

ATTACHMENT M

COMMERCIAL ARTIFICIAL INTELLIGENCE RISK AND IMPACT ASSESSMENT (CAIRIA)

SEPTEMBER 2025

CONTRACT TBD

1. DRD Title: Commercial Artificial Intelligence Risk and Impact Assessment (CAIRIA)		
2. DRD No.: 4	3. Data Type: 1	4. OPR: OCIO
5. Solicitation No.: 80GSFC27R0015		6. Contract No.: TBD
7. Date Issued: 09/30/2025	8. Date Revised: N/A	9. DRD Category: Technical ☒ Administrative ☐ S&MA ☐
10. Description/Use: To ensure that requirements for Artificial Intelligence (AI) impact and risk assessments are met for all AI systems utilized during work associated with this contract.		
11. Distribution: All deliverables required by this DRD shall be submitted electronically to the CO, COR, and NASA's Chief Artificial Intelligence Officer (CAIO) at agency-caio@mail.nasa.gov . Approval will be provided by the CAIO. Initial Submission: Contract Phase-In and/or Prior to Use of new AI components or capabilities to perform contract performance Submission Frequency: Updated annually after initial submission and when new AI components or capabilities are introduced. Format: Unless otherwise directed, the data requested in this DRD shall be delivered to the Government in soft-copy via an electronic transfer mechanism (e.g., electronic mail, flash drive, or file transfer protocol) in a format readable by a Government device utilizing the standards in NASA-STD-2804, "Minimum Interoperability Software Suite." Interrelationship: SOW Sections II.A.1, II.A.2, II.B.1, II.D.1, II.E.3 Applicable Documents: Executive Order 13960, Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government Executive Order 14179, Removing Barriers to American Leadership in Artificial Intelligence OMB Memorandum M-15-14, Management and Oversight of Federal Information Technology. OMB Memorandum M-25-21, Accelerating Federal use of AI through Innovation, Governance, and Public Trust. OMB Memorandum M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government. Scope: AI is defined as Artificial Intelligence is defined in Section 238(g) of the John S. McCain National Defense Authorization Act for Fiscal Year 2019. “(g) ARTIFICIAL INTELLIGENCE DEFINED. —In this section, the term “artificial intelligence” includes the following: (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.”

OMB Memorandum M-25-21 provides this additional technical context:

1. This definition of AI encompasses, but is not limited to, the AI technical subfields of machine learning (including deep learning as well as supervised, unsupervised, and semi-supervised approaches), reinforcement learning, transfer learning, and generative AI.
2. This definition of AI does not include robotic process automation or other systems whose behavior is defined only by human-defined rules or that learn solely by repeating an observed practice exactly as it was conducted.
3. For this definition, no system should be considered too simple to qualify as covered AI due to a lack of technical complexity (e.g., the smaller number of parameters in a model, the type of model, or the amount of data used for training purposes).
4. This definition includes systems that are fully autonomous, partially autonomous, and not autonomous, and it includes systems that operate both with and without human oversight.

AI Systems include the following capabilities:

CLASSIFICATION	DESCRIPTION	EXAMPLES
Generative AI	AI that generates new or synthetic content (e.g., images, videos, audio, text, code).	Chatbots, image generation, code completion, synthetic media, data augmentation, retrieval-augmented generation (RAG), deep research
Agentic AI	AI systems that perform tasks or make decisions autonomously with minimal human intervention.	Task automation, personal assistants
Classical/Predictive Machine Learning	Models trained on data to make predictions or classifications based on identified patterns or relationships.	Sales forecasting, weather forecasting, spam detection, recommendation systems, credit/fraud scoring, anomaly detection
Computer Vision	AI that processes and interprets visual data (e.g., images and videos).	Object detection, facial recognition, biometric identification, medical imaging, image segmentation
Natural Language Processing (NLP)	AI that processes, interprets, and shares information in human language.	Sentiment analysis, speech-to-text, text-to-speech
Reinforcement Learning	AI trained through trial and error using rewards and penalties to optimize decision-making policies.	Robotics operating in physical environments, strategic game play, autonomous systems (drones)

Contents:

Baseline and High-Impact AI Risk and Impact Assessments are included below. Formats and reporting processes and procedures will be provided annually by the CAIO based upon OMB and Agency requirements. This is a dynamic process, and formats, processes, procedures, and structures are subject to change.

Definitions

Baseline AI: Any AI product, tool, software, service, or system that does not meet the definition of “high-impact AI” is considered a **Baseline AI** system.

High-Impact AI: A system is considered **high-impact AI** if it meets the criteria in OMB M-25-21 section 4(a), which states that high-impact AI is any AI model or system with an output that serves as a principal basis for decisions or actions with legal, material, binding, or significant effect on:

1. an individual or entity's civil rights, civil liberties, or privacy; or
2. an individual or entity's access to education, housing, insurance, credit, employment, and other programs;
3. an individual or entity's access to critical government resources or services;
4. human health and safety;
5. critical infrastructure or public safety; or
6. strategic assets or resources, including high-value property and information marked as sensitive or classified by the Federal Government.

For the purposes of AI development and deployment at NASA, this includes AI systems that:

1. Perform or significantly support government decisions affecting rights, safety, benefits, entitlements, or protections
2. Pose significant risks of harm to individuals, groups, or public trust
3. Operate with a high level of autonomy or limited human oversight
4. Are used in safety-critical or mission-critical systems

Baseline AI Risk and Impact Assessment:

For any AI product, tool, software, service, or system considered as a baseline AI system, the contractor shall provide, at a minimum, plain-language documentation of the following elements:

1. The AI system's function and role.
2. Any automated decision-making involved.
3. All data sources and types used to develop or operate the system.
4. Characterization of model type, autonomy level, and model update mechanisms (e.g., clustering, reinforcement learning, deep neural networks, large language models, computer vision models, time-series forecasting models, recommendation systems). Also include the AI Classification from the table included above.
5. Known risks or limitations (e.g., performance boundaries, error rates, fairness or bias concerns, data quality limitations, explainability challenges, adversarial vulnerabilities, model drift, hallucination risks).
6. Human review or override mechanisms, if applicable.
7. Information on open-source or third-party models or datasets including data shared with or received from external systems or models.
8. Process for testing, validation, and performance monitoring procedures, including:
 - o Key performance measures or acceptance criteria
 - o Security and robustness testing (e.g., red-teaming, adversarial input resistance)
 - o Whether testing is manual, automated, or third-party validated
 - o Ongoing monitoring or alerting for degradation or unexpected behavior

High-Impact AI Risk and Impact Assessment:

For high-impact AI systems, the contractor shall provide, **in addition to the Baseline AI Risk and Impact Assessment**, plain-language documentation of the following elements:

1. Assessment of potential impacts on individuals' rights, safety, access to government services, or well-being.
2. Identification of all known or reasonably foreseeable risks **and the contractor's corresponding mitigation strategies**, including safeguards, operational controls, monitoring mechanisms, and corrective action procedures.
3. Process for testing, validation, and performance monitoring procedures, including:
 - a. Defined performance measures and system acceptance criteria.
 - b. Security and robustness testing (e.g., red-teaming, adversarial input resistance).
 - c. Identification of whether testing is manual, automated, or third-party validated.
 - d. Ongoing monitoring procedures and alerting for system degradation or unexpected behavior.
 - e. Iterative fairness and bias audits, including evaluation across affected populations.
 - f. Description of the testing objectives and the operational scope covered.
 - g. Conditions that trigger model re-testing or re-validation during the system lifecycle.
 - h. Methods used to document test results, failures, and corrective actions.
 - i. Evidence of independent review, validation, or audit readiness where applicable.
 - j. Approach to versioning and traceability of testing and model changes.
4. Transparency and user notice plans, including intended disclosures and user-facing explanations (if applicable).

Upon the Government's request, the contractor must provide any additional information to support internal AI risk tracking and lifecycle management. If the AI system's use, function, or impact materially changes during performance, the contractor may be required to support reassessment, including possible reclassification as a high-impact AI system.

Remarks: None.

Maintenance: CAIRIA plans shall updated annually, or as new AI components or capabilities are introduced.