


NASA

Office of Inspector General




NASA's Management of Its Artificial Intelligence Capabilities



May 3, 2023

IG-23-012



Office of Inspector General

To report, fraud, waste, abuse, or mismanagement, contact the NASA OIG Hotline at 800-424-9183 or 800-535-8134 (TDD) or visit <https://oig.nasa.gov/hotline.html>. You can also write to NASA Inspector General, P.O. Box 23089, L'Enfant Plaza Station, Washington, D.C. 20026. The identity of each writer and caller can be kept confidential, upon request, to the extent permitted by law.

To suggest ideas or request future audits, contact the Assistant Inspector General for Audits at <https://oig.nasa.gov/aboutAll.html>.

COVER IMAGES

Top Left—Mars 2020 Perseverance Rover uses AutoNAV, an artificial intelligence application.

Top Right—Artist's rendering of a portable Martian greenhouse. NASA is developing artificial intelligence to autonomously assess plant health beyond low Earth orbit.

Bottom— The International Space Station's ECOSTRESS mission utilizes artificial intelligence to automatically select science observations.

RESULTS IN BRIEF

NASA's Management of Its Artificial Intelligence Capabilities



May 3, 2023

IG-23-012 (A-22-12-00-MSD)

WHY WE PERFORMED THIS AUDIT

Artificial intelligence (AI) is generally thought of as the capability of a machine to imitate intelligent human behavior. Aspects of this technology are used in a wide variety of applications from medical devices to autonomous vehicles, with tools like ChatGPT capable of mimicking human thought processes. NASA is a leader in AI usage and innovation across government, with applications such as a storm prediction tool that uses image recognition technology to identify atmospheric conditions to provide early warnings for destructive hailstorms and space vehicles such as the Mars Perseverance rover that uses an autonomous navigation system. While NASA and other federal agencies are continually exploring ways to incorporate AI into their organizations to meet agency goals, its adoption across such a wide spectrum of disciplines raises challenges for regulating and managing risks such as cybersecurity threats and drives the need for more detailed federal governance.

To that end, in February 2019 the White House issued Executive Order (EO) 13859 to promote sustained investment in AI research and development to generate technological breakthroughs while bolstering the requirement for AI developers to minimize the vulnerability of attacks from malicious actors and reflect federal priorities for innovation, public trust, and public confidence in AI systems. Similarly, EO 13960 issued in December 2020 seeks to promote the continued expansion of AI research and development in the United States while introducing measurable requirements to promote its transparency and trustworthiness. Although these EOs establish baseline principles for agencies to adopt into their AI governance policies and practices, AI standards remain in their infancy across the federal government.

In this audit we examined NASA's progress in developing its AI governance framework and standards and assessed whether security controls are being considered and implemented to protect AI data and technologies from cyber threats. To complete this work, we reviewed existing and proposed federal and NASA policies, regulations, frameworks, and industry best practices for AI governance; assessed NASA's reporting efforts for AI requirements; and reviewed the Agency's AI cybersecurity practices. We also interviewed NASA officials and other government agency AI representatives to benchmark NASA's progress in implementing federal AI guidance and discuss best practices.

WHAT WE FOUND

NASA has made progress in establishing an AI framework through development of the *NASA Framework for the Ethical Use of Artificial Intelligence (AI)* in April 2021, which drew upon the principles of leading AI organizations to guide consideration of ethics for AI projects and provide initial recommendations for NASA governance, advice related to AI, and questions for AI practitioners to consider during their work. Additionally, development of *NASA's Responsible AI Plan* in September 2022 identified NASA's Responsible AI officials and outlined how NASA intends to implement requirements of EO 13960, including capturing and reporting use case inventories, establishing oversight of AI projects to ensure continuous monitoring efforts, and engaging the AI community on the Agency's ethical AI standards and how to implement them.

However, NASA has not adopted a standard definition of AI and instead has three separate definitions: one in the *NASA Framework for the Ethical Use of Artificial Intelligence (AI)*, one in *NASA's Responsible AI Plan* that utilizes the definition found within EO 13960, and one on NASA's internal Artificial Intelligence Machine Learning SharePoint collaboration

website. While all three definitions are similar, subtleties and nuances in each can alter whether a particular technology is properly considered AI. Personnel we interviewed stated they reported AI based on their own individual understanding of what the term means rather than a formal definition provided by the Agency. As a result, NASA does not have a singular designation or classification mechanism to accurately classify and track AI or to identify AI expenditures within the Agency's financial system, making it difficult for the Agency to meet federal requirements to monitor its use of AI. Moreover, at NASA AI is generally managed as part of a larger project rather than as its own project and therefore is not tracked separately. This impacted the Agency's response to EO 13960 to create an AI inventory as well as its response to EO 13859 to compile an estimated annual budget for AI expenditures. To compile such an inventory and budget, NASA uses a multi-faceted data call to gather individual responses from AI users, which takes significant time to compile, validate, and vet and runs the risk of clerical errors that could be significantly lessened using an automated process.

Further, EO 13859 requires that technical controls exist to minimize AI vulnerability from attack by a malicious actor. Agency officials believe NASA's existing processes should be adequate to address security concerns specific to AI including monitoring requirements and ensuring it is properly safeguarded from cybersecurity vulnerabilities. However, our prior work has shown that NASA's fragmented approach to information technology management puts the Agency at a higher-than-necessary risk from cyber threats. Without an AI-specific classification mechanism or means to appropriately categorize and classify AI within its system of records, the Agency faces increased challenges to implement potential future federal AI cybersecurity controls.

WHAT WE RECOMMENDED

To improve the governance, budgeting, and cybersecurity of NASA's AI capabilities, we recommended the Associate Administrator for Technology, Policy, and Strategy; Chief Scientist; Chief Information Officer; and Chief Financial Officer: (1) establish a standardized definition for AI within the Agency to include harmonizing the definitions in the *NASA Framework for the Ethical Use of Artificial Intelligence (AI)*, *NASA's Responsible AI Plan*, and NASA's Artificial Intelligence Machine Learning SharePoint website; (2) ensure the standardized AI definition is used to identify, update, and maintain the Agency's AI use case inventory; (3) identify a classification mechanism to assist in the rapid application of federal requirements for cybersecurity controls and monitoring practices; and (4) develop a method to track budgets and expenditures for AI use case inventory.

We provided a draft of this report to NASA management who concurred or partially concurred with our recommendations and described planned actions to address them. We consider management's comments responsive; therefore, the recommendations are resolved and will be closed upon completion and verification of the proposed corrective actions.

For more information on the NASA Office of Inspector General and to view this and other reports visit <https://oig.nasa.gov/>.

TABLE OF CONTENTS

Introduction.....	1
Background	1
NASA's Artificial Intelligence Governance Efforts Are Progressing but Further Refinement Is Needed	11
Lack of a Standard Definition for AI Hinders the Agency's Ability to Effectively Manage Its AI Inventory.....	11
NASA's Classification and Tracking of AI Is Inadequate to Fully Address Current and Future Federal Requirements and AI Cybersecurity Concerns	13
Conclusion	16
Recommendations, Management's Response, and Our Evaluation	17
Appendix A: Scope and Methodology	18
Appendix B: Management's Comments	20
Appendix C: Report Distribution.....	23

Acronyms

AI	artificial intelligence
AIML	Artificial Intelligence Machine Learning
EO	Executive Order
GAO	Government Accountability Office
ML	machine learning
NIST	National Institute of Standards and Technology
OCIO	Office of the Chief Information Officer
OIG	Office of Inspector General
OTPS	Office of Technology, Policy, and Strategy
RAI	Responsible AI

INTRODUCTION

Though definitions vary, artificial intelligence (AI) is generally thought of as the capability of a machine to imitate intelligent human behavior. Aspects of this technology are broadly used in a wide variety of applications ranging from medical devices to autonomous vehicles to automated maintenance for military systems. While NASA has been utilizing AI for the past several years on experiments in low Earth orbit to conduct weather modeling and in deeper space to map terrain hazards for future landing sites, standards for AI are still in their infancy across the federal government. Executive Orders (EO) issued in 2019 and 2020 established baseline principles for agencies to adopt into their AI governance standards.¹ But as with any emerging technology, a risk identification process must be undertaken to ensure the technology is properly regulated and managed. However, without clearly defined AI definitions and standards to use during such a process, agencies—including NASA—will likely struggle to appropriately identify, categorize, and protect AI. Additionally, NASA may be hindered in its ability to properly identify and manage its AI, while keeping pace with a rapidly evolving AI technological landscape. In this audit, we examined the Agency’s progress in developing its AI governance framework and standards and assessed whether security controls are being considered and implemented to protect AI data and technologies from cyber threats. See Appendix A for details on the audit’s scope and methodology.

Background


The concept of programming a machine to solve problems was first introduced in 1950 by Alan Turing. An English mathematician and code breaker, he is known as the father of theoretical computer science and AI through the “Turing Test,” which proposed that if the right variable was introduced to a machine, it would be able to imitate human reactions and establish its own sense of intelligence. Five years later the term AI was officially coined. Since that time, the flurry of innovation occurring over the course of nearly 70 years has been nothing short of remarkable. AI has morphed from an algorithm to teach computers how to make simple decisions into a multi-faceted tool (e.g., ChatGPT or Google Bard AI) capable of mimicking human thought processes.² See Figure 1 for a timeline of AI milestones.

AI is now found in a wide variety of technologies used in everyday life such as video games, web searching, spam filtering, and voice recognition. More broadly, AI applications exist in a variety of sectors including transportation, health care, education, finance, defense, cybersecurity, and space exploration. As the use of AI expands, multiple federal agencies including NASA are exploring ways to incorporate it into their organizations to meet agency goals.

¹ EO 13859, *Maintaining American Leadership in Artificial Intelligence* (February 11, 2019) and EO 13960, *Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government* (December 3, 2020).

² ChatGPT and Google Bard AI are language model chatbots capable of answering questions using their respective datasets. ChatGPT is a Generative Pre-Trained-3 Transformer reliant on preloaded large sets of data to review. Google Bard AI is a Language Model for Dialogue Applications still in development that will be able to use the internet itself as a dataset.

Figure 1: Timeline of Select Artificial Intelligence Milestones



Source: NASA Office of Inspector General (OIG) visualization of the history of AI using select events from the National Institute of Justice, *A Brief History of Artificial Intelligence* (accessed January 25, 2023) <https://nij.ojp.gov/topics/articles/brief-history-artificial-intelligence>; federal guidelines; and NASA missions.

AI Innovation—from the Office to Outer Space

NASA's core tenets—exploring the unknown in air and space and making scientific discoveries for the benefit of humanity—demonstrate the Agency's commitment to leading in research and innovation. Towards that end, NASA is a primary consumer of AI and, according to Agency leaders, openly embraces its usage in as many ways as it can introduce AI legally, ethically, and for a designed purpose. For example, NASA uses AI in a wide variety of applications ranging from administrative tasks using analytics to review standard operating procedures to botanical experiments in low Earth orbit to weather modeling and space exploration.


The Agency's Earth-based AI usage includes projects like a storm prediction tool that uses image recognition technology to assist with weather and environmental event tracking. The project seeks to train machine learning (ML) image recognition techniques to identify atmospheric conditions on satellite imagery to eventually contribute to an early warning system for destructive hailstorms (given that certain atmospheric conditions are a strong precursor to destructive hailstorms).³ Such an early warning system could be invaluable towards safeguarding human lives and property in the event a storm trigger can be identified. In another example, NASA is utilizing an ML program to analyze parking spot availability at a NASA Center to reduce overcrowding and the amount of time employees spend

³ ML is a type of AI in which software applications become more accurate at predicting outcomes without being explicitly programmed to do so. ML algorithms use historical data as input to predict new output values.

searching for parking. Strategically placed cameras use ML image recognition to identify full and empty spaces and display estimated counts of free spaces on a website.

NASA's use of AI is also expanding into missions based in low Earth orbit and beyond. The creation of space vehicles like the Mars Perseverance rover is one of NASA's most robust AI platforms for exploration, providing the Agency with data and imagery that will be crucial for the success of future planetary missions.⁴ Perseverance houses multiple AI applications within its rover including "AutoNAV." This feature enables the rover to autonomously plan a safe path of travel using numerous camera and stereo technologies, ML algorithms, and technology that performs safety checks to create 3D maps of the terrain ahead to identify hazards and avoid rocks and other obstacles without input from controllers on Earth. While the rover navigates autonomously, it conducts AI-driven research, experiments, and engineering to empower future rovers with onboard autonomy. In addition, onboard science, image processing, terrain classification, and fault diagnosis are all key AI components of the rover that will enhance future capabilities to explore Mars and other planets. See Figure 2 for examples of AI at NASA depicted in the media.

Perseverance Rover Drives Autonomously



Perseverance looks back with one of its navigation cameras toward its tracks on July 1, 2021, after driving autonomously 358 feet.

Source: NASA/Jet Propulsion Laboratory-Caltech.

⁴ The Mars Perseverance rover, which launched in July 2020 and landed in February 2021, seeks to better understand the geology of Mars, identify evidence of ancient life, collect Martian surface samples, and test new technologies.

Figure 2: NASA's Use of Artificial Intelligence

New NASA-Google Collaboration Will Help People Track Air Pollution at the Local Level Using AI

By IANS - 16 September, 2022 - TWC India



NASA Supports Research to Advance Earth Science

Feb. 23, 2022



A prize competition is designed to engage underrepresented academic institutions in helping NASA make advancements in machine learning, AI, and developing of autonomous systems.

NASA and AI SpaceFactory's vaulted lunar outpost will be 3D printed by autonomous robots

By Niall Patrick Walsh Jul 1, '22 12:09 PM EST




Image via AI SpaceFactory

NASA astronauts on Artemis could talk to a spaceship computer

Daniel Fraser - August 20, 2022

11 6 minutes read



NASA to Upgrade Computing in Space with RISC-V

September 14, 2022 by Agam Shah



NASA's Perseverance Mars rover. Image credit: NASA/JPL

NASA realizes its probes and computers in space are outdated, and a massive upgrade is needed as the agency starts unlocking the mysteries of unexplored frontiers.

NASA is now rebooting its spaceflight computing hardware, starting with RISC-V chips made by Microchip and SiFive. The Mars rover called Perseverance was sent into space last year with an unimpressive 20-year-old PowerPC chip.

NASA has said the new Microchip product will provide "at least 100 times the computational capacity of current spaceflight computers."

Sources: Clockwise from top left, The Weather Channel; US Today (top right); Jet Propulsion Laboratory, NASA (middle); Enterprise AI (bottom right); and Archinect News (bottom left).

Going one step further, NASA is considering ways AI applications can be used to advance deep space exploration. The Agency is funding projects for AI applications to automate image analysis for planet and star classifications, develop autonomous spacecraft that can avoid space debris without human intervention, and create communication networks that are more efficient and distortion-free by using an AI-based cognitive radio.⁵ For example, Callisto is a technology demonstration project that involves a customized version of Amazon's digital voice assistant Alexa to act as a virtual assistant to astronauts while onboard their spacecraft.⁶ This AI was incorporated into the Orion Multi-Purpose Crew Vehicle spacecraft in 2022 for the Artemis I unmanned mission to the Moon and is planned for inclusion on

⁵ A cognitive radio is a system of wireless communication that allows a device to automatically detect an open or unused radio frequency and switch to it in order to improve performance.

⁶ The Lockheed Martin Corporation, under agreement with NASA, partnered with Amazon to show how astronauts and flight controllers can use human-machine interface technology like Alexa to make their jobs simpler, safer, and more efficient while also advancing human exploration in deep space.

future Artemis missions.⁷ These types of AI projects and applications will be crucial for Artemis and other missions as NASA continues to explore low Earth orbit, our solar system, and beyond.


Across the federal government, agencies are highlighting the diversity of AI innovations and capabilities. Specifically, 13 agencies have reported a total of 351 AI public use cases as of January 2023.⁸ For example, AI is being used to process human resources-related information to assist the U.S. Department of Health and Human Services' Centers for Medicare and Medicaid in reviewing agency staff changes (e.g., promotions, reassignments, and changes in supervisors) and generating information to ensure disability-related accommodations follow employees. The U.S. Department of Veterans Affairs also has a wide range of medical AI uses including using it as a "perfusionist" during heart and lung surgery. In this situation, a perfusionist operates a heart-lung bypass machine and supplies needed blood and oxygen to the patient whose heart has been stopped to perform the surgery. Even expert human perfusionists can become overwhelmed, given the critical demands on their skills. Consequently, a computerized system using AI to partner with perfusionists can improve patient safety and potentially avoid adverse events such as kidney damage or a blood clot from a heart-lung bypass.

A comparison of AI use case inventories across the government demonstrates that NASA is a leader in AI usage and innovation and illustrates its wide adoption across the Agency. As shown in Figure 3, NASA has more published AI use cases than other federal agencies on record.

⁷ Artemis I is the first in a series of increasingly complex missions to build a long-term human presence on the Moon with the ultimate goal of crewed missions to Mars in the 2030s.

⁸ Use cases are defined within EO 13960 as AI designed, developed, acquired, or used specifically to advance the execution of agencies' missions, enhance decision making, or provide the public with a specified benefit. While this EO requires most federal agencies to develop a use case inventory, it excludes the U.S. Department of Defense and agencies or agency components with functions that lie wholly within the Intelligence Community.

Figure 3: Federal AI Use Case Information (as of January 2023)



Source: NASA OIG representation of public reported use case information downloaded from the National Artificial Intelligence Initiative Office, *Agency Inventories of AI Use Cases* (accessed January 27, 2023) <https://www.ai.gov/ai-use-case-inventories/> and the NASA Technical Reports Server (accessed January 27, 2023) <https://ntrs.nasa.gov/citations/20220014918>.

Note: Public AI use case inventory totals were analyzed from information provided at <https://www.ai.gov/ai-use-case-inventories/>. Agencies that reported zero use cases (U.S. Agency for International Development, National Science Foundation, and the U.S. Department of Housing and Urban Development) were not included in the graphic.

Challenges for Adopting AI

With the public and private sectors adopting AI across such a wide spectrum of disciplines, inherent challenges are associated with its implementation. AI processes are being introduced into human resources tasks, information technology program expansion, research and development, and organization-specific products. While multiple industries have embraced AI technologies, many have still not clearly defined an approach for regulating and managing the risks associated with these technologies. The concept of AI risk management is similar in many respects to other risk management practices already in place for other types of technology. For example, risks to software or information-based systems—including concerns related to cybersecurity, privacy, safety, and infrastructure—also apply to AI systems and similarly can be characterized as long- or short-term, high- or low-probability, systemic or localized, and high- or low-impact. These risks are relatively easy to recognize and address as they apply to current AI applications.

However, while the promise of AI is great, its evolution presents a unique set of risks that are not as well known or documented. According to the Institute of Electrical and Electronics Engineers, AI poses potential new social, legal, and ethical challenges and its adoption will require new requirements to

address issues related to systemic risk, diminishing trust, privacy challenges, and data transparency and ownership.⁹ Additionally, AI and ML technologies have become targets for modernized cybersecurity attack methods.

Examples of modernized cybersecurity attack methods include integrity and confidentiality attacks. Integrity attacks deliberately attempt to cause an AI model to make errors. The most popular approaches for these attacks include data poisoning and evasion. In data poisoning, attackers embed malicious patterns into the AI's training data for the machine to learn, causing the model to learn these wrong patterns and skew the intended results. On the other hand, attackers using the evasion approach attempt to discover vulnerabilities in the AI and apply extremely subtle changes to exploit these weaknesses to modify the output. For example, an image recognition and classification system known as MobileNetV2 was attacked to alter its outputs. Using a publicly available evasion attack style program, this ML algorithm was tricked into identifying an image of Georgetown University's Healy Hall as a triceratops instead of a historic building. Finally, confidentiality attacks—sometimes called model stealing techniques—are used to duplicate or steal training data models from AI. If an attacker can steal or recreate the training models, they can make the AI unintentionally reveal proprietary or sensitive information to the attacker. Model stealing can be used to steal stock market prediction models and spam filtering models for the attacker to utilize or be able to optimize more efficiently against those models. Regardless of the type of attack, every user of AI must be aware of the challenges new technologies pose to the security of the user or industry.

AI Governance Best Practices

Acknowledgment of the challenges for adopting AI so widely throughout the government is driving the need for more detailed federal governance. To this effect, EO 13859 was issued in February 2019 to promote sustained investment in AI research and development to generate technological breakthroughs while bolstering the requirement for AI developers to minimize the vulnerability of attacks from malicious actors and reflect federal priorities for innovation, public trust, and public confidence in AI systems. The EO directed the National Institute of Standards and Technology (NIST) to issue a plan for the development of technical standards and related tools in support of reliable, robust, and trustworthy systems that use AI technologies. NIST's initial plan, released in August 2019, emphasized the need for global and federal collaboration towards the continued development and shaping of AI standards.¹⁰

Similarly, EO 13960 was issued in December 2020 to promote the continued expansion of AI research and development in the United States while introducing measurable requirements to promote its transparency and trustworthiness. EO 13960 established the following nine principles for agencies to adhere to when designing, developing, acquiring, and using AI in the federal government:

- *Lawful and respectful of our Nation's values.* Design, develop, acquire, and use AI in a manner that exhibits due respect and is consistent with the U.S. Constitution and all other applicable laws.

⁹ Institute of Electrical and Electronics Engineers, *IEEE Position Statement, Ethical Aspects of Autonomous and Intelligent Systems* (June 24, 2019). The Institute of Electrical and Electronics Engineers is the world's largest technical professional society, advancing innovation and technological excellence for the benefit of humanity and serving professionals in the electrical, electronic, and computing fields and related areas of science and technology.

¹⁰ NIST, *U.S. Leadership in AI: A Plan for Federal Engagement in Developing Technical Standards and Related Tools* (August 9, 2019).

- *Purposeful and performance-driven.* Seek opportunities for designing, developing, acquiring, and using AI, where the benefits of doing so significantly outweigh the risks and the risks can be assessed and managed.
- *Accurate, reliable, and effective.* Ensure the application of AI is consistent with the use cases for which that AI was trained.
- *Safe, secure, and resilient.* Include resilience when confronted with systematic vulnerabilities, adversarial manipulation, and other malicious exploitation.
- *Understandable.* Ensure operations and outcomes of AI applications are sufficiently understandable by both subject matter experts, users, and others, as appropriate.
- *Responsible and traceable.* Ensure human roles and responsibilities are clearly defined, understood, and appropriately assigned throughout the implementation of the AI.
- *Regularly monitored.* Ensure AI applications are regularly tested against these principles and determine if the application should be altered or superseded should the outcome become inconsistent with the AI's intended use or EO 13960.
- *Transparent.* Disclose relevant information regarding the use of AI to appropriate stakeholders, including Congress and the public, to the extent practicable.
- *Accountable.* Implement and enforce appropriate safeguards for the proper use and functioning of AI applications.

To complement this federal guidance, in June 2021 the Government Accountability Office (GAO) published an AI accountability framework outlining key practices of governance, data protection, performance and continuous monitoring for the deployment of AI systems for federal agencies.¹¹ The publication is a widely accepted set of best practices for enabling third-party assessments and audits of AI systems that resulted from a 2017 forum led by the GAO's Comptroller General of the United States. This forum brought together experts across the federal, academic, and nonprofit sectors to discuss the potential implications of AI developments in four areas—cybersecurity, automated vehicles, criminal justice, and financial services—as well as implementation of existing frameworks and executive guidance towards a means for auditors to develop credible assurance assessments of AI and AI systems.

Follow on requirements and best practices to implement these EOs are still in development at a variety of federal agencies. In addition to the EOs and GAO framework, NIST issued additional AI guidance and best practices in January 2023:

- *AI Risk Management Framework (V 1.0).* A guidance document for voluntary use by organizations designing, developing, deploying, and using AI systems to help manage the risks of AI technologies. The framework includes instructions for how to think about, communicate, measure, and monitor AI risks and its potential positive and negative impacts, and includes four core functions: Govern, Map, Measure, and Manage (as shown in Figure 4).
- *AI Risk Management Playbook.* A companion document to the *AI Risk Management Framework* that suggests ways to navigate and use the framework to incorporate trustworthiness considerations in the design, development, deployment, and use of AI systems.

¹¹ GAO, *Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities* ([GAO-21-519SP](#), June 30, 2021).

Figure 4: NIST AI Risk Management Framework Core



Source: NIST AI 100-1, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (January 2023).

As AI technologies and usages continue to evolve, increased federal requirements and regulations are likely to follow.

Evolution of AI Program Management at NASA

In 2018, a NASA-commissioned Enterprise Digital Transformation study and research team recommended that AI and ML be included as one of six key strategic foundations for larger integrated digital transformation efforts across the Agency.¹² Operating under the auspices of NASA's Office of the Chief Technologist and Office of the Chief Information Officer (OCIO), the Agency created the role of an Artificial Intelligence Machine Learning (AIML) Transformation Lead in April 2020 as part of its digital transformation efforts (under the OCIO) whose initial responsibilities included collecting data as required by EO 13960 and reporting NASA's AI use case inventory to the Office of Management and Budget. The AIML Lead also helped author the *NASA Framework for the Ethical Use of Artificial Intelligence (AI)* (hereafter referred to as the Framework for the Ethical Use of AI), released in April 2021.¹³ This framework drew upon the principles of several leading AI organizations to help guide consideration of ethics for AI projects while also providing initial recommendations for NASA governance, advice related to AI, and questions for AI practitioners to consider during their work.

¹² NASA's other five foundations for digital transformation are related to the collaboration and cultivation of data.

¹³ NASA/TM-20210012866, *NASA Framework for the Ethical Use of Artificial Intelligence (AI)* (April 2021).

Additionally, in November 2021 NASA established the Office of Technology, Policy, and Strategy (OTPS) within the Office of the Administrator to provide leadership with data- and evidence-driven recommendations to develop and shape NASA technology, policy, and strategy.¹⁴ Since its establishment, OTPS in conjunction with the Office of the Chief Scientist initiated a review of the Agency's AI and ML efforts and created *NASA's Responsible AI Plan* (hereafter referred to as the RAI Plan) in September 2022. The plan outlines how NASA intends to implement the remaining requirements of EO 13960. The plan also identifies the Associate Administrator for OTPS and NASA's Chief Scientist as the Agency's Responsible AI (RAI) officials as well as their roles and responsibilities in implementing EO 13960. Further, it details plans for responding to EO 13960 including capturing use case inventories and reporting them to RAI officials annually, establishing oversight of AI projects to ensure continuous monitoring efforts, and engaging the AI community with educational opportunities and discussions about the Agency's ethical AI standards and how to implement them.

¹⁴ The Agency merged the Office of Strategic Engagements and Assessments and the Office of the Chief Technologist into OTPS.


NASA'S ARTIFICIAL INTELLIGENCE GOVERNANCE EFFORTS ARE PROGRESSING BUT FURTHER REFINEMENT IS NEEDED

NASA has made progress towards establishing an AI governance framework for the Agency's AI and ML efforts by developing the Framework for the Ethical Use of AI and RAI Plan. In addition, the Agency developed an initial AI use case inventory in response to EO requirements. However, NASA has yet to establish an Agency-wide standard AI definition and its ability to identify and track AI and AI spending is inadequate to meet federal requirements. As AI cybersecurity standards evolve, it is imperative for the Agency to best position itself to ensure protection of NASA's AI. Without a clear definition of AI or the ability to categorize AI and track related spending, NASA's efforts to capture the entire spectrum of AI used across the Agency will be incomplete and as a result may hobble its efforts to monitor and protect its AI use.

Lack of a Standard Definition for AI Hinders the Agency's Ability to Effectively Manage Its AI Inventory

NASA has not adopted a standard definition of AI that is consistently utilized throughout the Agency. While there is no universally accepted definition for AI, standard practice among researchers and the AI community is to tailor a definition specific to each organization depending on the scope, risk appetite, internal structure, culture, and implementation details of their AIML efforts. Such a definition is foundational to the effective management of AI and any related agency programs. NASA currently has three separate definitions of AI published within internal documentation. NASA's RAI Plan, published in September 2022, utilizes the definition for AI found within EO 13960. This differs from the definition published within the Framework for the Ethical Use of AI (which preceded the RAI Plan and was published in April 2021). Additionally, NASA's internal AIML SharePoint collaboration website includes yet another definition. The multiple definitions are shown in Figure 5.

Figure 5: NASA's Definitions of AI



Source: NASA OIG representation of Agency information.

While all three definitions are similar in many respects, subtleties and nuances in each can alter whether a particular technology should be considered AI. For example, usage of a word such as "robot" in the RAI Plan's definition and not in the other two can affect whether a use case or technology is reported as AI or as a robotic process automation.¹⁵ Similarly, use of "computerized capabilities" in NASA's AI framework definition versus "computer systems" on the AIML SharePoint site can alter the reporting metric of a piece of computer hardware reported as AI opposed to an AI capability itself. This may lead to duplicate reporting (both the hardware/computer itself and the AI program/software inside reported as AI) or differing reporting depending on the subjective views of the user. As such, we found the lack of a singular definition for AI makes it more difficult for the Agency to fully implement requirements to develop a use case inventory as outlined in EO 13960.

Specifically, NASA provided an initial Agency public use case inventory to the Office of Management and Budget based on the definition provided by the EO (which was later adopted into NASA's RAI Plan). Additionally, the EO stipulated that "use cases" should include:

- both standalone AI and AI embedded within other systems or applications;
- AI developed both by the agency or by third parties on behalf of agencies for the fulfilment of specific agency missions, including relevant data inputs used to train AI and outputs used in support of decision making; and
- agencies' procurement of AI applications.

¹⁵ Robotic process automation is a commercial off-the-shelf technology that can be used to automate repetitive (rules-based) tasks.

These stipulations are not explicitly reiterated across any of NASA's published definitions for AI, including the most recently published RAI Plan. Consequently, this omission amplifies concerns about the accuracy of NASA's reported AI inventory.

NASA submitted its initial AI use case inventory in March 2022, nearly a year after the Agency's Framework for the Ethical Use of AI was published (April 2021) but prior to issuance of the RAI Plan (September 2022). The absence of a standard definition may have impacted the original use case inventory as Center and Mission Directorate personnel that we interviewed stated they reported AI based on their own individual understanding of AI and not based on a formal definition provided by the Agency. When asked to define AI, one Center official stated, "if a computer was being taught to do something without specifying the task," then that should be reported as AI. An official at another Center suggested that individual project managers determine whether a project should be reported as AI. For example, robotic process automation is classified by some as robotics and not AI. Each of these Center officials demonstrated how individual perceptions may affect how a project or technology is reported.

A final version of the public use case inventory completed in October 2022 reported that the number of use cases decreased by 84 percent when compared to the original March 2022 submission. According to NASA officials, this was due to the removal of planned or proposed AI projects within the AI use case inventory. The absence of a standard AI definition at the center of NASA's AI governance efforts hampers the Agency's ability to thoroughly compile and track its inventory, a foundational step for the effective management of AI. Furthermore, it may limit NASA's ability to comprehensively respond to future AI-related federal requirements or inquiries.

NASA's Classification and Tracking of AI Is Inadequate to Fully Address Current and Future Federal Requirements and AI Cybersecurity Concerns

NASA currently lacks a process or mechanism to accurately classify and track AI within the Agency, an issue further compounded by its lack of a standardized AI definition. As a result, it is difficult for the Agency to monitor AI and meet federal requirements, including implementing and responding to EOs. NASA's AI representatives explained that the Agency does not have a singular designation or classification mechanism for tracking or monitoring AI. Generally, AI is not considered its own project, but rather is included as part of larger, individual projects and as such is reviewed and managed in accordance to processes specific to its purpose separate from the AI. For example, AI utilized as part of an engineering project is reviewed and managed in accordance with published engineering regulations and standards. Similarly, AI utilized as part of a scientific project or software development/information technology project is reviewed and managed under guidance specific to those respective areas.

Since AI is typically managed as part of a larger project and is not tracked separately, the Agency's response to EO 13960 to create an AI inventory was accomplished in an ad hoc manner using a data call—rather than an automated process. NASA's initial AI public use case inventory, dated March 2022, took 2 years to compile and was not finalized for public distribution until October 2022. The Agency AIML Lead within the OCIO stated that the process for developing the inventory required individual responses from AI users to a multi-faceted data call, a process that takes significant time to compile, validate, and vet for inclusion as the Agency's public response to the EO.

Similarly, NASA's management of AI as part of larger projects impacted its response to EO 13859, which requires agencies to compile an estimated annual budget for AI expenditures and ensure technical controls exist to minimize the vulnerability of AI from attack by a malicious actor. According to OTPS and OCIO officials, NASA does not have the ability to definitively identify the amount of money spent on AI because the funds are usually associated with larger project costs and no specific classification or identification exists for AI expenditures within the Agency's financial system. Therefore, the Agency is unable to compile the estimated annual budget without an Agency AI data call. By compiling AI inventory and budget data through the manually labor-intensive data call process, the Agency runs the risk of manual errors in addition to the time and resources staff spend working on a process that could be significantly lessened via automation.

With the lack of a classification mechanism or process for specifically identifying AI, the Agency will continue to face challenges accurately responding to federal mandates that require comprehensive reviews of NASA's AI inventories or spending reports—a lack of transparency that could undermine public trust in how NASA invests its funding. Moreover, the lack of a comprehensive inventory also hinders the Agency's ability to effectively monitor and respond to AI specific cybersecurity threats or vulnerabilities.

According to NASA officials, their explanation for not identifying a process specific to managing AI was to limit additional "bureaucratic" procedures within the Agency and promote innovation. The assumption is that existing processes for each department (for example, engineering, scientific, information technology, etc.) should be adequate to track AI inventory and address concerns that may be specific to AI including monitoring requirements and ensuring it is properly safeguarded from cybersecurity vulnerabilities. Our previous work has shown that the Agency's approach to information technology management is fragmented, with numerous separate lines of authority that have long been a defining feature of the environment in which cybersecurity decisions are made, putting NASA's overall cybersecurity posture at a higher-than-necessary risk from cyber threats.¹⁶ By using the same fragmented lines of authority and lacking a classification mechanism for the monitoring of AI in its environment, in our view NASA's AI is exposed to the same higher-than-necessary risk of cyber threats.

Implementing future federal AI cybersecurity controls will prove more difficult since NASA lacks a singular classification for AI within its system of record.¹⁷ The *NIST AI Risk Management Framework* states that risks to software and information-based systems should be applied to AI for risk management purposes. Similarly, GAO's accountability framework suggests that organizations should develop an AI-specific risk management plan to systematically identify, analyze, and mitigate cybersecurity risks. These best practices and guidance suggest that future actions will likely be introduced at the federal level.


The Agency's ability to implement, assess, and monitor cybersecurity controls within its system(s) of record currently utilizes the traditional Assessment and Authorization process, which is comprised of the six key tasks shown in Figure 6. The process consists of a review of security policies and procedures (management controls); physical facility infrastructure (operational controls); and network testing, server testing, application security testing, penetration testing, and scanning (technical controls).

¹⁶ NASA Office of Inspector General, *NASA's Cybersecurity Readiness* ([IG-21-019](#), May 18, 2021).

¹⁷ A cybersecurity control is a safeguard or countermeasure prescribed for an information system or an organization designed to protect the confidentiality, integrity, and availability of its information and to meet defined security requirements.

Assessment and Authorization end products include an authorization to operate the information technology system, risk-based decisions on the application of individual controls, and a plan of action and milestones to address identified deficiencies.

Figure 6: NASA Assessment and Authorization Process



Source: NASA OIG depiction of the Assessment and Authorization process.

For the “Categorize System” function, the Agency’s lack of an AI specific classification mechanism in its system(s) of record would directly encumber any federal requirements for the adoption of cybersecurity controls or requirements. This is critical as the system categorization determines the necessary actions for the rest of the Assessment and Authorization process. NIST’s *AI Risk Management Framework*, for example, suggests that agencies can benefit from applying the framework through enhanced processes for governing, mapping, measuring, and managing AI risk, and clearly documenting outcomes. But the application of this or any future framework is stymied by the Agency’s lack of an AI classification process because it is nearly impossible to map and govern a technology without the appropriate categorization and classification.

CONCLUSION

NASA continues to prove itself a leader in innovation within the federal government and leads federal agencies in AI use as shown in Figure 3. However, as AI applications continue to evolve federal regulations also will likely evolve to keep pace. As such, it is imperative that NASA standardizes an Agency definition for AI and implements an overall classification and monitoring process. Failure to do so will leave the Agency unable to effectively track and safeguard its AI. With comprehensive oversight and required reporting of its AI projects incomplete, the financial resources spent on their usage lack transparency and traceability. The Agency must ensure it is not only capable of responding to current federal requirements but is also positioned to implement future requirements. As the use of AI at NASA expands, the Agency will need to establish a strong governance foundation now to effectively manage AI in the future.

RECOMMENDATIONS, MANAGEMENT'S RESPONSE, AND OUR EVALUATION

To improve the governance, budgeting, and cybersecurity of NASA's AI capabilities, we recommended the Associate Administrator for Technology, Policy, and Strategy; Chief Scientist; Chief Information Officer; and Chief Financial Officer:

1. Establish a standardized definition for AI within the Agency, to include harmonizing the definitions in the *NASA Framework for the Ethical Use of Artificial Intelligence (AI)*, *NASA's Responsible AI Plan*, and NASA AIML SharePoint.
2. Ensure the standardized AI definition is used to identify, update, and maintain the Agency's AI use case inventory.
3. Identify a classification mechanism to assist in the rapid application of federal requirements for cybersecurity controls and monitoring practices.
4. Develop a method to track budgets and expenditures for AI use case inventory.

We provided a draft of this report to NASA management who concurred or partially concurred with our recommendations and described planned actions to address them. We consider management's comments responsive; therefore, the recommendations are resolved and will be closed upon completion and verification of the proposed corrective actions.

Major contributors to this report include Tekla Colón, Mission Support Audits Director; Scott Riggensbach, Assistant Director; Anu Bakshi; Joseph Cook; and Vincent Whitfield. Lauren Suls provided editorial and graphics support.

If you have questions about this report or wish to comment on the quality or usefulness of this report, contact Laurence Hawkins, Audit Operations and Quality Assurance Director, at 202-358-1543 or laurence.b.hawkins@nasa.gov.

**PAUL
MARTIN**

Digitally signed by
PAUL MARTIN
Date: 2023.05.02
13:21:27 -04'00'

Paul K. Martin
Inspector General

APPENDIX A: SCOPE AND METHODOLOGY

We performed this audit from June 2022 through March 2023 in accordance with generally accepted government auditing standards. Those standards require that we plan and perform the audit to obtain sufficient, appropriate evidence to provide a reasonable basis for our findings and conclusions based on our audit objectives. We believe that the evidence obtained provides a reasonable basis for our findings and conclusions based on our audit objectives.

The scope of our audit encompassed the management of NASA's current AI program as well as its capability to aptly monitor proposed policies and regulations that may soon be implemented at the federal level. During our audit we (1) identified and reviewed existing criteria (regulations, policies, plans and procedures); (2) reviewed NIST and GAO frameworks for comparative purposes to current practices; (3) assessed NASA's reporting efforts towards requirements published in the November 17, 2020, Office of Management and Budget M-21-06, *Guidance for Regulation of Artificial Intelligence Applications*, and EOs 13859 and 13960; and (4) assessed NASA's AI cybersecurity hygiene practices.

The objective of our audit was to determine the Agency's progress in developing its AI governance framework and standards and assess whether security controls are being considered and implemented to protect AI data and technologies from cyber threats. We reviewed applicable federal and NASA policies, regulations, frameworks, and industry best practices for AI governance. We interviewed responsible NASA officials from OTPS and OCIO, including OCIO's Cybersecurity & Privacy Division, and met with representatives from three NASA Centers to discuss their specific AI practices. Additionally, we met with officials directing AI efforts in other government agencies to benchmark NASA's progress in implementing federal AI guidance and discuss best practices.

Assessment of Data Reliability

The data used in our audit was primarily in the form of a publicly available use case inventory of federal agencies, as well as an audit-prompted inventory provided by the Agency to GAO. The NASA Office of Inspector General did not perform an independent data reliability assessment and instead relied on the AI inventories as presented. Due to the lack of a centralized system housing the AI data, we were unable to reproduce the inventory. Per the standards outlined in the GAO publication *Assessing Data Reliability*, we determined that the reliability of the data needed to support the findings, conclusions, or recommendations in the context of the audit objectives is satisfactory.

Review of Internal Controls

We assessed internal controls and compliance with laws and regulations necessary to satisfy the audit objectives of determining NASA's progress in developing its AI governance framework and associated security considerations. Control weaknesses are identified and discussed in this report. Our recommendations, if implemented, will improve those identified weaknesses.

Prior Coverage

During the last 5 years, the NASA Office of Inspector General has issued one report containing some relevance to the subject of this audit. Additionally, GAO has issued two ancillary reports of interest to this topic. Unrestricted reports can be accessed at <https://oig.nasa.gov/> and <https://www.gao.gov/>, respectively.

NASA Office of Inspector General

NASA's Cybersecurity Readiness ([IG-21-019](#), May 18, 2021)

Government Accountability Office

Artificial Intelligence: DOD Should Improve Strategies, Inventory Process, and Collaboration Guidance ([GAO-22-105834](#), March 30, 2022)

Artificial Intelligence: Status of Developing and Acquiring Capabilities for Weapons Systems ([GAO-22-104765](#), February 17, 2022)

APPENDIX B: MANAGEMENT'S COMMENTS

National Aeronautics and Space Administration

Mary W. Jackson NASA Headquarters
Washington, DC 20546-0001



April 26, 2023

Reply to Attn of: Office of the Chief Information Officer

TO: Assistant Inspector General for Audits

FROM: Chief Information Officer
Chief Financial Officer
Chief Scientist
Associate Administrator for Technology, Policy, and Strategy

SUBJECT: Agency Response to OIG Draft Report, "NASA's Management of Its Artificial Intelligence Capabilities" (A-22-12-00-MSD)

The National Aeronautics and Space Administration (NASA) appreciates the opportunity to review and comment on the Office of Inspector General (OIG) draft report entitled, "NASA's Management of Its Artificial Intelligence Capabilities" (A-22-12-00-MSD), dated March 23, 2023.

The OIG found that NASA has made progress in establishing an Artificial Intelligence (AI) framework through development of the NASA Framework for the Ethical Use of Artificial Intelligence in April 2021. However, NASA has not adopted a standard definition of AI and instead has three separate definitions: one in the NASA Framework for the Ethical Use of AI, one in NASA's Responsible AI Plan that utilizes the definition found within Executive Order (EO) 13960, and one on NASA's internal Artificial Intelligence Machine Learning (AIML) SharePoint collaboration website.

In the draft report, the OIG makes four recommendations addressed to the Associate Administrator for Technology, Policy, and Strategy, Chief Scientist, Chief Information Officer, and Chief Financial Officer intended to improve the governance, budgeting, and cybersecurity of NASA's AI capabilities.

Specifically, the OIG recommends the following:

Recommendation 1: Establish a standardized definition for AI within the Agency, to include harmonizing the definitions in the *NASA Framework for the Ethical Use of Artificial Intelligence (AI)*, *NASA's Responsible AI Plan*, and NASA AIML SharePoint.

Management's Response: NASA partially concurs. While NASA will endeavor to align its definitions across the Agency, other agencies have the lead with respect to AI for the Federal Government, particularly the National Institute of Standards and

Technology for standard definitions. NASA anticipates that such definitions will be forthcoming. When received, NASA will develop a unified definition across the Agency, and incorporate that definition within its AI policy.

Estimated Completion Date: April 30, 2024

Recommendation 2: Ensure the standardized AI definition is used to identify, update, and maintain the Agency's AI use case inventory.

Management's Response: NASA concurs. Once a standardized definition is in place in NASA policy it will be used in updating our use case inventory.

Estimated Completion Date: July 31, 2024

Recommendation 3: Identify a classification mechanism to assist in the rapid application of federal requirements for cybersecurity controls and monitoring practices.

Management's Response: NASA concurs. The existing NASA AI inventory contains multiple fields, such as "specific AI techniques," which can be used to identify relevant practitioners and issue alerts for emergent bugs, issues, or risks including cybersecurity threats. The inventory can be updated at any time as a living resource and will also leverage a yearly inventory update. NASA will further refine classification/identification processes to respond to bugs, issues, or risks, and will ensure the AIML community of practice is adequately engaged.

Estimated Completion Date: July 31, 2023

Recommendation 4: Develop a method to track budgets and expenditures for AI use case inventory.

Management's Response: NASA partially concurs. The Office of the Chief Financial Officer is in compliance with the data call referenced in EO 13859 to the Networking and Information Technology Research and Development (NITRD) Program (see latest report at <https://www.nitrd.gov/pubs/FY2022-NITRD-NAIO-Supplement.pdf>) and will participate fully in future NITRD data calls on AI research and development spending after the annual President's Budget Justification is submitted to Congress, in accordance with the EO. As the Federal Government and NASA refine AI definitional issues, we will continue to evolve our cost reporting accordingly.

Estimated Completion Date: Not Applicable.

We have reviewed the draft report for information that should not be publicly released. As a result of this review, we have not identified any information that should not be publicly released.


Once again, thank you for the opportunity to review and comment on the subject draft report. If you have any questions or require additional information regarding this response, please contact Jeremy Yagle at (757) 864-9622.

 Digitally signed by
ROBERTBINKLEY
Date: 2023.04.26
16:20:40 -04'00'


Jeffrey Seaton
Chief Information Officer

 Date:
2023.04.27
10:29:11 -04'00'

Margaret Vo Schaus
Chief Financial Officer

 Digitally signed by Bhavya
Lal
Date: 2023.04.27 10:54:58
-04'00'

Bhavya Lal
Associate Administrator for
Technology, Policy, and Strategy

 Digitally signed by
Katherine Calvin
Date: 2023.04.27
09:16:30 +07'00'

Dr. Katherine Calvin
Chief Scientist

APPENDIX C: REPORT DISTRIBUTION

National Aeronautics and Space Administration

Administrator
 Deputy Administrator
 Associate Administrator
 Associate Administrator for Technology, Policy, and Strategy
 Chief of Staff
 Chief Scientist
 Chief Information Officer
 Chief Financial Officer

Non-NASA Organizations and Individuals

Office of Management and Budget
 Deputy Associate Director, Climate, Energy, Environment and Science Division
 Government Accountability Office
 Director, Contracting and National Security Acquisitions

Congressional Committees and Subcommittees, Chairman and Ranking Member

Senate Committee on Appropriations
 Subcommittee on Commerce, Justice, Science, and Related Agencies
 Senate Committee on Commerce, Science, and Transportation
 Subcommittee on Space and Science
 Senate Committee on Homeland Security and Governmental Affairs
 House Committee on Appropriations
 Subcommittee on Commerce, Justice, Science, and Related Agencies
 House Committee on Oversight and Accountability
 Subcommittee on Government Operations and the Federal Workforce
 House Committee on Science, Space, and Technology
 Subcommittee on Investigations and Oversight
 Subcommittee on Space and Aeronautics

(Assignment No. A-22-12-00-MSD)