



AI in Flight: Engineering Artificial Intelligence into Human Spaceflight



Dave Marquette, JSC S&MA



About the Presenter



- AI work began as JSC CIO Technical Assistant for HSF
- Graduate education in mechanical engineering
- 15 years of mission operations experience
- 6 years software development experience
- 9 years IT architecture experience
- 12 years of cybersecurity experience
- 3 years CIO executive leadership experience
- **2025 graduate of University of Texas program in AI/ML**

“AI? It’s just more complicated software, right?”



Origin



- Commercial LEO Development Program solicited OCIO guidance for AI uses in commercial spacecraft proposals (Sep 2024)
- JSC OCIO Technical Assistant for Human Spaceflight consulted NASA's Chief Data & AI Officer, & created "AI in Flight" group
 - Named to convey its distinction as *AI in high-risk applications* (human-rated spacecraft)
 - Included CLDP & other human spaceflight programs (ISS, EHP, Gateway)
 - Added technical experts from OCE, OSMA, ARMO, SOMD, STMD, & JSC Engineering



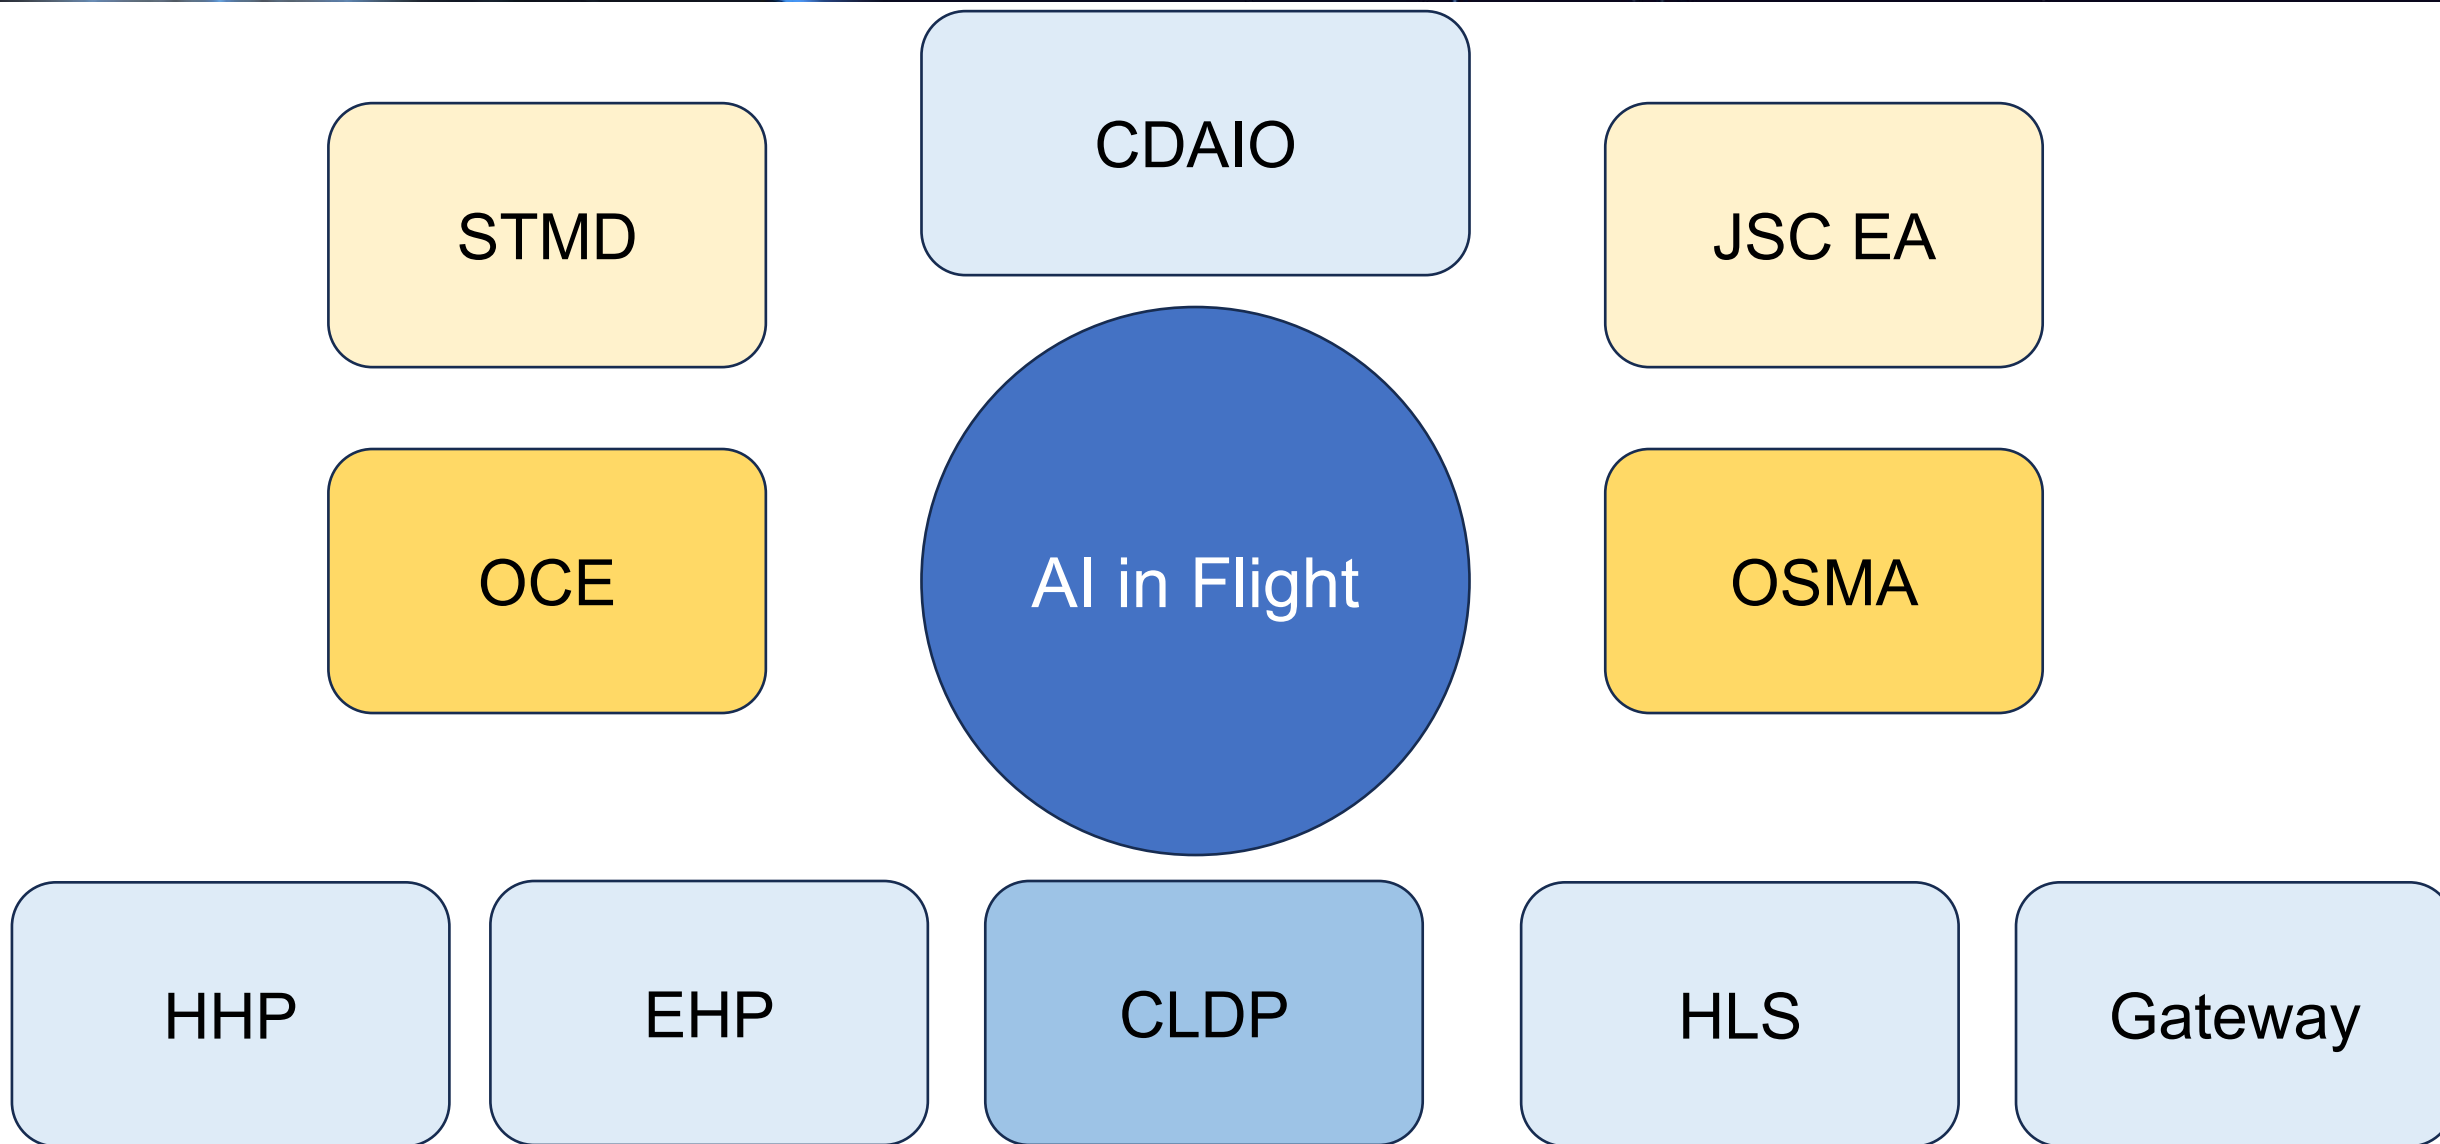
Approach



- Refined scope to Class A, safety-critical software in human-rated spacecraft/spaceflight
 - Primarily machine learning (adaptive) types of AI
- Reviewed all existing AI policy: US Law, Executive Orders, OMB Memos, NASA NPDs & NPRs & Handbooks, etc.
 - Identify holes & hurdles of NASA policy
- Researched applicable industry AI standards (ISO, NIST)
- Studied AI strategies of spaceflight's closest analog, aviation
- Identified promising space use cases
- Completed personal, hands-on training



Organizational Coordination





Charter Questions



1. Are the technologies that companies are proposing for commercial LEO spacecraft considered *artificial intelligence*?
2. Are existing engineering/safety requirements & processes suitable & sufficient for managing *onboard AI*?
3. If not, what additional requirements need to be formulated—and conveyed via contracts—to minimize risks associated with onboard AI?



Official Definitions



Webster dictionary:

The capability of computer systems or algorithms to imitate intelligent human behavior.

Oxford dictionary:

Computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.

NIST:

An engineered or machine-based system that can, for a given set of objectives, generate outputs such as **predictions, recommendations, or decisions** influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy.

ISO:

“a technical and scientific field devoted to the engineered system that generates outputs such as **content, forecasts, recommendations or decisions** for a given set of human-defined objectives”

NASA:

Any computerized capability to **perceive, reason, learn, & act.**”

From DoD AI Strategy of 2018:

“ability of machines to perform tasks that normally require human intelligence — for example, recognizing patterns, **learning from experience**, drawing conclusions, making decisions, or taking action — whether digitally or as the smart software behind autonomous physical systems.”

National AI Initiative Act of 2020:

“a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to (A) perceive real or virtual environments, (B) abstract such perceptions into models through analysis in an automated manner, and (C) use model **inference to formulate** options for information or action.”



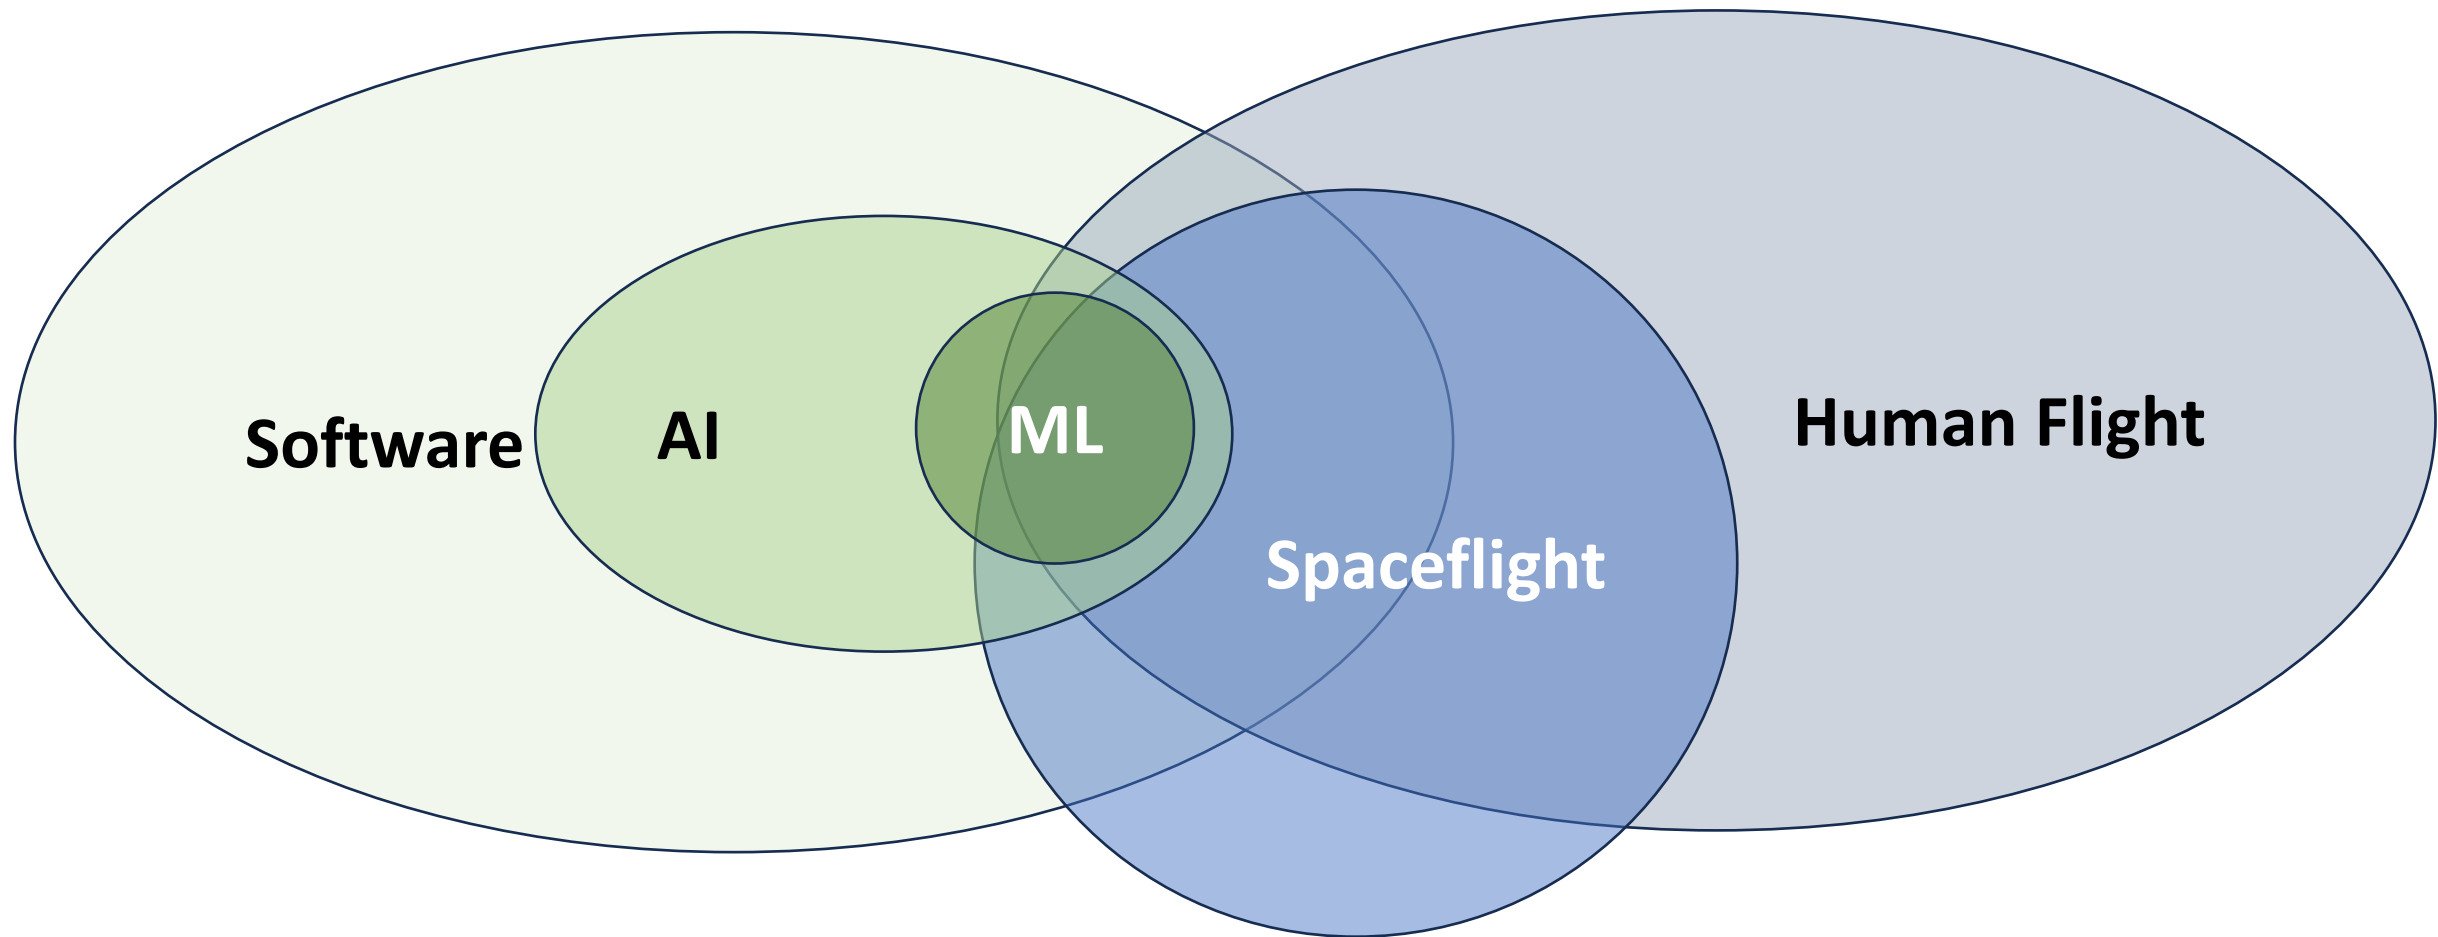
AI Fundamentals



- In essence, AI is mimicking human intelligence, Google-defined as:
 - “The ability [of a non-human entity] to acquire and apply knowledge and skills”— i.e., **to learn or adapt**
- AI & ML are not new, but they are new to human spaceflight
- The most fundamental parts of AI:
 - **Patterns** → **Predictions** → **Progression**
- The essential, powerful difference:
 - Traditional software: the algorithm manipulates the data,*
 - AIML systems: the data manipulates the algorithm*



NASA's Highest-Risk Application of AI: Human Spaceflight






Evolutionary Context



We are very far from both epochs represented below; much more advanced than one, much more primitive than the other:



ML represents another evolutionary step, as we can grow from very static, unchangeable systems to more agile, resilient, adaptable systems.



Goal



Our shared goal across disciplines, centers, program offices, mission directorates:

AI in Flight's Goal:

Responsible Innovation

Determine how to ensure crew **safety**, vehicle integrity, & mission success while realizing benefits of *learning* AI technology to accomplish that mission more **safely**, effectively, and efficiently.

To remain focused, this forum's scope does **not** include terrestrial AI or uncrewed spacecraft, as these environments have radically different risk postures.

Important intangibles include:

- building relationships
- sharing knowledge
- infusing AIML vernacular
- increasing stakeholder familiarity



Organizational Coordination





TECHNOLOGY INFUSION

(STMD/OCE)
new ML TRLs
NPR 7123, SP-20205003605

ENGINEERING / ASSURANCE

OCE/OSMA
new NIST 100-1
NPR 7150, STD 8739

 **Intelligent
Spacecraft** 

SYSTEM SECURITY

OCIO (CDAIO/CSPD)
new NIST 800-218A
NPR 2810, TM-20210012886?

HUMAN-RATING IMPLEMENTATION

HSF Programs
NPR 8705/STD 8719/CLDP-REQ-1130



Notable Guest Speakers



- OSMA's Technical Fellow for Software Assurance, Tim Crumbley
 - Introducing AI section 8.25 of the Software Engineering Handbook, "AI & Software Assurance"
- CLDP Systems Engineer, Conor Creagh
 - Current approach for AI adoption in CLDP human-rated spacecraft
- OCE's Deputy Technical Fellow for Software, Scott Tashakkor
 - Q & A session on SWEHB section 7.25, "AI and Software Engineering"
 - Current state of the art of space-based infrastructure
- FAA's Chief Scientific and Technical Advisor for Artificial Intelligence – Machine Learning, Dr. Trung Pham
 - Technical certification measures for AI on aircraft
- Chief Data & AI Officer, David Salvagnini/Krista Kinnard
 - Introducing the CDAIO



Notable Guest Speakers



- ARMD's Paul Nelson/Adham Jaber
 - will present the current FAA strategy for implementing AI on aircraft
- Drs. Roshan Patel, Shean Phelps, & Dave Hilmers (JSC Human Health & Performance)
 - AI uses in space medicine
- UCLA Chief Data & AI Officer, Chris Mattman (formerly JPL)
 - Technology Readiness Levels for Machine Learning systems
- LHP Engineering
 - Bridging Safety, AI, & Software Development
- JPL's Stephen Chien
 - AI/ML on interplanetary spacecraft



Unexpected FAA Conclusion



- FAA example: Aircraft Collision Avoidance System now in use
- “While considering the importance of the safety assurance of AI, it became apparent that there are **significant opportunities to use AI for safety**”
- “AI has the potential to add tremendous value in identifying safety risks and increasing operational and business process efficiency.”



NPR 7150 Findings



A deep analysis of NPR 7150 (pub. 3/2022) uncovered potential gaps in requirements. Some classes of these:

- Data curation & validation
- Hyperparameter tuning (batch size, epochs, learning rate, etc.)
- Model drift prevention (e.g., introducing/amplifying bias, esp. over long durations)
- Degree of model autonomy (human feedback, human override, etc.)
- Model behavior understanding & replication
- Critical (competing) processing protections
- Internal redundancy

In many cases, there is a general requirement for something, but not an AI-specific requirement (such as “performance monitoring” in a general sense only)



NPR 7150: A *Prohibitive* Example



There are also requirements that do not apply to AI/ML:

- 3.7.5 If a project has safety-critical software, the project manager shall ensure all identified safety-critical software components have a **cyclomatic complexity value of 15 or lower**. Any exceedance shall be reviewed and waived with rationale by the project manager or technical approval authority. [SWE-220]
 - Note: Cyclomatic complexity is a metric used to measure the complexity of a software program. This metric measures independent paths through the source code. The point of the requirement is to minimize risk, minimize testing, and increase reliability associated with safety-critical software code components, thus reducing the chance of software failure during a hazardous event. The software developer should assess all software safety-critical components with a cyclomatic complexity score over 15 for testability, maintainability, and code quality. For more guidance on this requirement, see NASA-HDBK-2203.



A Comprehensive Approach



While existing processes are being promulgated over the next few years, here is a 3-pronged, interconnected approach for AI/ML development in human-rated spacecraft to balance safety & innovation:

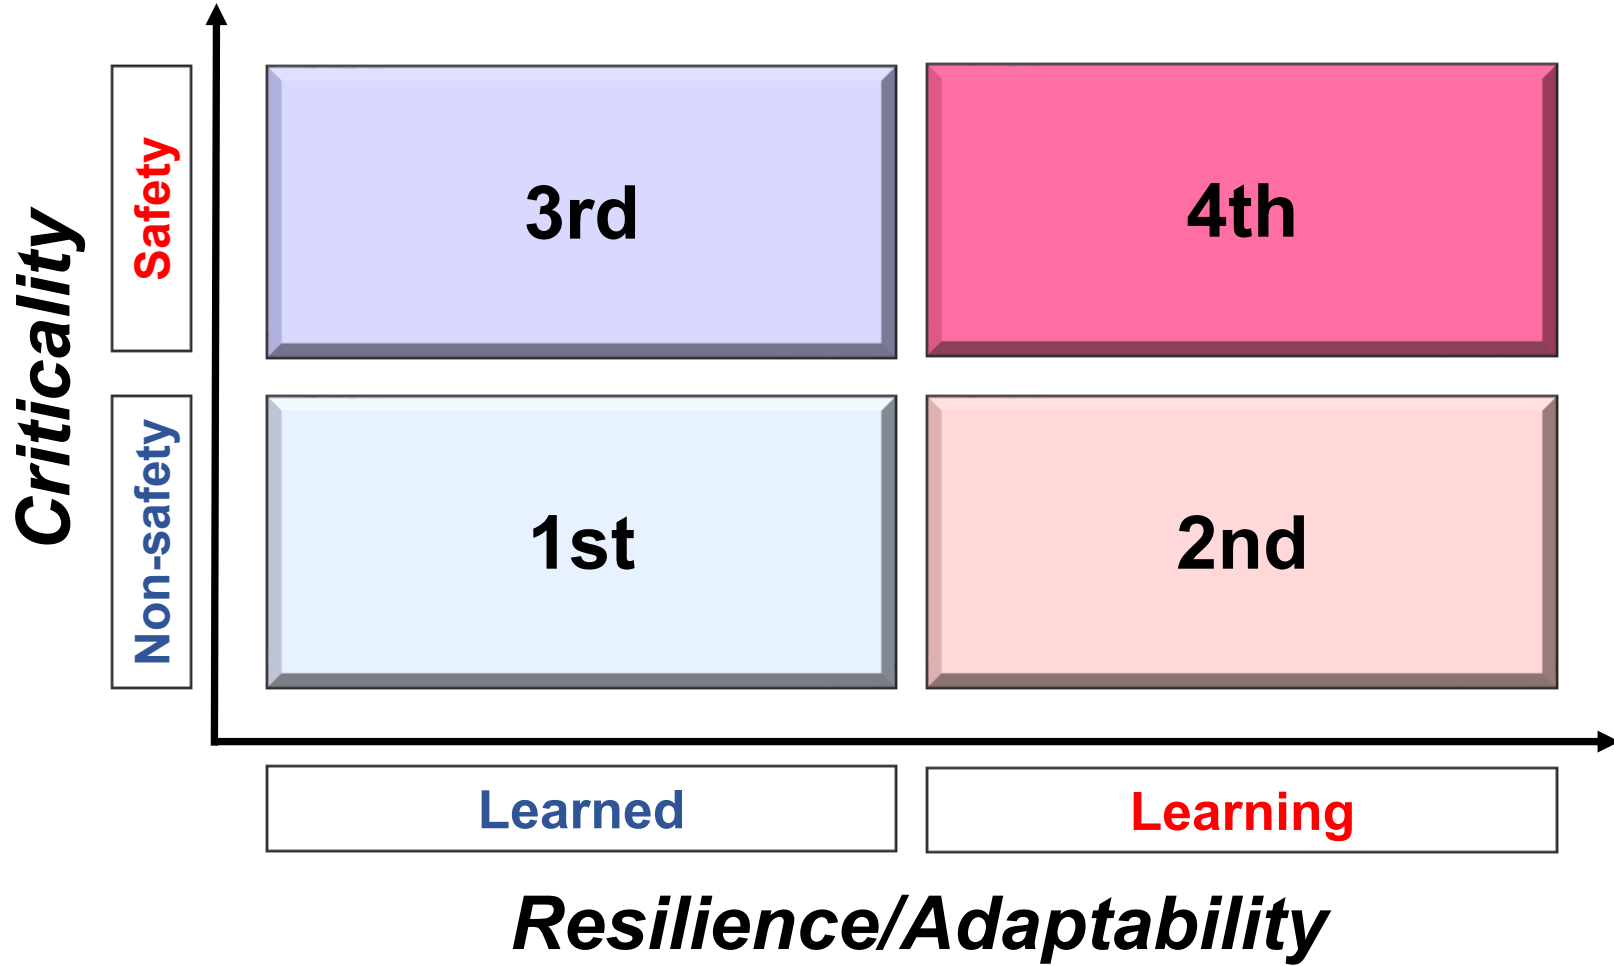
- 1. Risk reduction:** Select/accept “learning AI” (adaptable types) in *low-risk* applications first, moving into other ones as maturity, reliability, & trust improve.
- 2. Technology maturation:** For each application, leverage proven TRL paradigm (of NPR 7123) to design, build, & test the technology (Note: Level 6 typically required by PDR)
- 3. Organizational maturity:** Require suppliers to follow NIST’s “AI Risk Management Framework”—tailored as required—to ensure rigorous, comprehensive *organizational* maturity.



AI Implementation Risk Reduction



Evolving smartly along two dimensions of risk:



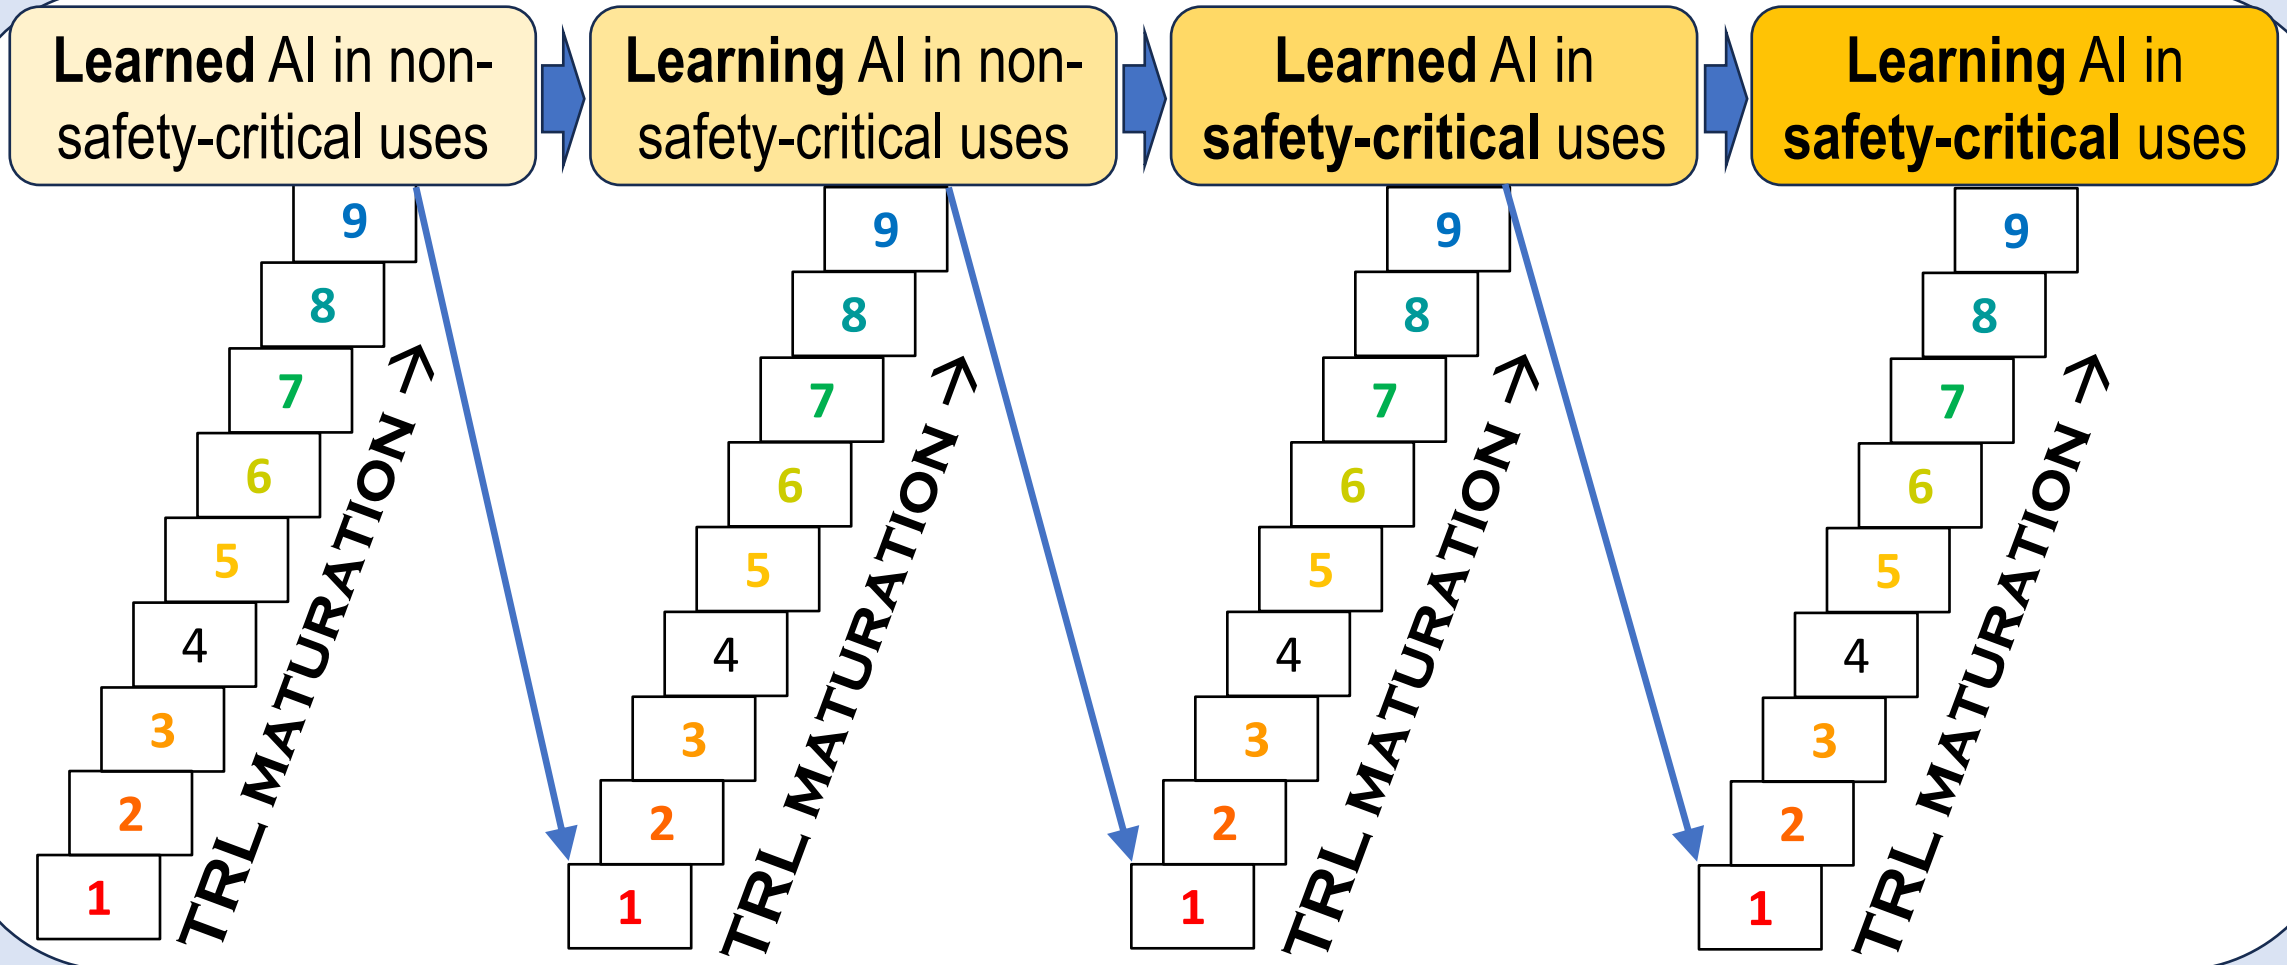
A possible 3rd dimension: near-Earth uses first



Technology Maturity Evolution



← Organizational Processes for AI Engineering →





Recent Efforts



- First review of human-spaceflight AI application (12/25)
 - Informally vetted ISS Axiom payload GRaceful Architecture for Mitigation of System failures (GRAMS)
- Joint SOMD/OSMA development of AI approval framework
 - Alternate tracks to (1) foster innovation & (2) ensure safety
- NASA AI Governance draft
 - Center-based AI Senior Officials
- AI Risk Quick Reference Guide
 - Documented for Mary Skow, ARMO
- Comm LEO (CLDP) AI Risk Analysis
 - Possible new programmatic requirements for commercial spacecraft



Recent Efforts



- Draft bare minimum AI/ML performance standards for Human Spaceflight
 - To bridge to NPR 7150 revisions expected in 2027/2028
- SOMD input for SMD RFI for industry AI concepts
 - Dynamic, individualized medical diagnoses/treatments for in-flight physiological changes
- JSC Engineering's LLM guidance for Class A systems
- *Cybersecurity* standards for AI
 - CSET STD-2661
- Human Research Program briefed on 12/25
- OSMA briefed on 2/11
- Chief AI Officer symposium in planning



Central Claims



- AI has the potential to improve the safety & success of human spaceflight
- Due to its inherent differences, it also poses new challenges & introduces new risks
- Existing NASA processes have not been adapted to account for these differences
- To ensure crew safety & mission success, three general philosophies emerge:

Existing engineering & safety processes are adequate to safely implement AI on human spacecraft.

New processes need to be determined to minimize the risks of AI in human spaceflight while leveraging its benefits.

AI introduces too much risk and should be prohibited from all human spacecraft.



Overarching Desire



- However we proceed, the primary desire for AI is this progression:
 - Valid, thorough **testing** to verify...
 - AI **consistency & accuracy**, to provide...
 - **Predictability**, to ensure...
 - **Reliability**, to foster...
 - **Confidence**, to enhance...
 - ***Safety & success***




Conclusions



- AI systems do not operate like traditional, static software
- As such, AI systems offer significant new capabilities & pose new risks
 - Human spaceflight represents the design case for AI (i.e., highest risk)
- While NASA & commercial interest in AI is high, resource commitment, depth of expertise, formal direction, & technical requirements are currently insufficient to reduce AI's risk to acceptable levels.



BACKUP INFO





AI/ML Requirement Classes*



Performance		Construction	Operations	Security
<i>Classification</i>	<i>Regression</i>	Data Curation	Monitoring	
Precision/Recall/ Accuracy	R-Squared	Pedigree (open-source reliance)	Replication (algorithm & data)	
Computational cost		Compatibility	Interventions (override, disablement)	
Long-term drift		Opacity vs. Insight	Recovery	

*There are also some existing requirements (for traditional code) that will need to be explicitly waived.



Data criticality: How do you perform sound data curation?

- How did you prevent any model bias?
- Or imbalanced data sets?
- Was any synthetic data used?
- How did you handle:
 - Outliers? [invalid data]
 - Missing values? [imputation]
 - Duplicates? [unintended influence]
 - Non-contributing variables/features? [non-descriptive influence]
 - Widely varying data scales? [feature dominance]
- How did you ensure no “data leakage” between data sets?
 - Training, validation, testing sets
 - Over-fitting & under-fitting



For model construction:

- What platform specifications does the model require?
 - GPUs, parallel processing, code versions, etc.
- How does CLDP certify AI's foundational code libraries?
- Were alternative models evaluated & what were the selection criteria?
 - Were ensemble models used?
 - Were convolutional, recurrent, transformer models used? Do we care?
- How were the model & its hyperparameters tuned?
 - What ranges were used for learning rate, tree depth & endpoints, etc.?



For model (“system”) performance:

- How did you measure the AI model's performance/accuracy?
 - AI industry uses established statistical metrics (“R-Squared”, “F1 Score,” Mean Squared Error, etc.)
 - What value-threshold was used in validation?
 - How did you guard against “overfitting”?
- How do you meet Executive Order/OMB/NIST requirements for model transparency & explainability?
- Most importantly, is the model built to learn in-flight or only on the ground?



Some Operational Safeguards



- System redundancy & hybridization
- Manual confirmation & override
 - STD 8739 requires 2 operator actions to override
- Output constraints
- Learning restrictions
- Automatic safing/fallback mechanisms
- Transparency for explainability



Supportability



- Current computing infrastructure may not yet support full-fledged AI beyond LEO
 - Computational requirements are high (especially in learning phase)
 - Current computers cannot withstand radiation effects beyond LEO
 - More severe Van Allen radiation belts begin around 400 NM
- GSFC's FAST (Foundational AI for Space Technology)
 - Planning to leverage power of HPSC (High Performance Space Computer) now available
 - Could allow space-based computational products prior to downlink
- GRC working on ASIC Tensor Processing Units (TPUs) chips, whose sole purpose is accelerating neural networks.
 - These chips hardcode tensor instructions, parallelizing neural inference time. SCAMP (Super Complicated Ai Mission Payload) is a three-part high-atmospheric experiment designed to evaluate TPU performance in flight-like systems, the first step towards integration with small-sats and sounding rocket payload development.
 - Being tested in high-altitude-balloon missions



Internal AI Guidance



- NASA/TM-20210012886 NASA Framework for the Ethical Use of Artificial Intelligence (AI)*
- NASA's Responsible AI Plan (2022)

*Being updated.



NIST AI RMF



From AI Risk Management Framework:

Includes 72 organizational best practices to accomplish sound AI risk management.



Fig. 5. Functions organize AI risk management activities at their highest level to govern, map, measure, and manage AI risks. Governance is designed to be a cross-cutting function to inform and be infused throughout the other three functions.



Principles: NIST AI RMF



Comparing Principles of NASA Framework & NIST AI Risk Framework:

NASA	NIST
(1) Fair	(1) Fair – With Harmful Bias Managed
(2) Explainable & Transparent	(2) Explainable & Interpretable
(3) Accountable	(3) Accountable & Transparent
(4) Secure & Safe	(4) Safe; (5) Secure & Resilient
(5) Human-Centric & Societally Beneficial	(6) Privacy-Enhanced
(6) Scientifically & Technically Robust	(7) Valid & Reliable



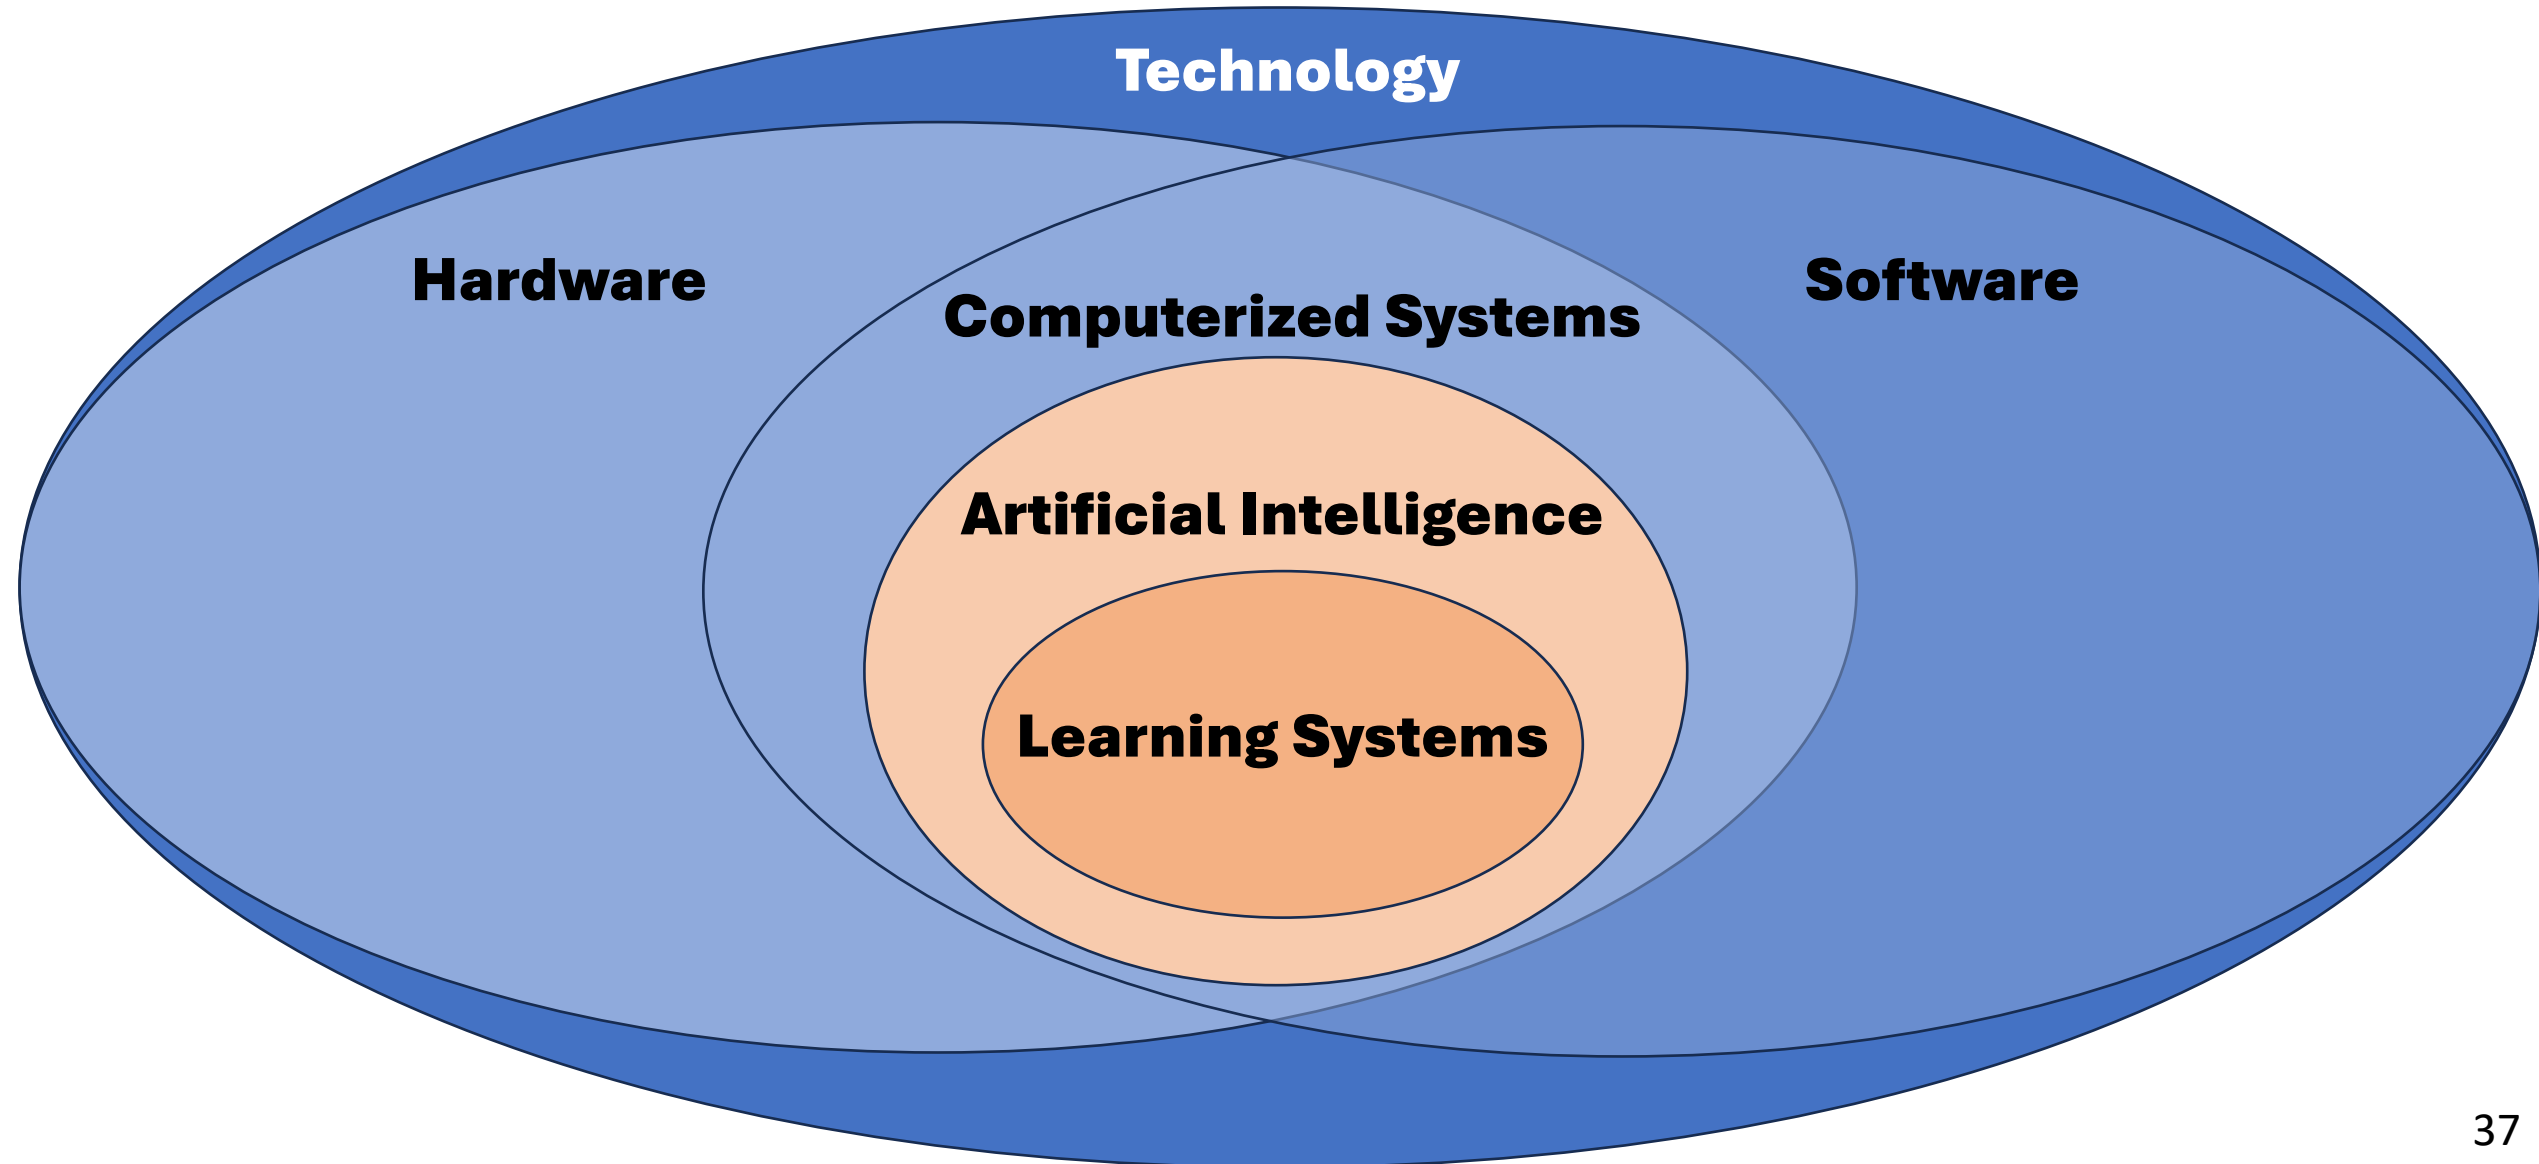
Newer Information



- Two new OMB memoranda on AI released on 4/3/25:
 - [OMB M-25-21](#), Accelerating Federal Use of AI through Innovation, Governance, and Public Trust
 - [OMB M-25-22](#), Driving Efficient Acquisition of Artificial Intelligence in Government



Technology & Learning





Distinguishing Security & Safety



A cybersecurity distinction from DHS/CISA:*

- (CYBER)SECURITY ASSURANCE: “Measure of confidence that the security features, practices, procedures, and architecture of an information system accurately mediate and enforce the security policy.” (i.e., standard cybersecurity)
- AI ASSURANCE: “A process that is applied at all stages of the AI engineering lifecycle ensuring that any intelligent system is producing outcomes that are valid, verified, data-driven, trustworthy and explainable to a layman, ethical in the context of its deployment, unbiased in its learning, and fair to its users.”
- To wit, FAA’s guiding document is titled *Roadmap for Artificial Intelligence Safety Assurance (July 2024)*

*Department of Homeland Security/Cybersecurity & Infrastructure Security Agency



AI/ML Differences



From AI Risk Management Framework, Appendix B: How AI Risks Differ From Traditional Software Risks:

1. The **data used for building** an AI system may not be a true or appropriate representation of the context or intended use of the AI system, and the ground truth may either not exist or not be available. Additionally, harmful bias and other data quality issues can affect AI system trustworthiness, which could lead to negative impacts.
2. AI system dependency and reliance on **data for training** tasks, combined with increased volume and complexity typically associated with such data.
3. Intentional/unintentional **changes during training** may fundamentally alter AI system performance.
4. **Datasets used to train** AI systems may become detached from their original and intended context or may become stale or outdated relative to deployment context.
5. AI system **scale and complexity** (many systems contain billions or even trillions of decision points) housed within more traditional software applications.
6. Use of pre-trained models that can advance research and improve performance can also increase levels of **statistical uncertainty** and cause issues with bias management, scientific validity, and reproducibility.



AI/ML Differences



Continued...

7. Higher degree of difficulty in predicting **failure modes** for emergent properties of large-scale pre-trained models.
8. **Privacy** risk due to enhanced data aggregation capability for AI systems.
9. AI systems may require more frequent **maintenance** and triggers for conducting corrective maintenance due to data, model, or concept drift.
10. Increased **opacity** and concerns about reproducibility.
11. Underdeveloped software testing **standards** and inability to document AI-based practices to the standard expected of traditionally engineered software for all but the simplest of cases.
12. Difficulty in performing regular AI-based software **testing**, or determining what to test, since AI systems are not subject to the same controls as traditional code development.
13. Computational **costs** for developing AI systems and their impact on the environment and planet.
14. **Inability to predict** or detect the side effects of AI-based systems beyond statistical measures.



TRLs for AI/ML: A Synopsis



TRL	Scale	Data	Integration	Conditions
0 1st principles	Novel idea	None	None	Research only
1 Goal-based research	Low-level algorithm	Small sample sets	None	Local computer
2 Proof of Principle Dev.	Active R&D	Simulated/surrogate, large-scale data sets	None	Simulated testbed
3 Systems Dev.	Prototype model	Dedicated test data sets	None	Tests of compatibility, scalability, reliability
4 Proof of Concept	Full application w/documentation	Real-world, use-case	Candidate practical applications	Stochastic, noisy



TRLs for AI/ML: A Synopsis



TRL	Scale	Data	Integration	Conditions
5 ML Capability	Production-type	Large-scale real-world	Specific capability	Real-world; V&V complete
6 Application Dev.	Robust, product-caliber code	Robust, multiply-sourced, incl. adversarial	Target use cases; autonomy rules defined	Users w/integration APIs, etc.
7 Integrations	Full scale, fully integrated	Baseline data set, data governance, & test data	Into productions systems	Production uses
8 Mission-ready	Final form	Real data measured vs. performance	Fully integrated	Operational environment
9 Deployment	Operational system	Continuous monitoring of data quality & model drift	Fully deployed w/continuous policy changes	Full operational use



7150: Risk-Based Analysis




- **Use Model-Based Systems Engineering (MBSE) to conduct gap assessment and manage traceability of standards, gaps, risk, and recommended fix actions.**
- **Identify & define key differences between AI-based software & traditional software**
 - Leveraging authoritative sources to define unique AI/ML requirements.
 - Consulting NASA standards for software engineering
- **Perform differential analysis to reveal potential gaps in NASA standard coverage of AI/ML software requirements.**
- **Associate potential risks to human-rated space systems if AI/ML software is developed using current NASA standards by assessing driving factors of:**
 - likelihood that AI/ML software will perform an action
 - driving factors of consequences that action will have
- **Produce recommended closure criteria to address the risk of identified gap.**



Federal Guidance



- Dept. of Commerce/NIST: “[Artificial Intelligence Risk Management Framework](#)”, NIST AI 100-1 [1/23]
 - “[Roadmap for the NIST AI Risk Management Framework](#)” [3/14/23]
 - “[AI RMF Playbook](#)” [3/30/23]
 - “[Artificial Intelligence Risk Management Framework: Generative AI Profile](#)”, NIST AI 600-1 [7/24]
 - “[Secure Software Development Practices for Generative AI & Dual-Use Foundation Models](#)” [7/24] (Note: Covers cybersecurity risks only)
- FAA “[Roadmap for Artificial Intelligence Safety Assurance](#)” [7/23/2024]





Risk Reduction: FAA Roadmap



- **Differentiate between **Learned and Learning AI****: Establish distinct safety assurance methodologies for learned (static) AI and learning (dynamic) AI, understanding the difference their respective operational and safety implications.
- **Take an Incremental Approach**: **Implement AI in aviation incrementally**, learning and adapting safety assurance methods based on real-world application and experience.
- **Leverage the Safety Continuum**: **Utilize the safety continuum, starting with lower-risk applications to gain experience and inform broader applications and safety methods.**
- **Leverage Industry Consensus Standards**: **Adopt industry consensus standards** for AI safety assurance in aviation, as appropriate, promoting global harmonization and adapting to technological changes while aligning with the principles in the roadmap.



Federal Guidance



- U.S. Law:
 - [“National Artificial Intelligence Initiative Act of 2020 \(P.L. 116-283\)”](#) [1/1/21]
 - [“Federal Artificial Intelligence Risk Management Act of 2023”](#) [11/2/23]
- Presidential Executive Orders:
 - 13859: [“Maintaining American Leadership in Artificial Intelligence”](#) [2/14/19]
 - 13960: [“Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government”](#) [12/3/20]
 - 14110: [“Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence”](#) [10/30/23]
- OMB Memos:
 - M-24-10: [“Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence”](#) [3/28/24]
 - OMB Memo M-24-18: [“Advancing the Responsible Acquisition of Artificial Intelligence in Government”](#) [9/24/24]



Federal Guidance



- Dept. of Homeland Security Dept/CISA: “[CISA Roadmap for Artificial Intelligence](#)” [11/23]
- Dept. of Transportation: “[An Overview of AI Assurance for Transportation](#)” [4/2024]
- State Dept. “[Risk Management Profile for Artificial Intelligence & Human Rights](#)” [7/25/24]
- Dept. of Defense: “[Data, Analytics, and Artificial Intelligence Adoption Strategy](#)” [6/27/23]
- Government Accountability Office: “[Artificial Intelligence: An Accountability Framework for Federal Agencies & Other Entities](#)” [6/21]



NASA Engineering Guidance



- NPR 8705.2C Human-Rating Requirements for Space Systems [7/10/2017]
 - NASA-STD-8719.29 NASA Technical Requirements for Human-Rating [12/11/2023]
- CCT-REQ-1130 Crew Transportation & Services Requirements Document
- NPR 7120.5F NASA Space Flight Program and Project Management Requirements
- NPR 7120.7A NASA Information Technology Program & Project Management Reqmts
- NPR 7123.1D NASA Systems Engineering Processes and Requirements [7/5/2023]
 - NASA/SP-6105 NASA Systems Engineering Handbook, Rev. 2 [2016]
 - NASA/SP-20210010952 Human Systems Integration Handbook [2021]
- NPR 7150.2D NASA Software Engineering Requirements [3/8/2022]
 - NASA-HDBK-2203 NASA Software Engineering Handbook [4/20/2020]
- NASA-GB-8719.13 NASA Software Safety Guidebook [3/31/2004]
- NASA-STD-8739.8B Software Assurance and Software Safety Standard [9/8/2022]
- NASA-STD-7009 Standard for Models and Simulations [12/7/2016]



External Standards



Possible alternatives to NIST's *AI Risk Management Framework* include*:

- ISO/IEC 42001: IT—AI —Management system [12/2023]
 - “provides requirements for establishing, implementing, maintaining, & continually improving an AI management system”
- ISO/IEC 23894: IT—AI —Guidance on risk management [2/2023]
 - “provides guidance on how organizations that develop, deploy, or use products, systems and services that utilize AI can manage **risk specifically related to AI**”

*not yet fully reviewed



External Standards



- ISO 26262: Dynamic safety assurance
 - Object detection, collision avoidance, path planning
- ISO 5469
- ISO 21448
- ISO 5338
- 24028: Overview of Trustworthiness in Artificial Intelligence
- ISO 25010: Systems & software quality monitoring
- ISO 62443: Fail-safes & Redundancies
- ISO 8800: Safety & AI
- ISO 25053: Framework for AIML Systems



AI Requirements Gap Analysis



NASA Standards Applicable to AI

Fully Applicable
NASA-NPR-7120.7A
 NASA Information Technology Program & Project Management Requirements

Fully Applicable
NASA-NPR-8705.2C
 Human-Rating Requirements for Space Systems

Fully Applicable
NASA-NPR-7150.2D
 NASA Software Engineering Requirements

Fully Applicable
NASA-STD-8739.8B
 Software Assurance & Software Safety Standard

Existing CLDP Content

Requirements for IT systems
 IT Lifecycle and Project Management

Human-Rating Requirements
 V&V Requirements for Critical Software

General software requirements

Safety-critical software determination
 Human-rated software classification
 Provides lifecycle requirements applicable to AI systems

AI SW Content Gaps

- AI-tailored risk management
- Training Data
- Long-Duration Autonomy
- Security of AI Models and Parameters
- AI-specific lifecycle
- Real-time Adaptability and Learning
- V&V and Testing of nondeterministic software
- AI integration with existing space systems
- AI system management

AI Industry Standards

- NIST AI 100-1**
AI Risk Management Framework
- NIST SP 800-218A**
Secure Software Development Practices for Generative AI and Dual-Use Foundation Models
- ISO/IEC 5338**
Information technology — Artificial intelligence — AI system life cycle processes
- ISO/IEC 23053**
Framework for Artificial Intelligence (AI) Systems Using Machine Learning (ML)
- ISO/IEC 42001**
Information technology — Artificial intelligence — Management system

Note 1:
 Many of the industry documents address multiple gaps listed. This trace is not comprehensive.

Note 2:
 AI **Hardware** expected to be addressed by existing NASA processes.

Although AI standards have been initially identified, none have been reviewed, approved, or vetted for human spaceflight applications by NASA software authorities.

[Link: Crosswalk between AI standards](#)




Why “Learning AI” is Harder



From NIST 100-1, AI Risk Management Framework:

“While there are myriad standards and best practices to help organizations mitigate the risks of traditional software or information-based systems, the risks posed by AI systems are in many ways unique (See Appendix B). AI systems, for example, may be trained on data that can change over time, sometimes significantly and unexpectedly, affecting system functionality and trustworthiness in ways that are hard to understand. AI systems and the contexts in which they are deployed are frequently complex, making it difficult to detect and respond to failures when they occur. AI systems are inherently socio-technical in nature, meaning they are influenced by societal dynamics and human behavior. AI risks – and benefits – can emerge from the interplay of technical aspects combined with societal factors related to how a system is used, its interactions with other AI systems, who operates it, and the social context in which it is deployed.”





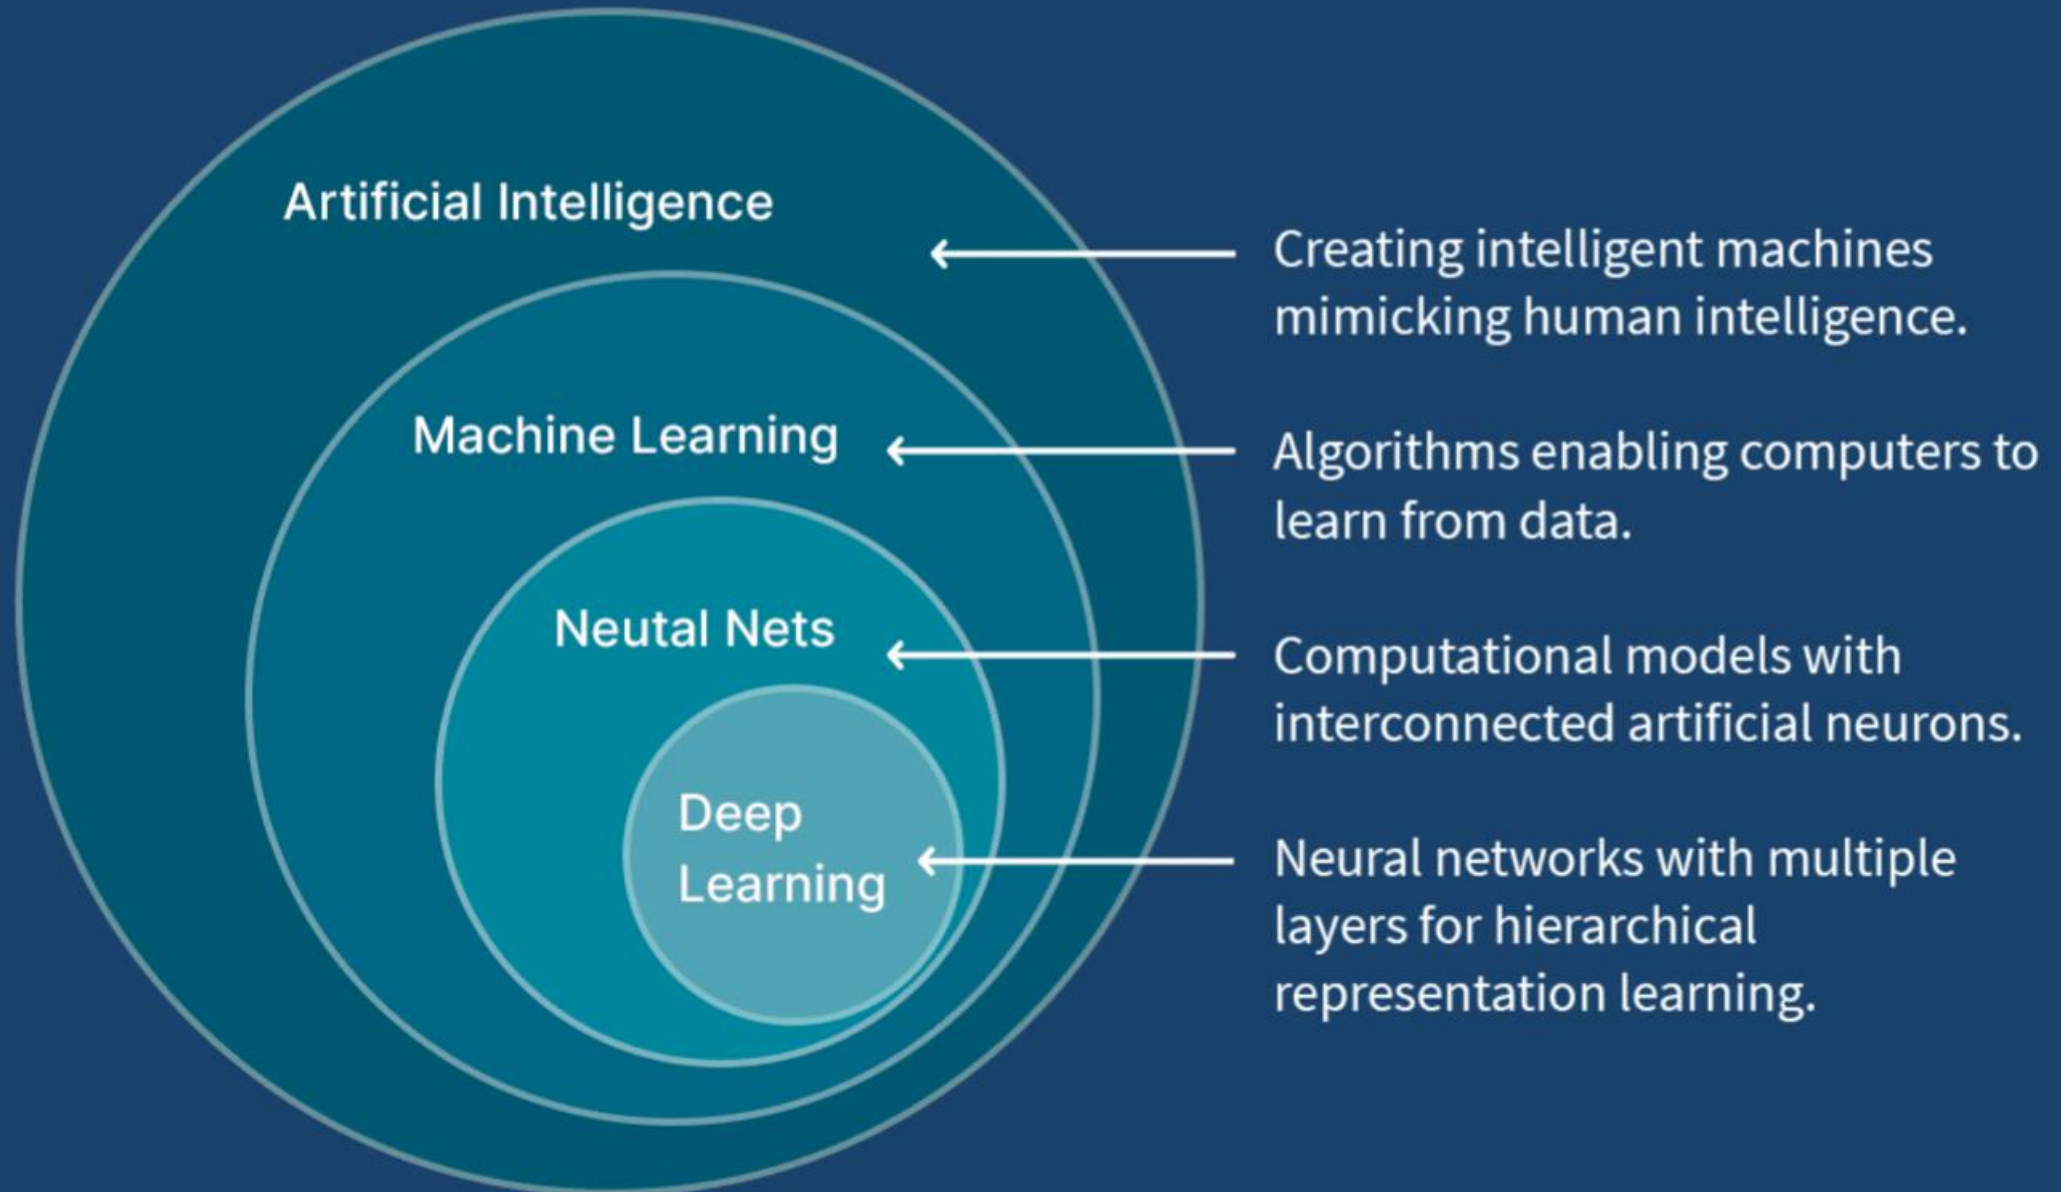
AI is not coming—it's here!



- Credit product approvals
- Petroleum production predictions
- Aircraft collision avoidance (ACAS (airborne) vs. TCAS (traffic))
- Aerospace engineering (fluid dynamics, combustion, acoustics, structural)
- Insurance risk & fraud analysis
- Sales vs. demographics patterns
- Star cluster identification
- Investment market analysis
- Sports performance predictions
- Disease predictions
- Manufacturing equipment failure predictions
- Emergency room situational awareness, X-ray/ultrasound image analysis
- Software development

NASA itself had
over 600 uses
as of Dec. 2024

Relationship of Technologies





Software vs. AI/ML



Traditional Coded Systems



“Learning AI” Systems

