. A TIA is an agreement executed in accordance with the DoD Grant and Agreement Regulations (DoDGAR), see 32 CFR Part 21, 10 U.S.C § 2358, and 10 U.S.C. § 2371. They support commercial involvement in defense and national security research. TIAs are applicable when objectives are unlikely or cannot be met through other agreement or contract methods.

AI Context: R&D agreements allow the government to work with academia, labs, industry, and non-profits for various levels of R&D. For AI, if the requirement is to develop fundamental research and drive forward the science, then it is possible that this vehicle is the right path for the project.

See Table 1 for a snapshot of the contracting options described above:

Table 1 - FAR vs non-FAR options

	Contract Option		Links	AI Context
FAR Based	Federal Supply Schedule		FAR 8.4, DAU, GSA eLibrary	Use for products and services available commercially and for volume procurement
	Commercial Items		FAR 12, DAU, Commercial Item Guide	Use for COTS products and services for IT, cyber, cloud, software, telecommunications, etc.
	Simplified Acquisition		FAR 13, DAU, AcqNotes	Use for repeated commercial/non-commercial products and services like engineering or studies
	Negotiated Contracts		FAR 15, DAU	Use for technology that is not commercially available or dual use especially for attracting large national security contractors that provide unique solutions for the government
	Modular Contracts		FAR 39, DAU, 41 U.S.C. § 2308	Use for FAR-based major AI acquisition. FAR 39 enables rapid contracting procedure that can keep pace with the AI technology evolution
	IDIQ		FAR 16.5, DAU	Use for easy access to products and services at pre-negotiated prices
	Basic Ordering Agreements		FAR 16.7, DAU	Use for uncertain product and service requirements
	Small Business		FAR 19, DAU	Use when a small business can support, due to sole source and quick contracting method
Non-FAR Based	Other Transactions	Research	10 U.S.C. § 2371, DAU	Use for basic, applied, and advanced research projects, to advance new technologies or evaluate feasibility
		Prototype	10 U.S.C. § 2371b, DAU, Contract Attorney Deskbook Ch. 35	Use for proof of concept, model, novel application of commercial technology, demonstration of operation utility, etc.
		Production	10 U.S.C. § 2371b(f)	Use only as a follow-on production project to a Prototype OT contract
	Procurement for Experiments		10 U.S.C. § 2373, DAU	Use when a quick turnaround for test or experimentation purposes is needed
	CRADA		15 U.S.C. § 3710a, DAU	Federal labs use these to partner with Federal and/or non-Federal agencies to work on a combined research project

	PIA	15 U.S.C. § 3715, DAU	Federal labs use these to partner with State or local government or nonprofit entity to support technology transfer
	TIA	10 U.S.C. § 2371, 10 U.S.C § 2358, 32 CFR Part 21, DAU	Use when no other contract or agreement type works to promote new relationships in researching future defense needs

Step Four: Contract Deliverables

Machine learning algorithms inherently require metrics to inform development. DoD acquisitions must consider technical metrics, such as those listed below, to ensure model success and government quality assurance. Therefore, metrics should be included as contract deliverables throughout the entire AI lifecycle.

Regression Algorithms
For regression algorithms, common metrics include: Mean Squared Error (MSE), Root MSE (RMSE) and Mean Absolute Error (MAE).
Classification Algorithms
For classification algorithms, common metrics include: Accuracy, Precision, and Recall, F1 Score, Logarithmic Loss (Log loss), Categorical Crossentropy, and Area Under Receiver Operating Characteristic (ROC) Curve (AUC).
Reinforcement Learning
Evaluation criteria for reinforcement learning (RL) depend on how each RL algorithm's rewards are classified; however, the common metrics for RL are the total reward scores.

3.3 Development and Integration

Data fuels the development of AI and must be recognized as a critical component of AI acquisition. There is not an explicit number of data sets required for any AI solution; however, the program manager should realize that data density is important. Data density helps eliminate bias in a model and can improve modeling predictions. The USAF AI Annex to the DoD AI Strategy advises to treat data as a strategic asset. The DAF must become a data-ready organization in order to effectively acquire AI systems:

> "AI requires a vast and diverse amount of representative data whether derived from an authoritative source or generated through digitizing our daily workflows...For AI to influence operations, data must be provided to algorithms and archived in near real-time. As such, it is a priority to architect solutions that provide the shortest path between development and operational events. This crucial enabling activity, however, is only valuable if we trust how this data was obtained, generated, and/or trained.[1]"

The DoD needs to be primed to receive and distribute data so that any of its AI models can learn from reliable and relevant data sources. The following subsections will discuss how to collect data with respect to machine learning models. Without proper data grooming techniques, the model is susceptible to risk leading to unintended effects or unsuccessful model performance.

[1] Department of the Air Force Annex to the 2019 Department of Defense Artificial Intelligence Strategy.

AI-Ready Data Focus

The importance of high-quality data requires accentuation, gathered with a purpose and with specific guidelines in mind. Accordingly, program offices should consider the following 2020 DoD Data Strategy Data Principles to validate data quality:

> **Data is a Strategic Asset**: "DoD data is a high-interest commodity and must be leveraged in a way that brings both immediate and lasting military advantage."
>
> **Collective Data Stewardship**: "DoD must assign data stewards, data custodians, and a set of functional data managers to achieve accountability throughout the entire data lifecycle."
>
> **Data Ethics**: "DoD must put ethics at the forefront of all thought and actions as it relates to how data is collected, used, and stored."
>
> **Data Collection**: "DoD must enable electronic collection of data at the point of creation and maintain the pedigree of that data at all times."
>
> **Enterprise-Wide Data Access and Availability**: "DoD data must be made available for use by all authorized individuals and non-person entities through appropriate mechanisms."
>
> **Data for Artificial Intelligence Training**: "Data sets for A.I. training and algorithmic models will increasingly become the DoD's most valuable digital assets and we must create a framework for managing them across the data lifecycle that provides protected visibility and responsible brokerage."
>
> **Data Fit for Purpose**: "DoD must carefully consider any ethical concerns in data collection, sharing, use, rapid data integration as well as minimization of any sources of unintended bias."
>
> **Design for Compliance**: "DoD must implement IT solutions that provide an opportunity to fully automate the information management lifecycle, properly secure data, and maintain end-to-end records management."

The DoD Data Strategy acknowledges AI-enabled systems are written and trained by data; making (good, clean, useful, x y z) data paramount to the successful performance of any AI-enabled system. Therefore adequate planning is necessary, as data acquisition and consolidation can often take longer than anticipated and can be overwhelming if not adequately planned for. In most cases, data consolidation requires significant time (e.g., often months and sometimes years) of proactive effort to achieve a robust and diverse dataset. Consequently, a program office should view data as a strategic asset, with a potentially long-term shelf-life.

Nevertheless, just gathering every piece of "data" does not solve our data asset crisis. Data can come in many forms, which may not be appropriate for the specific mission at hand. The PM, to the maximum extent possible, should request data in common formats and sizes to facilitate efficient data supply pipelines. In general, beneficial formats for AI projects include: image, video, tabular, comma-separated value (CSV), and/or tab-separated values (TSV). The proper data format will be dependent upon your AI algorithm and will be influenced by the type of model being trained (e.g., convolutional neural network, feed-forward processing, computer vision, etc). Even with the proper file format, data needs to be curated and conditioned prior to AI training. The data conditioning process can be more easily digested using the diagram below, illustrating the need for early user engagement:

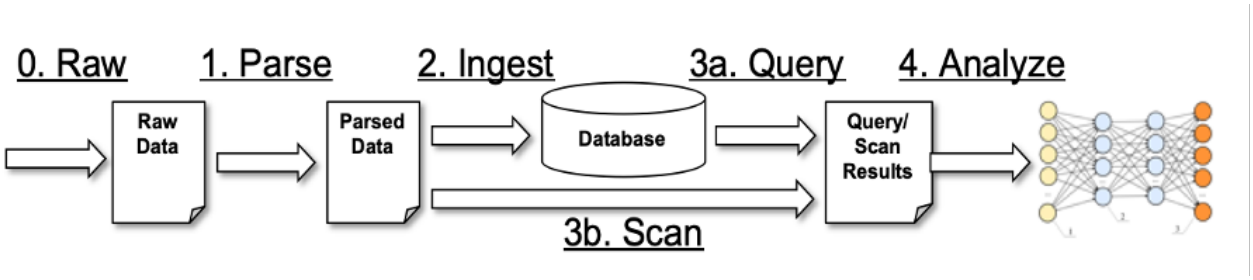

Figure 5: Standard data collection and management steps for AI

Figure 5 depicts the data conditioning lifecycle.

- Step 0: Most data arrives as unparsed raw text or imagery.
- Step 1: Each data segment must be purposely prepared for AI analysis and converted to the right size and/or format.
- Step 2: Once the data is correctly formatted, it will be ingested into a database, storing the data for later use. However, some information is best left within file in the file system.
- Step 3a / Step 3b*: Then, the data is queried from a database or scanned from the file system for retrieval and manipulation.
- Step 4*: Results from the query or scan are analyzed to determine if the model is meeting the desired intent
  *Steps 3 and 4 will be repeated many times to continually train the model until requirements are met.

Having AI-ready data (data that is ingestible and capable of being processed into a model; usually in tabular form) is the main determinant of effective AI organizations. As the DoD aims to acquire AI systems effectively, programs must work to provide a foundation for efficient data conditioning, which involves:

Early User Engagement: Program offices must engage stakeholders early and often to initiate data collection planning. If data collection requires any type of mission planning changes or other unforeseen policy implications, it may cause delays to the acquisition life cycle. Likewise, consistent communication throughout raw data collection will increase the likelihood of high-quality data within desired parameters.

Collecting Sufficient Data to Re-train Algorithms: Throughout the data life cycle, the developer will need to analyze results and re-train the model. Program offices should develop an acquisition strategy for future data procurement once the model is in production. Program offices should plan to constantly collect and replenish data lakes to ensure sufficient data is available for continuous re-training of models.

Parsing During Initial Data Collection: Parsing during the initial data collection step is much more efficient than performing later, as the programmer has context the program office does not. From discussions with MIT researchers, "Data Wrangling" has shown to account for 80% of the data conditioning lifecycle. By limiting "data wrangling" on the front-end, the data life cycle progresses more efficiently through its entirety.

Acknowledging the fact that data is paramount, the DoD has instituted several data repositories that can facilitate data collection and distribution to users. One example of an open-source data solution is the Unified Data Library (UDL) which provides a centralized data warehouse that handles multiple data types and data

sources. Centralized data repositories provide a trusted source of high-quality data program offices can use to improve the overall data lifecycle. Understanding how to access such sources is critical to the efficiency and effectiveness of AI system development. Appendix D provides a more robust description of the UDL and instructions for access.

Another example of a centralized data repository is the Chief Data Office's (CDO) VAULT (Visible, Accessible, Understandable, Linked, and Trustworthy) Platform. VAULT provides Airmen with cyber secure, cloud-based tools to connect, find, share, and learn from Air Force data to improve readiness and mission success. It has a set of tools to support an entire lifecycle of data exploitation activities from the ability to ingest data, manage storage, manage metadata, manipulate, cleanse and experiment with data and visualize analytics results. The VAULT platform is hosted on AWS and allows for members of the Air and Space forces to load and access datasets. There are a variety of commercial tools such as Trifacta, Tableau, Databricks, Spark, Python, and Informatica. A user must be on a .mil network computer/VMOnce on the site, the initial access requires you to create an account. Given a proper justification, there are current lists of data in which the user can access; however, the VAULT does have certain fees associated with using and storing its data.

The VAULT SharePoint site is also a helpful resource.

Development

Data is the key component to developing successful AI solutions to problems. The actual development of the AI system will follow any number of software development strategies, such as Agile, Cognitive Project Management for AI (CPMAI), Cross-Industry Standard Process for Data-Mining (CRISP-DM), etc. This section highlights, as shown in Figure 6, an Agile approach where the software is planned, coded, built, and tested as increments of the full capability. Once a capability is created and deployed, more data is gathered and conditioned, then more capability is developed and/or refined. The cycle continues until the AI system successfully meets the requirements.

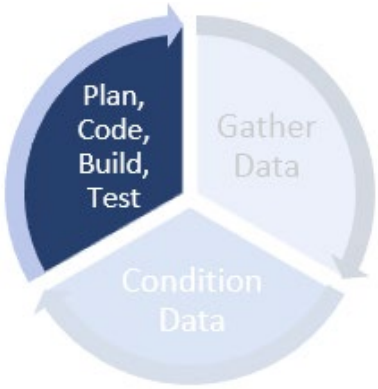

Figure 6: Agile development model for AI code, subset of AI Lifecycle

Integration

In the case of the AI system as only part of a larger system, integration with the larger system is required for end-to-end test and evaluation. There are no standard methods of integrating the AI system with the whole program. Touch points between the AI-subsystem and larger system will be program dependent. To the maximum extent possible, the Program Manager should advocate for common interfaces, e.g. REST. Early discussion on interface requirements and system architecture is crucial for future subsystem deployment and interoperability.

3.4 Test and Evaluation (T&E)

Artificial intelligence systems require testing like any other system; though AI/ML systems possess several unique characteristics and behaviors that present challenges for test and evaluation. These unique characteristics may include verifiable autonomous behaviors across dynamic environments, assurance of safety and security, instrumentation and measuring autonomy, and establishing user trust. The following distinction between autonomy and AI should be recognized – autonomy refers to an agent or machine being delegated to perform a task, while AI is a means to achieve that goal. Because AI is software-intensive, parallels can be drawn with agile software development, where developmental testing (DT) and operational testing (OT) run parallel. This enables a fail fast approach, allowing applications of solutions to root out early fault detection. Further, both the Air Force and Space Force continue to modernize legacy systems through integration of new AI structures. This forces a greater need to development test methods for assessing data transformation, incoming or out-going data streams, and algorithms feeding the AI/ML construct (analog-to-digital).All software and hardware, legacy or new, is subject to AI/ML modernization and will require testing.

Two key program management principles with regard to T&E are early tester involvement and integrated DT/OT. For autonomous systems these are especially relevant. Because autonomy relies on trust and the ability of the agent or machine to perform a task once performed by a human, testers are essential for the calibration of trust during system evaluation. Therefore, as early as possible testers (DT and OT alike) must be included so that a holistic approach can be taken in order to inform the design of the system and the scoping of the overall test program. Without such inputs, there will be an elevated risk of under-delivering on capability, running short on time for the test program, which ultimately, drives program costs. This risk will either be incurred by the program with delays or transferred to the user as a less-robust, measured AI/ML product.

A major difficulty in testing AI/ML is the concept of explainability. While we have witnessed an increase in the implementation of AI/ML models within the acquisition, engineering, and science communities, these models have correspondingly increased in their complexity and explainability. Explainability refers to AI/ML "black-box" models where the user has no access or knowledge of the model parameters.

Several approaches exist to expedite test-specific middleware for safety, security, effectiveness and efficiency, while also addressing explainability. For autonomy of air vehicles, these approaches are listed in the 412th Test Wing's Technical Handbook:

- Run Time Assurance (RTA)
- Modeling Simulation (M&S) to include Live Virtual Construct (LVC)
- Open Systems Architecture (OSA)
- Applications of Surrogate Test Platforms[2]

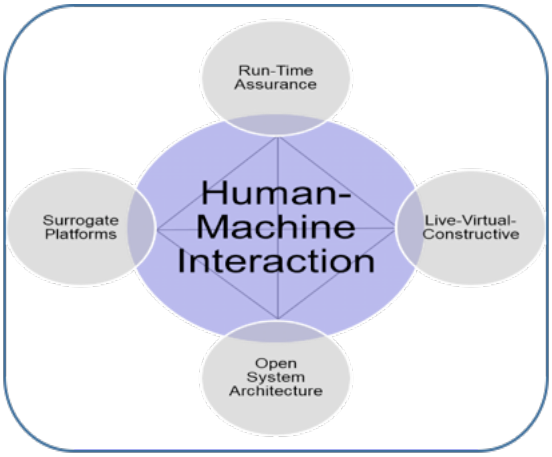

Figure 1: The Five Pillars of Testing Autonomous Systems

RTA can be summarized as a deterministic wrapper that serves two primary functions: (1) reliably detect problems (software or environmental) and (2) switch to a recovery/safe mode in the event of a failure. Other examples can be seen in Testing Autonomy in a Complex Environment (TACE), safety pilot, and manual flight mode.

LVC, a form of M&S, is a growing concept that is being applied for engineering and test vernacular. It merges real-world test assets (live), human-operated, simulated entities (virtual), and computer-controlled, simulated entities (constructive) to interact seamlessly with each other. This can provide a cost effective approach and create a pathway for testing when there are real-world AI/ML constraints across multiple domains.

The Air Force Test Center (AFTC) provides resources through the Test Flag Enterprise, a conduit for LVC testing and data collection for AL/ML efforts:

[2] 412th Test Wing, Test and Evaluation of Autonomy for Air Platforms, Jul 2020

Emerald Flag – focuses on providing an exercise for testing advanced long-range kill chain-enabling programs (96 TW).
Orange Flag – aims to bring test as close as possible to the warfighter through combat relevant testing early in the development cycle (412 TW).
Black Flag – delivers a high-end trial for new technology and tactics that are under operational evaluation (53 WG).

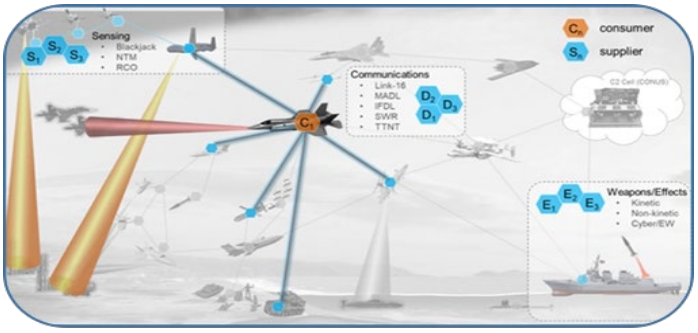

Figure 2: Emerald Flag

OSA allows for specific testing of autonomy services and it enables the decoupling of the autonomy from a specific platform. The Core Framework is built upon OSA, but the AI/ML powered capability is implemented at the app layer. This enables a program to add/subtract efficiently, without totally reconstructing the entire system to install a new capability.

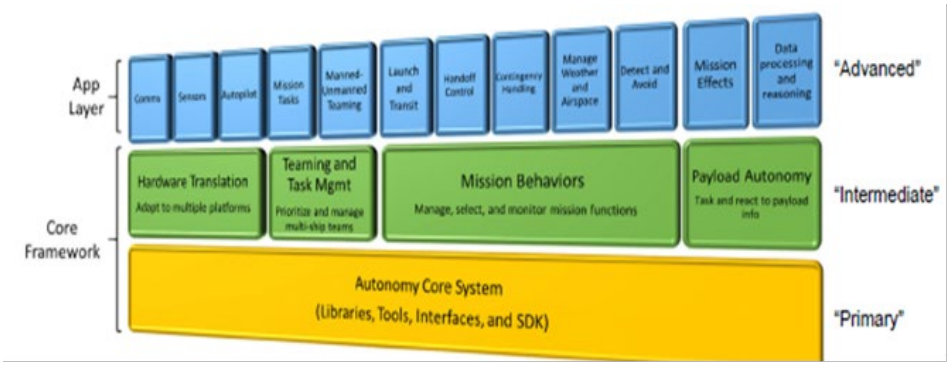

Figure 3: Open Systems Architecture - Cake Layer

Surrogate Platforms allow for greater cost-savings, lower risks, and greater flexibility to meet AI/ML testing schedules for test ranges and asset availability (UAS vs actual aircraft). This methodology also enables the separation of novel hardware from novel software, while the traditional problems of vehicle design are solved through RTA and human-machine interface means. The PM can develop novel algorithms in parallel in a cost-effective and low risk stable platform.

The Arnold Engineering and Development Complex (AEDC), Arnold AFB, Tennessee , 412th Test Wing, Edwards AFB, California, and the 96th Test Wing, Eglin AFB, Florida are some of the locations that support AI/ML test and evaluation for the Air Force. The 17th Test Squadron out of Schriever Space Force Base, CO is one of the USSF's main test and evaluation units. The Test Resource Management Center (TRMC) also looks across the entire test and evaluation infrastructure to align test and evaluation efforts with DoD modernization goals and ensure ranges are ready to test new capabilities as they emerge. The Center delivers AI/ML strategic priorities to harden digital environments, promote digital transformation, and develop T&E tools across the DoD. Program managers should be mindful that the degree of test and evaluation is subject to the nature of your program and operational requirements. In some cases where AI/ML is embedded into an already fielded system, an Operational Utility Evaluation may be appropriate.

Programs requiring a Product Support Strategy should account for the entirety of the system's or software's lifecycle, not the just particular acquisition and sustainment phases. The assurance that AI systems will function

as expected will require a holistic approach that tests and evaluates while the AI system continues to learn and adapt in its operating environment. Therefore, whenever possible, the AI portion of the program should seek a continuous authority to operate (cATO) and incorporate continuous testing and evaluation and cybersecurity to allow for continuous delivery of capabilities.[3]

The program team should seek to answer the following questions as early as possible:

- Will this capability eventually evolve or merge with other networks or AI/ML rooted systems?
- Does the vendor provide the necessary tools to secure, control, and test the presented AI/ML capability?
- Does the contract require the vendor to provide a method or dictionary to translate algorithms, models, and reports?
- What workarounds are required to elicit a response to get the required data?
- Is flexibility for the tester being sacrificed in order to design too quickly to the desired end state?
- What impact will this have on the test program overall?
- Are there test-specific software/interface requirements to consider?

3.5 Deployment

Following test and evaluation of the entire system, the deployment follows an official assignment of the system to a unit for operational use. Like traditional systems, AI systems have certain requirements and limitations of deployment that program teams need to consider when transitioning from T&E to operations. There are two standard cases from which AI systems are developed (some considerations for deployment are enumerated below). The first case is that the system was developed on a government system and granted a cATO during the development phase. The second case is that the system was developed on a contractor's system and needs to be transitioned to the government; whether the actual hardware along with the AI model and data or just the AI model and data.

For either case the program team should seek to answer the following questions as early as possible:

- Where will the algorithm reside and run?
- Where will the training data reside?
- Where will the operational data come from and where will it reside (possibly a higher classification)?
- Is there funding available to pay for a cloud compute and store solution, if applicable?
- Is there funding to move the algorithm and data from the T&E system to the operational system?

3.6 Sustainment

Algorithm sustainment is the maintenance and continuous improvement of an AI/ML system that is operational. Once the algorithm becomes operational, it will continually be fed new and more diverse data sets than in initial development and training. This process is required to enable the AI/ML system to meet the needs of the operators in a dynamic operational environment.

Closely tied to sustainment is the opportunity to proliferate the algorithm to other DoD programs. Currently, there is no enterprise-level solution to make entire AI/ML systems available for the DoD to use; however, there are paths in place for storing DoD-only datasets such as the Unified Data Library. This very important part of the AI lifecycle enables other organizations to benefit from the research, development, and datasets created by one program. If in the future there is an enterprise-level solution available, then the PM should seek that out for consideration.

[3] DoD, Cybersecurity Test and Evaluation Guidebook, 10 Feb 2020

3.7 Ethics

The DoD uses the phrase Responsible AI to describe operationalizing the DoD's Ethical Principles across its people, processes, partnerships, and policy. In this, the DoD's goal is to develop and operate AI systems responsibly to build trust with all who are affected by it.[4]
The DoD's five ethical principles for AI are:

1. **Responsible**. DoD personnel will exercise appropriate levels of judgment and care, while remaining responsible for the development, deployment, and use of AI capabilities.
2. **Equitable**. The Department will take deliberate steps to minimize unintended bias in AI capabilities.
3. **Traceable**. The Department's AI capabilities will be developed and deployed such that relevant personnel possess an appropriate understanding of the technology, development processes, and operational methods applicable to AI capabilities, including with transparent and auditable methodologies, data sources, and design procedure and documentation.
4. **Reliable**. The Department's AI capabilities will have explicit, well-defined uses, and the safety, security, and effectiveness of such capabilities will be subject to testing and assurance within those defined uses across their entire life-cycles.
5. **Governable**. The Department will design and engineer AI capabilities to fulfill their intended functions while possessing the ability to detect and avoid unintended consequences, and the ability to disengage or deactivate deployed systems that demonstrate unintended behavior.

In support of developing and operating AI systems responsibly, The Deputy Secretary of Defense instantiated the six foundational tenets for responsible AI (RAI):

1. **RAI Governance**: Ensure disciplined governance structure and processes at the Component and DoD-wide levels for oversight and accountability and clearly articulate DoD guidelines and policies on RAI and associated incentives to accelerate adoption of RAI within the DoD.
2. **Warfighter Trust**: Ensure warfighter trust by providing education and training, establishing a test and evaluation and verification and validation framework that integrates real-time monitoring, algorithm confidence metrics, and user feedback to ensure trusted and trustworthy Al capabilities.
3. **AI Product and Acquisition Lifecycle**: Develop tools, policies, processes, systems, and guidance to synchronize enterprise RAI implementation for the AI product throughout the acquisition lifecycle through a systems engineering and risk management approach.
4. **Requirements Validation**: Incorporate RAI into all applicable AI requirements, including joint performance requirements established and approved by the Joint Requirements Oversight Council, to ensure RAI inclusion in appropriate DoD AI capabilities.
5. **Responsible AI Ecosystem**: Build a robust national and global RAI ecosystem to improve intergovernmental, academic, industry, and stakeholder collaboration, including cooperation with allies and coalition partners, and to advance global norms grounded in shared values.
6. **AI Workforce**: Build, train, equip, and retain an RAI-ready workforce to ensure robust talent planning, recruitment, and capacity-building measures, including workforce education and training on RAI.

[4] Deputy Secretary of Defense, *Implementing Responsible Artificial Intelligence in the Department of Defense*, 26 May 2021

The DOD AI ethical principles are evident most notably in the human-machine teaming realm of AI systems. Human machine teaming is the dynamic between the AI/ML system and its user which attempts to create explainable and traceable predictions. AI is not always correct. Humans should monitor the system and provide quality assurance, ensuring user understanding of the model output and recommendation. For example, it is critical for a pilot to trust an artificially intelligent navigation system when it predicts collision and presents the pilot with the option to avoid it. If AI presented an unexplainable and untraceable prediction, the pilot may fail to execute the proper procedure. To ensure similar situations do not occur, the DoD should strive to ensure that AI systems abide by the DoD AI ethical principles.

3.8 AI Governance

Recognizing the need for AI Governance, the Deputy Secretary of Defense signed a memo on 8 December 2021 establishing the Chief Digital and Artificial Intelligence Officer (CDAO) position. The position will be reporting through the Deputy Secretary of Defense to the Secretary of Defense. The CDAO will be responsible for strengthening and integrating data, artificial intelligence, and digital solutions in the DoD. Furthermore, the Office of the CDAO shall: serve as the successor organization to the JAIC; serve as an intervening supervisor between Defense Digital Service and the Immediate Office of the Secretary of Defense and Deputy Secretary of Defense; and shall operationally align the Chief Data Officer to the Office of the CDAO.

The Department of the Air Force designated the Office of the Chief Information Officer (SAF/CN) as the lead for artificial intelligence and responsible artificial intelligence. The DAF also designated SAF/CN as the office of primary responsibility (OPR) for DAF AI Governance. An official AI governance plan is still in development as of February 2022.

IV. Funding Guidance for Artificial Intelligence Systems

The obligation of appropriated funds is provided by Congress in the Appropriations Act or Continuing Resolution. The basic principles of fiscal law that govern how executive agencies obligate funds are: 1) purpose, 2) time, and 3) amount. Appropriated funds may only be obligated for purposes for which the appropriation was made, except as otherwise provided by law. 31 U.S.C. § 1301. Where a particular expense is not specifically provided for in the appropriations act, it is permissible if it is necessary and incident to the proper execution of the general purpose of the appropriation. This "Necessary Expense Doctrine" has a three-part test to determine whether an expenditure can be obligated to a particular appropriation: the expenditure must bear a logical relationship to the appropriation sought to be charged; the expenditure must not be prohibited by law; the expenditure must not be otherwise provided for. If an expenditure meets the purpose requirement, it is available for a bona fide need during the appropriations period of availability. 31 U.S.C. § 1502(a). Finally, the Anti-Deficiency Act prohibits any federal officer or employee from obligating funds in excess of the amount available or in advance of an authorization unless authorized by law.

Funding efforts related to AI follow this fiscal law analysis. Before obligating funds towards an AI effort, it is recommended that you consult with your legal counsel and budget officer. To assist in determining the funds available for a particular AI effort, the DAF-MIT AI Accelerator, with coordination and approval from SAF/GC, SAF/AQ, and SAF/FM, developed the guidance below. Specific examples for funding AI systems can be found in Appendix E.

	A - Development Efforts (e.g., RDT&E)	**B** - Acquisition of Goods or Services (e.g., O&M, Procurement)
1 Software and/or Algorithm Activities	• Developing a new system that will be owned by the AF • Activities involved in bringing a program to its objective system • "Major upgrades increasing the performance envelope of existing systems…" • Performing design, integration, test, and evaluation on a Commercial-off-the-shelf (COTS) system that will be owned by the AF in order to modify that COTS system to meet objective performance or to include the COTS as a component in a system that is in development to meet its objective performance.	Software or application coding or labor activities that constitute maintenance or similar activities: • Costs incurred in continuing operations and current services • Minor improvements in software functionality which are accomplished during routine maintenance • "Software releases categorized as iterations on the basic release and not involving significant performance improvements or extensive testing" • Integration of COTS software into current AF systems, not requiring changes to the COTS software's or the current system's objective performance.
2 Data Activities	• The acquisition of data or data rights for use in the development of Air Force AI software (system will be owned by AF). • The examination, formatting, creation or use of existing data architecture solely for the development or testing of AI software/algorithms.	• The examination/study of existing data architecture preceding acquisition or deployment of an operational AI system. • The examination/study of existing data architecture for incorporation or optimization within an existing, operational AI system. • The formatting and management of existing data for inclusion or optimization in an operational AI system.
3 Funding	Where development is a direct Government cost (vice SaaS) it is funded out of an RDT&E appropriation.	Services shall be funded out of appropriations available for operations and maintenance. The purchase of equipment shall be funded out of the appropriate procurement account unless the effort falls below the expense investment threshold (whereby funding available for operations and maintenance shall be used).
	Where an IT effort cannot be funded by a single appropriation, it must be broken into discrete tasks, each budgeted under the proper appropriation.	

4.1 Funding for Artificial Intelligence Systems – Decision Tree

The numbers and letters in figures 7 and 8 refer to the corresponding row and column of funding guidance above, for example, 1A refers to Row 1, Column A. The question to ask when looking at the left side of each figure below is, "will the changes to the AI model or system result in…[software performance improvement / planning documents or onet-ime study on feasibility of data] "

Software/Algorithm Activities

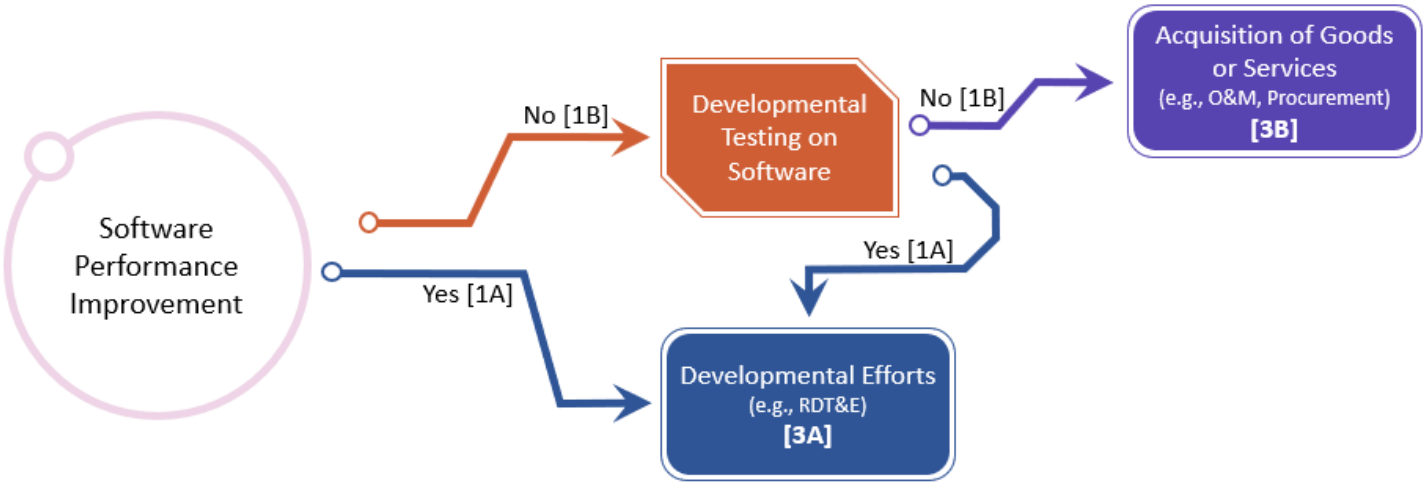

Figure 7: Decision tree to direct funding of software/algorithm activities

Data Activities

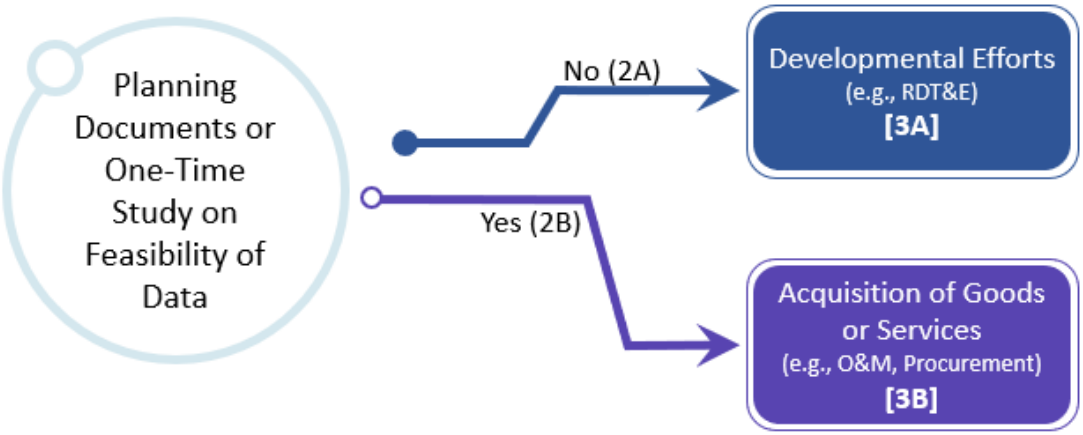

Figure 8: Decision tree to direct funding of data activities

V. Resources

AI Ecosystem

- Joint Artificial Intelligence Center (JAIC) provides
  - Access to the Joint Common Foundation
  - Technical Assistance on AI programs
  - Testing and Evaluation Support
  - AI Product Lines
  - AI Policy and Ethics Guidance
  - Visit their website for more information: https://www.ai.mil/
- Defense Innovation Unit (DIU)
  - Accelerate DoD adoption of commercial technology
  - Transform military capacity and capabilities
  - Strengthen the national security innovation base
  - Visit their website for more information: https://www.diu.mil/
- Dept. of the Air Force - MIT Artificial Intelligence Accelerator (AIA)
  - Strategic partnership between MIT and the Department of the Air Force
  - Designed to make fundamental advances in artificial intelligence to improve Department of the Air Force operations while also addressing broader societal needs
  - Research involves interdisciplinary teams, including Airmen, who collaborate across disparate fields of AI to create new algorithms, technologies, and solutions
  - Visit their website for more information: https://aia.mit.edu/
- Army Artificial Intelligence Integration Center
  - Partnered with Carnegie Mellon University
  - Connects with the broader artificial intelligence community
  - Basic research, force application, infrastructures & platforms, force integration, command & coordination, net-centric capabilities, sustainment, corporate management & support, battlespace awareness, and protection
  - Visit their website for more information: https://armyfuturescommand.com/ai2c/
- U.S. Naval Research Laboratory
  - Basic and applied research in artificial intelligence, cognitive science, autonomy, and human-centered computing
  - Adaptive systems, intelligent systems, interactive systems, perceptual systems
  - Visit their website for more information: https://www.nrl.navy.mil/itd/aic/
- MITRE Acquisition in the Digital Age (AIDA)
  - Transforming the federal acquisition environment using digital strategies and tools
  - Their website has many features to help PMOs think differently about acquisitions
  - Visit their website for more information: https://aida.mitre.org/demystifying-dod/
- Federally Funded Research and Development Centers (FFRDCs) and University-Affiliated Research Center Laboratories (UARCs)
  - Visit their website for more information: https://defenseinnovationmarketplace.dtic.mil/ffrdcs-uarcs/ or here https://rt.cto.mil/ffrdc-uarc/

5.1 Education Opportunities

- DAF's DigitalU and Udemy
  - The DAF has partnered with Udemy to provide a significant number of courses for free, sign up at https://digitalu.af.mil
  - Once you have your login information, navigate to https://digitalu.udemy.com/ and search for artificial intelligence to find over 500 courses available
  - Here are a few notable courses:
    - Udemy Executive Brief on ML
    - Udemy Executive Brief on AI (Includes Case Studies)
    - Machine Learning A-Z: (In depth – long, hands on with code templates)

- Coursera
  - This site features courses from over 200 leading universities and companies
  - They offer many free courses, but also offer paid projects, certifications, and even degrees.
  - Visit their website for more information: https://www.coursera.org/
- edX
  - This site features courses from a 160-member consortium of higher education institutions (i.e., MIT, Harvard, Berkely)
  - Most courses are free, though you can pay for a certificate
  - Visit their website for more information: https://www.edx.org/
- Miscellaneous
  - Commander's Read Ahead for AI
  - JAIC Understanding AI Tech
  - AI 101 Video (on MilTube)

5.2 Publications

- The MIT AIA publication page contains more than 100 scholarly works by the MIT faculty and students and DAF members supporting the AIA.
- 2021 GAO report on an Artificial Intelligence accountability framework: GAO-21-519SP
- National Security Commission on Artificial Intelligence https://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital-1.pdf

5.3 Miscellaneous Websites

- The Air Force Digital Guide SharePoint site (CAC-enabled): https://usaf.dps.mil/teams/afmcde/
  - The purpose of this guide is to aid programs and organizations in their journey toward meaningful enterprise digital capability and to provide Airmen our digital transformation information.

Appendix A - AI Lifecycle

AI is more than a complex neural network or regression algorithm. Each specific algorithm requires strategic planning, preparation, development, and data. The infographic below attempts shows the constant development cycle of gathering data, conditioning data, then developing the AI algorithm/model. Once development is complete and the AI system is operationalized, it must be sustained and may be proliferated to other programs.

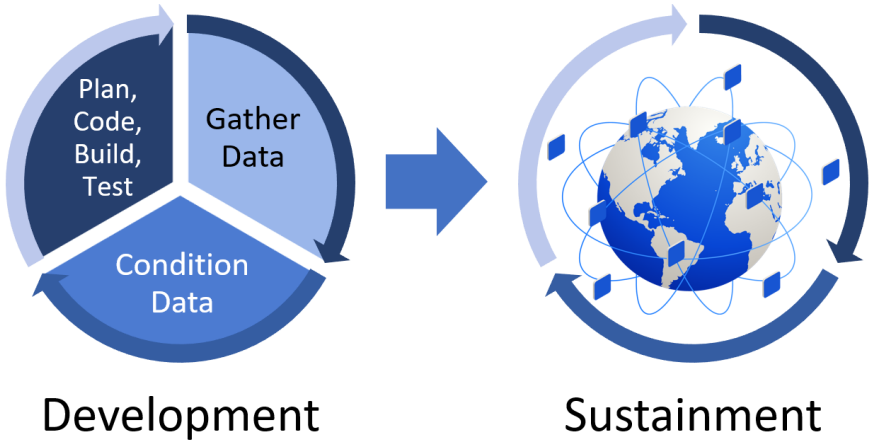

The simple process shown above is broken out below into its various parts with additional description to aid in understanding.

Gathering Data

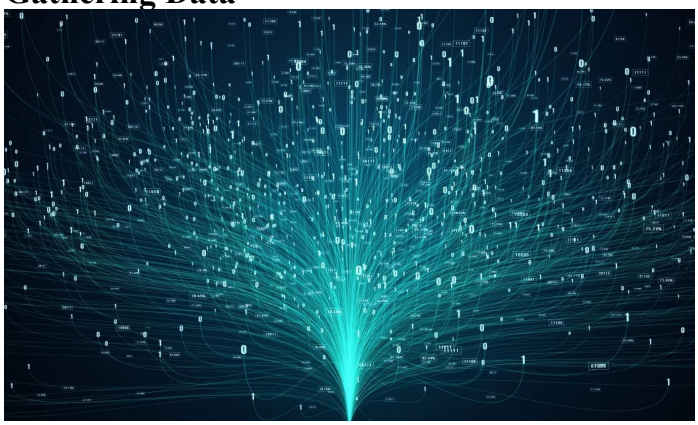

Data is critical to the overall lifecycle. Without diverse data sets from various sensors and/or sources, an algorithm is brittle and will not stand the tests of operational deployment. Algorithms are often developed without adequately formatted and parsed data and are required to wait until the dataset(s) have been developed. Without data, there is no operational AI. As such, gathering data is a key AI lifecycle stage that requires preemptive data aggregation in preparation for algorithm development, ultimately enabling rapid AI deployment and operational use.

Conditioning Data

Algorithms are sensitive and require specific types of data depending on the category of machine learning. Data needs to be conditioned, or prepared, for algorithm testing and training. The figure to the left is a high-level visualization of the data input, conditioning, and data output process of a generic algorithm. The data conditioning process begins with gathering raw data, which is then parsed into the correct size and format. The data is then directly used to test or train an AI/ML algorithm or is stored for later use. Conditioning also includes ensuring diverse, balanced, and robust data sets

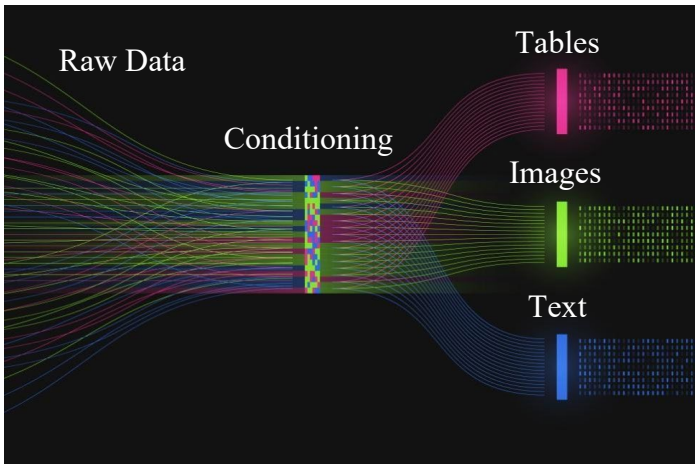

to ensure the AI/ML algorithm is trained adequately for all environments.

Model Development – Plan, Code, Build, Test

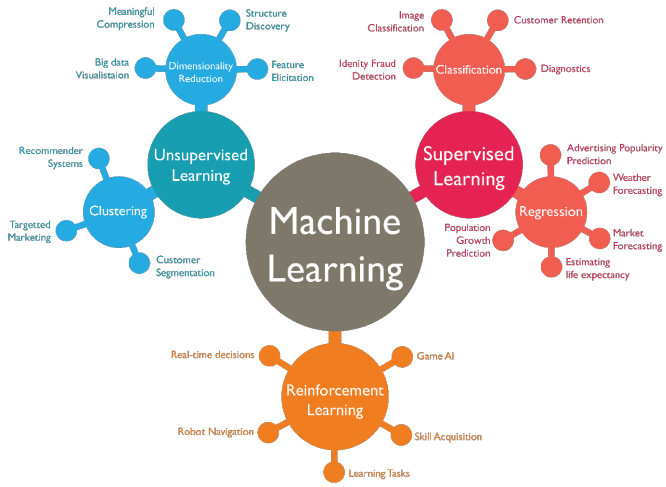

Artificial intelligence is composed of many unique algorithms. The first step is to determine which algorithm applies to your specific requirement. Once the type of AI is determined, the algorithm (mathematical function) is cyclically trained, evaluated, and tested using the specifically applicable data sets previously mentioned. Once the AI is fine-tuned, it is ready for deployment and operational use.

Model Sustainment

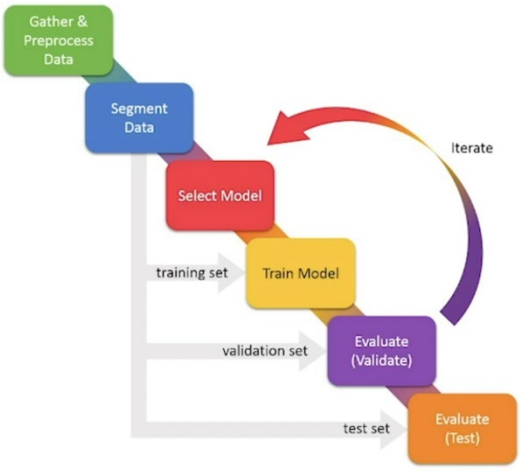

Algorithm sustainment is the maintenance and continuous improvement of an AI/ML system that is operational. Once the algorithm becomes operational, it will continually be fed new and more diverse data sets than in initial development and training. This process is required to enable the AI/ML system to meet the needs of the operators in the operational environment.

AI/ML Program Proliferation

Closely tied to sustainment is the opportunity to proliferate the algorithm to other DoD programs. Currently, there is no enterprise-level solution to make entire AI/ML systems available for the DoD to

use; however, there are paths in place for storing DoD-only datasets such as the Unified Data Library. This very important part of the AI lifecycle enables other organizations to benefit from the research, development, and datasets created by one program.

Appendix B - Software Acquisition Pathway

NDAA FY20, Subtitle A-- Section 800

The Secretary of Defense has authorized Section 800 authority for continuous integration and delivery of software applications and upgrades to embedded systems. This pertains directly to the continuous development of AI systems. Section 800 established a Software Acquisition Pathway to provide for efficient and effective acquisition of software to include:

> **Use of Proven Technologies:** Pathway is established to leverage the speed and talent of industry and academia. Leverage existing software solutions to streamline the development and integration of proven technologies.
>
> **Use of Authority:** The Secretary of Defense intends to initiate the engineering of new software capabilities quickly, demonstrate the viability and effectiveness of such capabilities no later than one year after funds are obligated to acquire software, and allow for continuous updating and delivery of new capabilities not less frequently than annually to iteratively meet a requirement.
>
> **Treatment Not as Major Defense Acquisition Program:** Software acquired using the authority of Subtitle A--Section 800 shall not be treated as a major defense acquisition program under Section 2430 of Title 10, US Code or DoD Directive 5000.01 without the specific direction of Under Secretary of Defense for Acquisition and Sustainment or a Senior Acquisition Executive.
>
> **Risk-Based Approach:** The Secretary of Defense shall use a risk-based approach for the consideration of innovative technologies and new capabilities for software to be acquired under the authority of Section A--Section 800.

Further, Section 800 enables an expedited process to streamline requirement, budget, and acquisition process to support rapid fielding of software applications, collection of data on software fielded, and continuous engagement with users to support engineering activities and delivery of software. Section 800 acknowledges the inapplicability of the Joint Capabilities Integration and Development System (JCIDS) manual and Defense Acquisition System (DAS) directive, while implementing specific elements to tailor requirements for the acquisition of software systems:

1) **Iterative requirement development:** Continuously define and update priorities through user and engagement and operational feedback. Requirements must be iteratively developed to ensure the AI system is meeting stakeholder requirements.
2) **Early Identification of User Need:** Identify appropriate stakeholders early in the development process and include rationale for how AI capabilities will support increased efficiency.
3) **Initial Contract Requirement/Format**: Include summary-level lists of problems and shortcomings in existing software. Then include how the new or upgraded system will meet requirements.
4) **Continually Refine and Prioritize Contract Requirements:** Informed by user engagement, continually re-evaluate and refine contract requirements through an evolutionary process.
5) **Continuous Consideration of Life cycle costs, technical data rights, and System Interoperability:** Throughout the life cycle of the program, program managers should continuously evaluate the cost of the system acquisition, ensure the protection of Government Purpose Rights, and maintain communication with the user to ensure system interoperability.
6) **Support Plan for Software Capabilities in Case of Cease of Developer Support:** Design a plan to ensure there is minimal gap support in the case that the software developer can no longer support the system acquisition. Determine a contingency plan to prevent lost development.

7) **Rapid Contract Procedures:** Use efficient contracting mechanisms, as discussed in part 5 of this guidebook, to expedite timeframes and make awards. Define teaming arrangements and options to enable timely award.

8) **Program Execution Process:** Define efficient processes to support development and test infrastructure, automation and tools, digital engineering, and data collection and sharing across DoD. Ensure the program office understands the roles of developmental and operational testing activities throughout program execution. Pre-determine key decision-making and oversight events throughout the life cycle of the program. Lastly, define support processes such as independent costing activity, operational demonstrations, and performance metrics.

9) **Cybersecurity Metrics:** Define and understand metrics that will be used to ensure the system meets all cybersecurity requirements. Such metrics should match wide-recognized standards and industry best practices. Metrics should be generated and made available to DoD and congressional defense committees.

10) **Administrative Procedures:** Include who may initiate and approve an acquisition under this authority, the roles of product teams, team selection and supporting process, governance and oversight roles and responsibilities, and appropriate independent assessments, testing, and cost estimation.

11) **Flexibility Mechanisms and Waivers:** Include use of other transaction authority and broad agency announcements.

12) **Congressional Reporting Mechanisms:** Include reporting procedures to include initiation of Section 800 pathway, data regarding individual programs or acquisition activities, how acquisition activities are reflected in budget justification, and compliance with other reporting requirements.

Full access to Subtitle A--Section 800 (Page 133 STAT. 1478): FY2020 NDAA
To find out more about how this pathway works within the DoD see https://aaf.dau.edu/aaf/software/

Appendix C - MITRE Innovation Tool Kit: Problem Framing Canvas Instructions

This approach helps teams:

- Establish consensus about the team's purpose
- Gain a sense of what "done" will look like
- Define the scope of a team's initial activities and goals
- Reduce the likelihood of working at cross purposes

HOW

Begin in the upper left corner and capture some words about the problem area. Work through the remaining boxes on the canvas and return to this box throughout the process as your understanding of the problem develops. Answer the questions that are most relevant to your problem area. Feel free to skip any questions that do not seem to apply. Use your inputs to build a problem statement in the bottom box and turn it into an actionable "How might we…" question.

Use the canvas by yourself or in a group. Doing some quick research to collect any necessary information, statistics, or data may be helpful prior to or during the activity.

FACILITATION TIPS

The Problem Framing canvas helps create an inclusive environment, where all members of the team are able to contribute to building a shared understanding of the problem at hand. Facilitators should take care to build a high-trust environment and ensure each participant is encouraged to speak up. This tool helps foster an adaptable mindset among groups, and they often end up coming up with new ways to describe the problem they are working on.

Produce an effective problem statement by avoiding problem statements that:

1. Assign a cause
2. Contain the solution
3. Are based on conjecture or belief rather than fact
4. Are too long
5. Do not describe actual current condition or problem condition
6. Do not describe the ideal or desired condition
7. Are not measurable
8. Are unclear
9. Are not specific
10. Refer to issues outside of the scope of the actual problem

QUESTION BANK (Additional Qs to help guide the discussion)

1. What is the scope? How can we expand the scope?
2. What are the elements of the problem? (Physical, social, emotional, professional, primary, secondary)
3. What prevents/ hinders/ blocks other attempts, or groups, at solving this problem?
4. What are we avoiding when looking at or solving this problem?
5. What are the stigmas associated with this problem? Is there a stigma we're trying to avoid?
6. What are the concerns or worries others may have about this problem?
7. What is unspoken about this problem?
8. What language or vocabulary do we not have and need to establish?
9. Who denies that this is a problem?
10. How does the environment or surrounding "landscape" impact this problem? What historical contexts may be influencing or biasing (our view on) this problem?
11. How are we making time to do the work that needs to be done to ensure we are designing equitably?

12. Whose perspective or experience isn't captured yet? Who doesn't immediately come to mind but also has this problem?
13. What types of inequities (uneven distribution or imbalance) exist in this problem? Power, financial, access, etc.

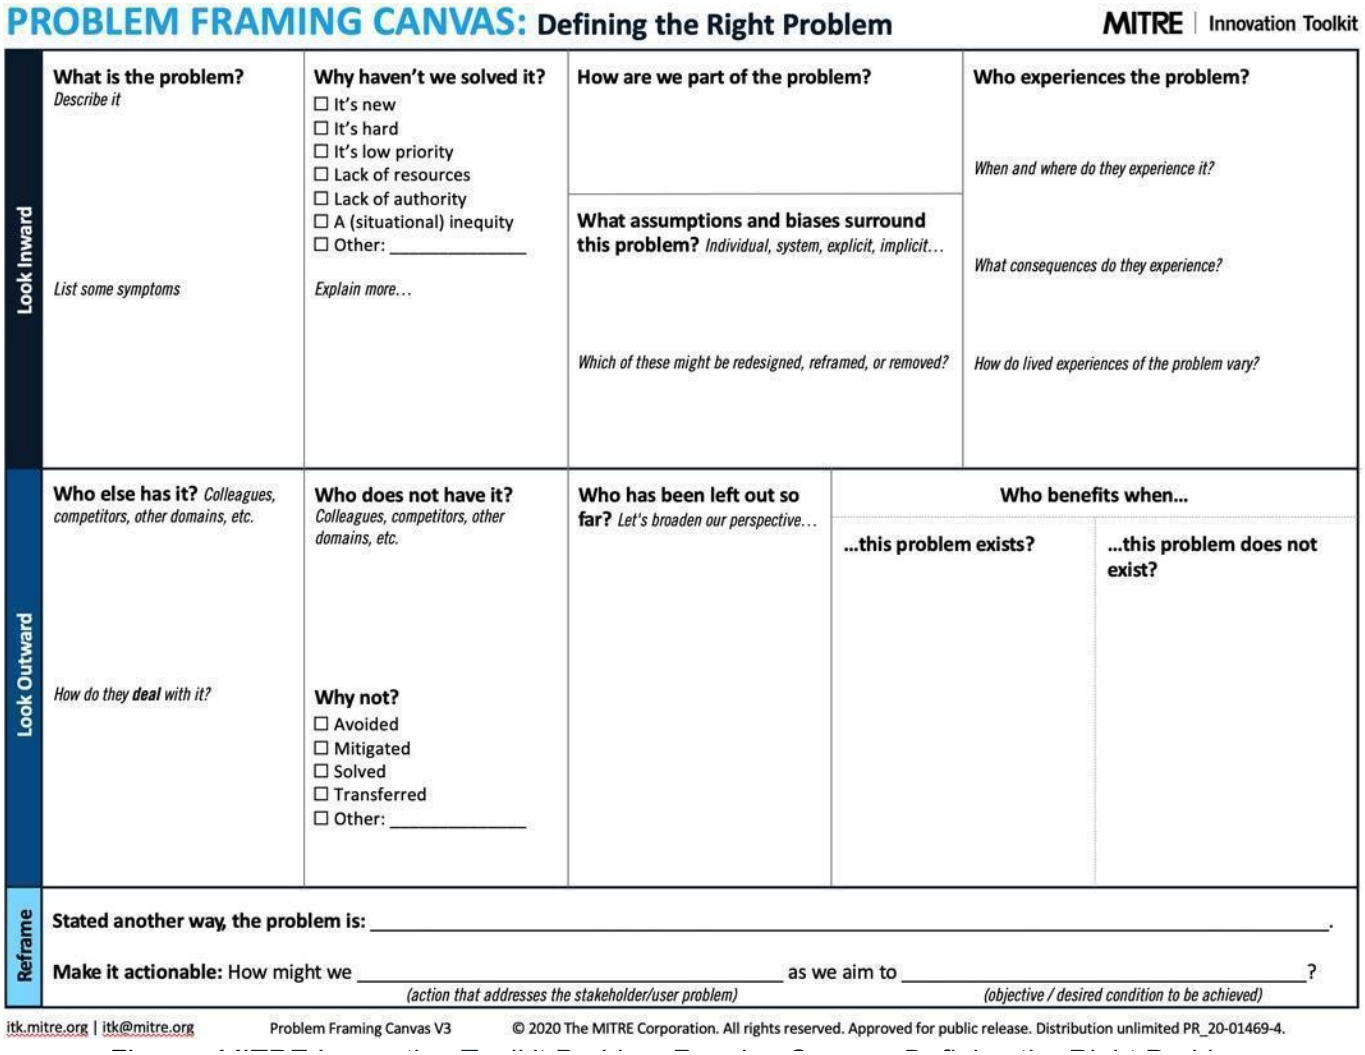

PROBLEM FRAMING CANVAS: Defining the Right Problem — MITRE | Innovation Toolkit

Look Inward

What is the problem? *Describe it*

List some symptoms

Why haven't we solved it?
- ☐ It's new
- ☐ It's hard
- ☐ It's low priority
- ☐ Lack of resources
- ☐ Lack of authority
- ☐ A (situational) inequity
- ☐ Other: _____

Explain more…

How are we part of the problem?

What assumptions and biases surround this problem? *Individual, system, explicit, implicit…*

Which of these might be redesigned, reframed, or removed?

Who experiences the problem?

When and where do they experience it?

What consequences do they experience?

How do lived experiences of the problem vary?

Look Outward

Who else has it? *Colleagues, competitors, other domains, etc.*

*How do they **deal** with it?*

Who does not have it? *Colleagues, competitors, other domains, etc.*

Why not?
- ☐ Avoided
- ☐ Mitigated
- ☐ Solved
- ☐ Transferred
- ☐ Other: _____

Who has been left out so far? *Let's broaden our perspective…*

Who benefits when…

…this problem exists?

…this problem does not exist?

Reframe

Stated another way, the problem is: _____.

Make it actionable: How might we _____ as we aim to _____?
(action that addresses the stakeholder/user problem) *(objective / desired condition to be achieved)*

itk.mitre.org | itk@mitre.org Problem Framing Canvas V3

Note: It's common for teams to iterate through multiple Problem Framing canvases, or that certain sections of this canvas would be complemented by another ITK tool to go more in-depth.

Appendix D - Unified Data Library (UDL)

On January 8, 2021, the Chief of Space Operations (CSO), Gen John W. Raymond, signed a memorandum stating that the Unified Data Library (UDL) will be the single source for accessing and managing all data in support of USSF operational systems. Additionally, the memorandum notes that the UDL will provide a central location to find and access data, enabling superior analytics through data integration and the strengthening of data security and trust.

Today, the UDL operates at the unclassified, Secret, and Top Secret levels. Those instantiations are live today while a Special Access Program (SAP) repository is currently under development. The multiple classification instances of the UDL are synchronized in near real time. The UDL is hosted in a single Amazon Web Service (AWS) region, but hosted across 3 availability zones for unclassified, and two availability zones for Secret and Top Secret. The public endpoint automatically switches if any availability zone goes down.

To create an account and access the UDL, you can visit the URL below.
https://unifieddatalibrary.com/storefront

Once you have created an account, to request data source integration with the UDL you can submit a data onboarding form. Once logged on, click the person icon in the bottom left of the page, then click "contribute data to the UDL."

Additionally, UDL API formats and schemas are kept in the Storefront at this location:
https://unifieddatalibrary.com/storefront/#/openapi

Currently, there are more than 30 API services for manipulation and querying of specific data sets.

To serve the wide variety of users and use-cases, the UDL has developed several different data distribution mechanisms to include:

1) **REST API**: Most commonly used service for users or processes that need data to support some specific analytical function
2) **History REST API**: Optimized for access to historically archived data in support of forensic analysis.
3) **Bulk Data Request API**: Preferred for Big-Data and AI data needs that require large volumes of data. Allows for synchronous data requests that notify the user via email or text message upon completion; supports CSV and JSON output formats.
4) **Secure messaging**: Allows for easy-to-implement real-time data streaming. The caller queries the offset for a particular service. The caller then uses that offset to query the topic and retrieve data.

Additionally, there are several labeled datasets on the UDL today. To access these datasets, under the Provider Discovery portal on UDL, click on SSC, then find SatNet and request access. For FOUO access, log into the UDL with a CAC. The data sets include annotated, rate-tracked GEO-targeted Resident Space Object (RSO) sky images. The objects were collected from four Raven-class sensors, from three locations, over two years. Additionally, there are labeled simulated sky images from AFRL's SatSim tool, of annotated breakup and collision RSOs in LEO, MEO, and GEO data.

Appendix E - Funding for Artificial Intelligence Systems - Examples

Software/Algorithm Activities Examples

Scenario 1 (1A)

- The Air Force is interested in an AI system in predictive maintenance that would predict when aircraft need to undergo maintenance. There are current software algorithms already developed that address car maintenance. The software code does not readily transfer to aircraft maintenance without significant changes in the software code. This is because the car maintenance data and aircraft maintenance data are so different that the underlying software code needs to change, while the overall principle remains the same. The end result leads to significant improvements in performance of the software code or changes to the basic release version of the software algorithm in order to be compatible with aircraft maintenance.

- Read more here: https://www.nationaldefensemagazine.org/articles/2018/9/26/army-investing-in-predictive-maintenance-for-bradleys

- The Air Force is interested in an AI system in predictive maintenance that would predict when aircraft need to undergo maintenance. There are currently existing software algorithms that predict when major mechanical machines like power plant turbines need maintenance before they fail. The underlying software algorithms must be changed in order to accommodate the new aircraft data. The aircraft algorithm works similarly to the turbine algorithm, but the data ingested, and the software algorithm need to undergo significant testing and changes.

Scenario 2 (1B)

- The Air Force is interested in an AI system in predictive maintenance that would predict when aircraft need to undergo maintenance. There are software algorithms developed that already address the problem in other areas, such as vehicle maintenance. The software does not need to undergo any significant changes and maintains its basic release. Additionally, there is no need for significant testing as the software algorithm has already been shown to work. It is only the data input into the system that changes with minor software updates and results in a different application of the AI system.

- A real-world example of this is the effect coronavirus has had on recommendations for shopping– think of the Google advertisements you get when you browse the internet. The pandemic has resulted in a massive change in people's buying habits, so the AI systems are not displaying optimized results. Instead, the AI system needs to be re-trained on updated data (i.e., buying habits) post-COVID outbreak to return accurate results. The underlying software code and algorithms remain unchanged. It is the data being given to the system that the AI is learning from that changes.

- Read more here: https://www.technologyreview.com/2020/05/11/1001563/covid-pandemic-broken-ai-machine-learning-amazon-retail-fraud-humans-in-the-loop/

Data Activities Examples

Scenario 1 (2A)

- The Air Force wants to develop magnetic navigation as an alternative to GPS. Artificial Intelligence and Machine Learning are needed to compute the position of a vehicle in relation to the earth at speed (that is at the speed of an aircraft) and to cancel out interference from other magnetic sources. To do this the Air Force gathers data from existing magnetic maps and creates new data sets from test flights and magnetic readers. In the context of RDT&E, Data Architecture, sometimes called data wrangling or data storage, is used to identify, store, structure, access and set policies for useful data sets in order to develop software code. In this scenario the data sets created, and gathered from existing sources, are used to inform the development of the software code. In addition, when the software code is used to inform how the data sets should be formatted for ingestion by the software code to be used. In this manner the data and the software are constantly informing each other and constantly evolving.

Scenario 2 (2B)

- Air Force has a need to determine operability of airfields for various aircraft platforms. Current artificial intelligence and machine learning software systems (either COTS, Government, or open source) have already developed algorithms that can ingest in data sources, such as aircraft size and runway length, to identify whether the airfield is appropriate for the aircraft platform. The Air Force has access to internal and external data sources that can be given to the software algorithm and identify results. The software algorithm gets more and more accurate with the more data provided; however, the underlying software code does not change. Data Architecture is used to identify useful data sources, format the data for software use, and set rules/policies for data storage, formatting, and identification. This can be done as a one-time step or every several years as a maintenance step to ensure data sources and policy are both usable by and optimizing the algorithm.

Scenario 3 (2B)

- Air Force has a need to determine operability of airfields for various aircraft platforms. Current artificial intelligence and machine learning software systems (either COTS, Government, or open source) have already developed algorithms that may be able to take in data sources, such as aircraft size and runway length, to identify whether the airfield is appropriate for the aircraft platform; however, it is unclear whether the algorithm will be able to make use of the data owned by the Air Force. The Air Force has a one-time study done to examine whether current data owned by the Air Force can be used with this algorithm and what, if any, formatting or changes can be done so the data properly interacts with the algorithm for the intended result.


PHANTOM