AIAA

Autonomous Aerospace Systems Research & Development Priorities

An AIAA White Paper

June 2025

Authored by the AIAA Autonomy Research and Development Task Force, approved by the AIAA Board of Trustees

TABLE OF CONTENTS

STATEMENT OF ATTRIBUTION

This paper was drafted by the AIAA Autonomy Research and Development Task Force, under the auspices of the AIAA Aerospace R&D Domain, over fall 2024 and winter 2025, and approved by the AIAA Board of Trustees in summer 2025. The Task Force consisted of members from academia, industry, and government who, collectively, have a breadth of experience in Autonomous Aerospace Systems Research & Development.

Contributing Members of the AIAA Autonomy Research and Development Task Force

Lead: Kerianne Hobbs, Vehicle Autonomy and System Trust
The Aerospace Corporation, AIAA Associate Fellow.

Co-Lead: Wes Ryan, Fellow for Airworthiness of Autonomy and AI
Northrop Grumman, AIAA Associate Fellow.

Greg Zacharias, former Chief Scientist of the United States Air Force
National Research Council Member, and AIAA Aerospace R&D Domain Lead.

Kristi Morganson, Chair, Department of Aeronautics and Astronautics
University of Washington, AIAA Fellow.

Natasha Neogi, Senior Technologist, Assured Intelligent Flight Systems
NASA Langley Research Center, AIAA Fellow.

John-Paul Clarke, Ernest Cockrell, Jr. Memorial Chair in Engineering
Aerospace and Engineering Mechanics, University of Texas, AIAA Fellow.

Danette Allen, Executive Director, Old Dominion University Modeling
Analysis, & Simulation Center, AIAA Associate Fellow.

Daniel Selva, Professor, Department of Aerospace Engineering
Texas A&M University, AIAA Associate Fellow.

The authors would also like to acknowledge and thank the following for their reviews and feedback:

- Warren Dixon, Dean’s Leadership Professor and Department Chair, University of Florida
- Jon Clauss, Technical Fellow, Lockheed Martin
- Kim Wasson, Autonomy Assurance and Certification, Joby Aviation
- Pamela Wood, Senior Leader in Advanced Autonomy for Space Systems, Aerospace Corporation
- Enrique Medina, Founder and Managing Director, Enrique Medina & Associates
- Constance Dunn, Computer Science Professional Development and (K-12) Curriculum Development Specialist, BootUp Professional Development
- Andrea Helbach, Middle School Science Teacher, Dayton Regional STEM School
- Felix Knutson, International Program Officer, Air Force Office of Scientific Research
- Joshua Moser, Research Aerospace Engineer, NASA Langley Research Center
- Hal Greenwald, Program Officer, Air Force Office of Scientific Research
- Fred Leve, Program Officer, Air Force Office of Scientific Research
- Erik Blasch, Program Officer, Air Force Office of Scientific Research
- Trey Smith, AI and Autonomy Solver-in-Residence at NASA Headquarters

EXECUTIVE SUMMARY

The aerospace industry is on the cusp of a revolution, with aerospace systems evolving from manual operation coupled with simple automated tasks to sophisticated, intelligent platforms capable of sensing, acting, and collaborating in complex environments. This position paper focuses on autonomy applications onboard air and space vehicles, also referred to as autonomy "at the edge" (e.g., deploying locally to enable real-time decision making). The AIAA Autonomy Research and Development Task Force has identified two critical challenges for research and development that demand immediate attention: 1) ensuring the safety of these increasingly complex autonomous aerospace systems and 2) equipping the current and future workforce with the specialized skills needed to drive and continuously sustain this rapid innovation. In this position paper we delve into the current state of autonomous aerospace systems to set the context, identify key research and development gaps, and propose actionable recommendations to propel the field forward. We begin by outlining the fundamental functional requirements of autonomous systems, maintaining a safe state condition related to the system and environment, encompassing self- and situational awareness, intelligent reasoning and action, and collaboration between humans and autonomous entities. We then explore the unique considerations that define aerospace autonomy, including the intricacies of safety-critical and mission-critical systems, the differing demands of civil and military aviation, evolving risk tolerances, economic realities, and the distinct challenges of air versus space applications. Furthermore, we analyze the entire lifecycle of autonomous systems, from initial design to rigorous testing and operational deployment, to uncover critical gaps and formulate targeted recommendations. Finally, we address the vital need for workforce development and training across all levels, from inspiring K-12 students to shaping undergraduate and graduate programs and providing continuing education for seasoned professionals. The paper concludes with a list of recommendations, recognizing that this work is part of a larger, ongoing dialogue reflected in the supporting documents cited in the appendices.

PURPOSE AND SCOPE

The purpose of this position paper is to bring aerospace academic, government, and industry stakeholders to a consensus on aerospace research and development priorities to enable the responsible development and operation of trustworthy and assured ***autonomy onboard air and space vehicles, also referred to as autonomy "at the edge."*** Additionally, considerations are provided primarily for a single vehicle, which might be used within a larger mission of many vehicles; however specific considerations for multi-agent autonomy (e.g., swarms) are out of scope. Autonomy technologies, including (but not limited to) machine learning (ML) and artificial intelligence (AI), have significant potential to enhance safety, efficiency, and mission capabilities over the next decade. An AI system or system that employs ML might enable a particular implementation of autonomy; however, ***AI and ML are not the only path to autonomy*** and benefits should be carefully weighed against verification and lifecycle management challenges. Additionally, it is important to note that not all AI systems are implementations of autonomy. Some AI systems may serve as perception and decision aids, providing guidance, visualization, and other representations of knowledge. However, challenges and concerns regarding safety, design, development, testing, airworthiness certification, and ethical standards must be addressed. This report identifies *safety assurance* and *workforce development* as the two most critical and challenging gaps for the research and development of autonomous aerospace systems.

The possible scope of autonomy applications in aerospace is large and could encompass many areas where new technologies like machine learning could contribute. This paper specifically focuses on autonomy technologies onboard air and space vehicles. Emerging approaches such as the use of generative AI aides in software development, design, system and test documentation and test case generation is out of the scope of this paper. Additionally, AI to handle massive scale aerospace challenges such as AI aides for air traffic control and space traffic management is out of the scope of this paper. The field of AI, ML, and autonomy is evolving rapidly and no brief position paper can possibly address all applications, gaps, and recommendations. The authors and AIAA agree that there are additional viewpoints and perspectives beyond the AIAA Autonomy Research and Development Task Force authorship, including but not limited to several references at the end of this report, that facilitate a more comprehensive and holistic understanding of aerospace research needs.

FUNCTIONAL REQUIREMENTS

It is likely impossible to have a discussion on autonomy without an argument on the definition of autonomy. Does an autonomous system need to be “intelligent” or use some form of machine learned components? Does autonomy simply define advanced control beyond what can be accomplished now? For the purposes of this publication, an ***autonomous aerospace system*** is defined as *an air or space system that achieves goals within delegated and bounded authority while operating independently or with limited external control or direct human intervention*. This section describes broad functional requirements to achieve this objective, specifically in the categories of being informed of its own state and the state of its environment, reasoning and acting, and collaboration and interaction.

2.1 Sensing and Perception

The first requirement of any autonomy function is for it to be informed of its own state and the state of its environment to manage objectives and safe outcomes. This is sometimes referred to as sensing and perception or self- and situational awareness, which encompasses a larger philosophical grouping as described in Table 1. This process involves representing information in a machine-readable format. In aerospace, we often think of physics-based representations (e.g., position, velocity); however, the field of AI has defined several frameworks for describing assertions about the world (e.g., formal logic or domain-specific languages). A gap in this research area is consistent identification of the knowledge representation framework used and associated assumptions or abstractions. The problem is further complicated as autonomous systems become more collaborative and boundaries extend beyond traditional definitions of an individual agent to the system of systems case. Several different knowledge representations may be used in different hierarchical levels of an autonomous system and for systems of systems. For example, mode switching may be defined in a logic language, while lower-level inner loop control may be defined using physics-based states and relationships described using differential equations. These various knowledge representation frameworks enable autonomous agents to reason about their environment and make informed decisions.1

Effective sensing and perception hinges on an autonomous agent’s ability to assess the accuracy of its information. This understanding of the current state and its potential evolution enables the agent to make informed decisions to achieve desired future states, even in rapidly changing environments. Sensing and perception describes activities such as the interrogation, identification, and evaluation of the state, a collection of variables at a point in time that describes the environment and/or the system. This larger group of tasks can be further divided into applications described in Table 1.

Table 1: Sensing and Perception Application Categories[2]

Sensing and Perception	State Estimation and Monitoring	Knowledge and Model Building
Collection and processing of information internal and external to the system from sensors and instruments	Estimation, ascertainment, and continuous comparison of internal and external states from raw or processed inputs generated by multiple sensors or instruments.	Creation of information sources about the environment or system from sensing, perception, and human interaction (i.e., turning data into context information)
Hazard Assessment	**Event and Trend Identification**	**Anomaly Detection**
Evaluation of whether the state of the environment, state of the system, and/or their interaction poses a threat to the safety of action or inaction that could compromise the system or mission	Analyses of data to identify events or trends that may affect future states, operations, or decision making	Determination that the environment or system is not exhibiting expected characteristics

[1] Russell, Stuart J., and Peter Norvig. Artificial Intelligence: A Modern Approach. Pearson, 2016.

[2] Adapted from: Fong, Terrence W., Jeremy D. Frank, Julia M. Badger, Issa A. Nesnas, and Michael S. Feary. "Autonomous Systems Taxonomy." NASA Autonomous Systems Capability Leadership Team, Report no. ARC-E-DAA-TN56290. 2018.

2.2 Reasoning and Acting

Once sensing and perception is achieved in a particular knowledge representation, an autonomous agent needs to reason about that information to make a decision and then act on that decision. Reasoning and acting is the process of making decisions and taking actions to achieve a goal based on the current understanding of the state of the environment, agent, system, and/or system of systems. An exceptionally important question in safety and mission critical autonomous aerospace systems is how we know an autonomous agent will take actions in an acceptable manner, i.e., in accordance with its own individual objectives and the collaborative autonomy objectives of the overall system.

Decision Aids and Alerting Systems

Decision aids may provide reasoning information and recommendations to a human or agent in an informative or assistive manner but stop short of taking action. It is important to note the difference between a decision aid and the reasoning and acting process discussed here. A decision aid provides information and recommendations to a human or automation, which then acts. Decision aids can help identify options available, as well as the likely outcomes, risks, benefits, and other tradeoffs between each option. While decision aids can improve the efficiency and effectiveness of the decision-making process, the decision and action are made by a human or automated system rather than the aid itself, recognizing there may be limitations to the efficacy of the decision aid. Alerting systems fall into this category. For example, traffic alert and collision avoidance system (TCAS) alerts pilots of traffic and may issue a resolution advisory that recommends actions such as climbing or descending to avoid a collision.

Decision Making, Control, and Layered Architectures

In autonomous aerospace systems, reasoning and acting usually takes place within a layered architecture,[3,4] where each layer or level has its own assumptions and abstractions. Note that the following is not by any means the only definition of levels, and these high levels may be refined. For example, layers could include objective level, mission level, system level, vehicle level, and component level. Application-specific layers might include specific commands and information sent between layers from an operator all the way down to sensor and actuator level inputs and outputs. Although there have been efforts to standardize the definition of architectural layers in autonomous systems, a universally accepted understanding of how to define these architectures and precisely pinpoint where autonomous decisions and actions occur remains elusive. This is because autonomy is highly application-specific and can occur concurrently across multiple layers and among collaborating agents. The connections and assumptions between the different layers, as well as the ability to trace requirements and dependencies across them, are all crucial aspects to consider.

To explain this layered architecture further, consider the following Army metaphor, visualized in Figure 1.

- At the top level, we have mission-level decision making, which includes selection of goals, objectives, and constraints, with associated requirements derived at that level. Depending on the design, this level could also include some activity and resource planning, including a selection and ordering of activities to be performed while managing resources to achieve goals at an individual and collaborative level. Think of this like a general officer who has high decision-making autonomy, is responsible for the decisions and actions of a large group under their command, and is expected to make decisions given objectives from higher ranking officials. They translate higher level goals and requirements into a plan of action that may include a big picture description of the movement of troops and equipment to achieve specific goals. In a complex autonomous decision system, this mission-level decision making entity may receive some bounded, and delegated authority to achieve goals of the human operator. The knowledge representation used for reasoning and acting may be some form of semantic logic or domain specific language that enables reasoning and optimization over complex options.

- The next level is a guidance, navigation, or planning layer, which we can think of as a Lieutenant who directs the platoon during combat with specific requirements to achieve the mission and instructions of higher commanders like the General. The Lieutenant has a moderate level of decision autonomy and is expected to follow a small set of orders and top-level requirements from higher ranking officials. This level translates higher level mission specifications and tasks into paths and trajectories for each agent to reach a desired end state (e.g., a physical location) subject to applicable constraints and reports back a status and feasibility for the task.

[3] Matni, Nikolai, Aaron D. Ames, and John C. Doyle. "Towards a theory of control architecture: A quantitative framework for layered multi-rate control." *arXiv preprint arXiv:2401.15185* (2024).

[4] J. Boskovic and R. Mehra, A Multilayer Architecture for Intelligent Control of Unmanned Aerial Vehicles, J. of Aerospace Computing, Information, and Communication, Dec 2004

- The lowest level is some form of feedback execution and control, which includes changing the system state to follow the path given, according to limitations, capabilities, and achievable requirements. This is like a private who has very little decision-making autonomy, is generally not responsible for decisions and actions of anyone else, and is expected to follow orders. In an autonomous aerospace system, most people would consider this level to be "automated," and likely using physics and mathematically defined tools like differential equations to describe the physics. This layer may use a simple control technique and may assure constraint satisfaction with run time assurance techniques such as control Lyapunov and barrier functions.

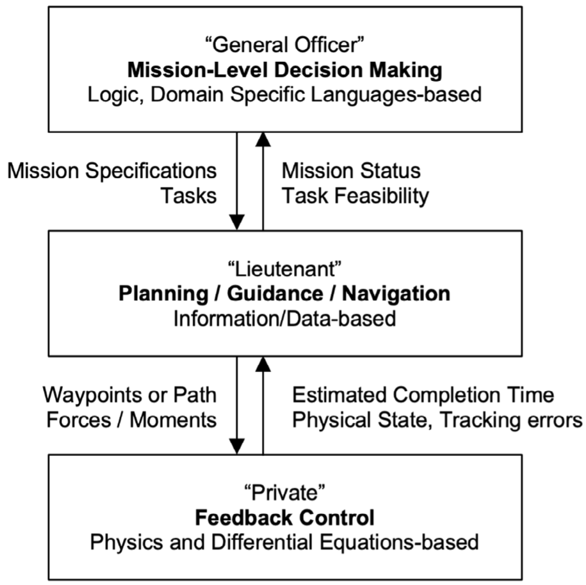

Figure 1: Levels of Autonomous Reasoning and Acting

Just as there are many levels between a General and Lieutenant and between a Lieutenant and Private, there can be many levels between mission-level decision making, trajectory planning, and feedback control with requirements and assumptions passed between each layer of abstraction. A similar hierarchy can be used to describe the abstraction necessary for evaluating a single agent versus a collaborative system of systems.

Architectural Mitigations and Controls for Autonomy

Safety interlocks, monitor and recovery concepts, fault detection and isolation, and other methods of architectural mitigations for safe implementation of automaton functions have long been part of system architecture to enforce safe outcomes by design. Safety interlocks prevent the simultaneous execution of two mutually exclusive actions to prevent an unsafe event, state, or condition. Common examples include a safety interlock that an elevator should not be in motion when the door is open, and the door should not open when an elevator is moving. In each case, the presence of one condition triggers a lockout of the other condition. Fault detection, isolation, and recovery (FDIR) is a form of autonomous reasoning and acting that identifies and isolates faults within the system and implements corrective actions to maintain safety and mission success. Traditional FDIR enables autonomous aerospace systems to respond to predicted failure states with predefined actions. Run-time assurance (RTA)[5] is a related concept that is designed to monitor an advanced controller to ensure its behavior is within predefined bounds and intervene when necessary to modify the control signal to ensure safety. An automatic collision avoidance system that detects an imminent collision, maneuvers to avoid, and returns control back to a human operator or autonomous control system is an example of run-time assurance.

In traditional systems design paradigms, performance requirements, safety constraints, and delineation between normal and failed states are explicitly defined. This means that conditions defining safety interlocks, FDIR, and RTA are straightforward and are explicitly defined. Today's aerospace systems largely rely on humans to handle complex failures and unforeseen events that were not envisioned or specifically planned for in the design. The response to accidents is to either increase training of aircraft pilots or spacecraft pilots and/or add requirements to the design and test of the system. Autonomous systems must be able to ensure acceptable outcomes beyond predicted failure states, as human intervention is not always possible in complex environments. The problem is exacerbated for collaborative systems where unintended behavior of one agent can negatively impact the ability of other agents to achieve safe outcomes. This is a significant gap in the development of autonomous systems that may be addressed with a variety of mechanisms, which are largely dependent upon the specific intended function of the system. Examples include fail-safe design, fail-functional design, multiple and/or alternate sources of function, safe recovery functions, contingency management, and alternate operating procedures.

[5] Hobbs, Kerianne L., Mark L. Mote, Matthew CL Abate, Samuel D. Coogan, and Eric M. Feron. "Runtime assurance for safety-critical systems: An introduction to safety filtering approaches for complex control systems." IEEE Control Systems Magazine 43, no. 2 (2023): 28-65.

2.3 Collaboration and Interaction

Collaboration and interaction describe when two or more systems or elements of a system work together to achieve a desired outcome. In the context of autonomous aerospace systems, collaboration and interaction may take place between autonomous systems and humans, between multiple autonomous aerospace systems, or any combination thereof. Some of the biggest challenges in teaming of multiple agents, whether human or autonomous, include synchronizing distributed decisions and actions to maximize team effectiveness and scaling local decisions that result in well-coordinated global behavior. Regardless of whether collaborators are human or autonomous, at least four key elements are important to the success of the collaboration, as defined in Table 2.

Table 2: Key elements of successful collaboration between agents[6]

Joint Knowledge and Understanding	Behavior and Intent Prediction	Goal and Task Negotiation	Operational Trust Building
Collection, assembly, sharing and interpretation of information and proper function among elements to solve problems and plan actions or responses	Forecasting the actions of other elements or system to support collaboration and interaction	Agreement on current and future activities, priorities and disposition among elements or systems	Assurance that the system is operating in a manner consistent with expectations of all elements.

In addition to the key elements above, the application of traditional system safety and safety assurance methodologies to collaborative autonomy of "systems of systems" also poses a challenge to safe design, development, airworthiness, and operation of future autonomous systems. Some of these challenges are already being discussed by industry standards groups, such as SAE S-18, "Aircraft and System Development and Safety Assessment Committee." While collaboration across large numbers of autonomous vehicles such as aircraft swarms or spacecraft constellations is outside the scope of this paper, some of the key elements are addressed by Recommendation 1.

Human-Autonomy Teaming

Contrary to fears that autonomy will replace humans, automation and autonomy are perhaps best utilized in augmenting humans by *extending* human ability, *relieving* some of the burden from humans, *supplementing* human strengths, or serving as a less capable backup. Human-autonomy teaming leverages strengths of each to come to the problem, including the human's ability to manage and respond to unexpected situations, and the autonomous system's strengths in speed and scale of decision making.

The autonomous agent's need to determine its own safe state in relation to its intended operating environment was discussed earlier. In the context of collaboration, the autonomous agent may also need some information to determine the safe state of its human or robotic teammates. The human collaborator must also have awareness of the state of the agent and the intended operating environment. The awareness of a human's teammates with respect to autonomous agent is divided into three levels: 1) what the system is doing, 2) why it acted in a certain way or made a decision, and 3) what the system will do next or do differently in other circumstances.[7,8]

---

[6] Adapted from: Fong, Terrence W., Jeremy D. Frank, Julia M. Badger, Issa A. Nesnas, and Michael S. Feary. "Autonomous Systems Taxonomy." NASA Autonomous Systems Capability Leadership Team, Report no. ARC-E-DAA-TN56290. 2018.

[7] Sanneman, Lindsay, and Julie A. Shah. "The situation awareness framework for explainable AI (SAFE-AI) and human factors considerations for XAI systems." International Journal of Human–Computer Interaction 38.18-20 (2022): 1772-1788.

[8] Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2008). Situation awareness, mental workload, and trust in automation: Viable, empirically supported cognitive engineering constructs. Journal of Cognitive Engineering and Decision Making, 2(2), 140–160. https:// doi.org/10.1518/155534308X284417

CONSIDERATIONS FOR DEVELOPMENT, ASSESSMENT, AND OPERATION OF AUTONOMOUS AEROSPACE SYSTEMS

This section explores important considerations throughout the autonomous aerospace system lifecycle and addresses what makes aerospace different from other domains implementing autonomy-enabled capability and AI-enabled capabilities (e.g., AI as an implementation choice) in potential safety critical use cases vs. consumer use cases.

3.1 Complexity

Aerospace systems can be highly complex and safety-critical, sharing characteristics with industries like nuclear power and medical devices. Their escalating complexity makes development and verification increasingly costly and unsustainable.[9,10,11,12,13] While autonomy promises advancements, it further complicates system design and safety assessment. Both aircraft cost and software complexity have grown exponentially. Despite efforts to shift verification earlier in the development process, including investments in formal methods and digital engineering, substantial reductions in verification costs have not materialized. This suggests that traditional linear verification methods are insufficient, and a new paradigm is needed. Autonomy introduces unique challenges: defining decision boundaries for autonomous systems, developing robust safety architectures (such as run-time assurance), establishing metrics for resilience against uncertainty and failures (especially for AI-based autonomy), and ensuring safe human-system interactions.[14] Moreover, autonomy development often occurs independently of aircraft development, creating a need for retroactive certification processes, exemplified by the Automatic Ground Collision Avoidance System (Auto GCAS) integration into existing aircraft or advanced autopilot functions such as the Garmin Autoland system.

3.2 Verification, Validation, Test, Evaluation, Certification, Assurance, and Airworthiness

Words in the title of this section are often used incorrectly or interchangeably to describe the process of gaining shared confidence in safe function and accepting an aerospace system design; however, each captures slightly different activities. *Verification* addresses whether *the system is built right* and ensures that a product or system is implemented correctly at the design/development level to meet applicable requirements and specifications. *Validation* addresses whether *the right system is being built for the intended use* and fulfills the functional needs of the end user. The detailed use and application of these terms depends greatly on the frame of reference. For example, they are applied slightly differently for development assurance of software, systems safety engineering, model-based engineering, digital transformation of processes, and/or airworthiness. *Test and evaluation* are often used together to describe the process of exercising a system or components and analyzing performance against expected performance specifications and intended/proper function. *Certification* is a determination by the proper authority that a system is compliant with the applicable legal and design related standards or criteria. *Assurance* is certainty or confidence gained through methodical gathering and exchange of information that a system will perform its intended function or operation safely and repeatedly. *Airworthiness* describes the process of gaining the shared confidence an air system can safely achieve, sustain, and terminate flight according to applicable criteria and/or standards. The level of rigor expected for satisfying each activity listed in the title of this section is dependent upon the criticality of function of the system, its intended use case, and its intended area of operation. For

[9] Hobbs, K. "Evaluation of verification approaches applied to a nonlinear control system." Air Force Institute of Technology Master's Thesis. 2016.

[10] Feiler, P. H., "Supporting the ARP4761 Safety Assessment Process with AADL", Software Engineering Institute Carnegie Mellon University, 2014.

[11] DoD Autonomy Test and Evaluation, Verification and Validation (TEVV) Investment Strategy. Technical report, Assistant Secretary of Defense / Research and Engineering (ASD/R&E), 2015.

[12] Clark, M. Kearns K. Overholt J. Gross-K. Barthelemy B. and C. Reed. Air Force Research Laboratory Test and Evaluation, Verification and Validation of Autonomous Systems Challenge Exploration", 2014. Air Force Research Laboratory, Wright-Patterson AFB, OH.

[13] Brat, G., "Reducing V&V Cost of Flight Critical Systems: Myth or Reality." AIAA Information Systems-AIAA Infotech@ Aerospace. 2017.

[14] Jonathon Phillips, Carina A. Hahn, Peter C. Fontana, Amy N. Yates, Kristen Greene, David A. Broniatowski, and Mark A. Przybocki. "Four principles of explainable artificial intelligence." (2021): 09.

example, small uncrewed systems may achieve airworthiness certification and operational approval while meeting an acceptable level of safety by applying these terms differently than larger passenger-carrying aircraft. Note, while this section does not address challenges specific to the test, evaluation, verification, and validation of systems with machine learning components that may continue to adapt over the lifecycle of the host vehicle, that gap is addressed further in Section 5.2.1 of this document, and recommendations are made in Section 6.

There are currently sound engineering processes (mainly focused on the design phase of product development) to ensure safe operation in military aviation, civil aviation, and space systems (e.g., SAE ARP 4754B, MIL-STD-516, NASA NPR 8705.2C Human-Rating Requirements for Space Systems, etc.). These processes are largely based on assessment of system failures and their potential impact on the system. Their application poses a challenge as system boundaries extend to systems of systems, and for evaluation of interactions between a system and its environment, and for AI enabled capabilities where bidirectional traceability of requirements into the AI model itself may be impossible. While complying with the intent of these processes and requirements is fundamental to the goal of development assurance and system safety, it is important to remember that safety and airworthiness are not solely composed of a design review against requirements and performance criteria. System safety and airworthiness are also predicated on demonstrated performance and proper function, predicated on a shared confidence between the designer of the vehicle, the operator of the vehicle, and the certifying authority that the system is safe to operate and can perform its mission to the level of expected performance and reliability. This state of agreement is reached through an interactive process involving the systematic, iterative sharing of information. This includes coming to agreement upon criteria such as the expected level of safety for the system's intended use and involves a demonstration of the system's operational readiness and capability maturity level. Note that the expected safety level and capability maturity level will depend on the intended function of the system (e.g., prototype vs. commercial product).

Risk

Civil airworthiness regulations generally focus on mitigating design, operational, and airspace related hazards, and have evolved over time based on past accidents and incidents, such as hull loss and loss of separation between aircraft (and other obstacles). The airworthiness regulations enforce design and operational requirements that mitigate risk and protect first parties connected to the operation (i.e., passengers) and by proxy protect third parties not connected to the operation (i.e., bystanders on the ground). Similarly, designs and operations that meet applicable airworthiness requirements help protect against loss of separation between aircraft (and other entities/obstacles) to protect second parties connected to the operation (i.e., other users of the airspace). Given the excellent safety record of commercial aviation (i.e., Part 121 operations), the risk tolerance of these operations is extremely low, and the airworthiness requirements and means of compliance are stringent. It's important to acknowledge that risk tolerances vary significantly across aviation sectors, with general aviation, package delivery, and other non-emergency use cases operating under different safety expectations compared to the stringent Part 121 standards for commercial flights. While there are public service/public good missions that have higher risk tolerances, such as wildland firefighting, hurricane relief and recovery, emergency medical evacuation, etc., they are still carried out in the context of a civil airspace, and much of the risk is borne by the mission participants (e.g., emergency response workers, etc.) and not by other participants in the airspace or the public.

Military aviation missions are frequently geared toward carrying out a strategic objective that weighs a set of risks relative to the overall success of the action. Tradeoffs may be made with respect to risks related to third parties not connected to the operation, and, in fact, if those third parties are hostile, this may not even be deemed a risk. Furthermore, risks to second parties connected to the operation, such as other airspace users, may be deemed acceptable if warnings have been put in place to alert these entities that using the airspace requires them to bear the attendant risk. The same set of risks may be weighted differently depending on the mission context. Or it may be preferred that a vehicle be destroyed rather than allowing an adversary to capture key technologies. Finally, risks to first parties connected to the operation will need to be weighed against the larger picture of the endeavor, resulting in tradeoffs related to risks being disproportionately borne by smaller groups or entities that are deemed trained to be able to shoulder the burden. For example, some missions may accept a sacrificial loss of an autonomous vehicle as long as mission objectives are accomplished. Thus, the success of the mission may outweigh conventional risk tolerance concerns related to the traditional parties regarded as being protected (either in full or by proxy) by civil airworthiness regulations.

Regardless of the level of acceptable risk for a particular user base, a considerable gap in assessing risk of autonomous aerospace systems is formally defining that risk in a measurable way. Traditionally, guidance such as MIL-STD-882E and NASA S3001 define risk in a matrix with axes of probability/likelihood and severity/consequence. The severity/consequence definitions are still valid for autonomous aerospace systems, but depending on the design of the autonomy, probability/likelihood may not be an effective measure. Probability makes sense for physical systems in which a value such as "mean time between failures" can be measured. The complexity of software, by contrast, makes probabilistic failure analysis more challenging, and specific failure models used for hardware components may not be applicable for software. Even though a neural network control system may be stochastic during training, trying actions based on a distribution while it learns an optimal state-action pair, the final design is likely to be deterministic: given a particular state (system and environment), the autonomous system will produce the same response. Verifying the safety of autonomous systems is challenging because their complex responses, even when deterministic, are not readily analyzed using probabilistic methods, necessitating the development of alternative assessment measures

The risk definition challenge comes in the potential complexity of that mapping of state/observation to control input/action, in which a slightly different state could result in a drastically different control action. A substitute metric is probably needed to define risk for autonomous aerospace systems. It's possible a general metric does not exist, and specific metrics must be defined based on the implementation of the autonomy (e.g., neural network, algorithm, differential equation, etc.).

Mission-Critical vs. Safety Critical

Even though human lives or welfare may not be directly at risk, mission-critical autonomous systems share many of the same high assurance needs as safety-critical autonomous aerospace systems. So, what is a mission-critical autonomous aerospace system? For autonomous instrumented missions that go into deep space or robotic missions that autonomously explore extraterrestrial worlds, the risk of harm to humans is no longer the driving concern. However, due to the prolonged duration and high cost of these missions, it is vital that these platforms succeed in their endeavors. Similarly, military missions may prioritize achieving a mission over the return of a vehicle, like collecting and transmitting imagery prioritized over survival of a reconnaissance drone.

The issues that dominate high-confidence mission-critical systems are akin to those that influence safety-critical systems, but the emphasis falls on different hazards, failures, paths toward loss, and effective mitigations. In the space exploration or reconnaissance drone examples, the most valuable product from these missions is the data generated, and the highest criticality hazard may be loss of data. Thus, mitigating risk requires high levels of redundancy in data acquisition, processing, storage, and communication. Additionally in uncrewed and uncontested environments, an acceptable failure mitigation might be to use a safe mode from which recovery from upset can be made. A common safe mode strategy is to cycle power systems (spacecraft) or loiter (aircraft) and wait to re-establish contact. The strategy of waiting for contact and risking running out of fuel and crashing is clearly not viable in crewed operations or in contested domains.

Cost and Scale

Each civil aircraft type brought to market must eventually turn a profit for its manufacturer and/or its operator. A military aircraft must be capable of providing a desired capability in a military portfolio of systems for a desired operational theater. These inherently differing mission objectives lead to different cost-benefit tradeoffs being made in terms of aircraft performance, efficiency/sustainability, safety, and security. For instance, an autonomous selection of aircraft cruise altitude (possibly via use of an ML implemented component) may fractionally reduce emissions and fuel burn for a commercial aircraft but be costly to implement and certify. However, that cost can be amortized over an entire fleet of aircraft flying a multitude of routes (for multiple airlines) and the fuel (and cost) savings will accrue enormously over the lifespan of the fleet.

While military aircraft are produced at some scale, they are not operated at the same scale as commercial aircraft, and the cost-benefit ratio of such tradeoffs can skew significantly differently. For instance, it may be a given that multiple aircraft of a certain type be lost over the fleet lifespan due to hostile operating conditions. Military acquisition accounts for some attrition of their vehicles. This leads to qualities like robustness, durability, and reliability to dominate the concern of these vehicles in terms of airworthiness, especially given that maintenance may not be as readily available in all operational scenarios.

Civil and military aircraft applications both need more cost-effective certification for different reasons. While civil autonomy's large scale means nonrecurring certification costs can be spread across more users, these costs still reduce profits and may deter technology adoption. The military needs autonomy to stay competitive, but diverse autonomy needs across its aircraft may hinder competitiveness unless costs are reduced. This issue may not be as pressing today in the less-regulated space domain; however, it will likely become more important as space becomes increasingly congested and contested.

Reuse

Reuse is an expected feature of commercial and military aircraft, in that most aircraft types are based strongly upon a predecessor aircraft, with only one or two changes being made between models in commercial aviation. Similarly, each vehicle will have a 20–30-year lifespan and can undergo multiple operations a day. While many space platforms are based on prior designs, they are rarely manufactured at scale, and the knowledge surrounding their hazard landscape is founded on a limited sample size. More tellingly, most platforms designed for exploration missions are unique by their nature, as they are designed for the evaluation of an unknown environment in a first-of-its-kind mission. The ability to reuse successful autonomy solutions across multiple platforms, let alone across multiple mission environments demands a careful evaluation of the required critical properties to achieve mission success. Notably, this paradigm may be changing for commercial low Earth orbit missions such as launch capabilities and satellite constellation deployment. It remains to be seen whether the scale of these operations will yield opportunities for reuse of autonomous control software. Finally, novel capabilities such as autonomy will be best enabled by taking an engineering centric approach to design (e.g., vehicle, roles and responsibilities, etc.) and mission-oriented (e.g., concepts of operation, environmental uncertainty, etc.) approach to operation. The right balance of process-based transactional airworthiness approaches with novel system assurance-based approaches must be struck to mature autonomy technologies in a safe manner.

WORKFORCE DEVELOPMENT AND TRAINING

The increasing automation and autonomy of "flying robot" aerospace systems necessitate a workforce with a foundational understanding of systems, robotics and electrical engineering, as well as computer science and data analytics. The multidisciplinary nature of 21st-century aerospace engineering, as well as the shift toward digital engineering practices, resulting in the need for more systems engineers, highlights the need for expertise beyond traditional "controls" disciplines. The expanding use of autonomy across all aspects of engineered systems requires education and training at all levels, from K-12 to the existing workforce and C-suite,[15] to ensure a workforce equipped with the necessary skills and expertise to drive and sustain innovation.

4.1 K-12 STEM Educational Considerations

While K-12 STEM education provides a solid foundation in science, mathematics, and some technology, it often lacks sufficient focus on engineering and computer science. We describe a number of these shortcomings in the next section. The resultant lack of exposure to engineering and computer science fundamentals in K-12 leaves many students unprepared for the rigors of an engineering education and unaware of the opportunities this field presents.

Although some middle school projects introduce students to the engineering design process using materials like cardboard, Legos, and basic drafts, and additional kits with microprocessors, servos, and sensors are available, these experiences are not consistently integrated into the curriculum. The Computer Science Teacher Association has developed grade-appropriate standards,[16] but their implementation varies widely. There is no such aerospace K-12 set of standards. In many elementary schools, computer science exposure is limited to basic computer usage or typing skills, with "tech lab" time often rotated with other activities like physical education, music, and art. Middle and high schools may offer elective courses in computer science, robotics, or related autonomy topics, but these are not universally available and can be dependent on school or district resources. Even where offered, these topics might only be covered briefly within a core class, in a separate STEM lab, or as part of an after-school club, leading to inconsistent learning opportunities. A lack of exposure to engineering and computer science fundamentals in K-12 leaves many students unprepared for the rigors of an engineering education and unaware of the opportunities this field presents.

[15] M. Cummings, K. Morgansen, B. Argrow, and S Singh, "Transforming Aerospace Autonomy Education and Research," *IEEE Aerospace Conference*, March 2021.

[16] https://csteachers.org/k12standards/interactive/

4.2 Undergraduate Aerospace, Aeronautics, and Astronautics Requirements

Undergraduate curricula in aeronautical, astronautical, and aerospace engineering in the United States are almost exclusively accredited by ABET.[17] AIAA, specifically its Committee on Higher Education (CHE), is the lead society for management of ABET requirements in aerospace and aeronautical engineering technology. The current disciplinary requirements with ABET accreditation for these degrees has the following requirements:

> *Aeronautical engineering or similarly named engineering programs must include the following curricular topics in sufficient depth for engineering practice: aerodynamics, aerospace materials, structures, propulsion, flight mechanics, and stability and control.*
>
> *Astronautical engineering or similarly named engineering programs must include the following curricular topics in sufficient depth for engineering practice: orbital mechanics, space environment, attitude determination and control, telecommunications, space structures, and rocket propulsion.*
>
> *Aerospace engineering programs or similarly named engineering programs, which combine aeronautical engineering and astronautical engineering topics, must include all curricular topics in sufficient depth for engineering practice in one of the areas—aeronautical engineering or astronautical engineering as described above—and, in addition, similar depth in at least two topics from the other area.*
>
> *The major design experience must include topics appropriate to the program name.*

Most engineering undergraduate programs have a limit on the maximum number of credit hours that can be built into the program (180 credit hours for quarter-based programs, 120-130 for semester-based programs). This limit in combination with the extensive list of necessary material to be covered has led to a situation where new material cannot easily be incorporated into aerospace curricula in accredited programs. To address this issue, the AIAA CHE has worked to update the criteria for undergraduate degree programs. The following new language will shortly be sent out for public comment before being adopted as the new set of guidelines:

> *The curriculum must include modeling, simulation, computing, and testing applied to the design and analysis of aerospace systems or subsystems and their operations. Aeronautical engineering or similarly named engineering programs must cover atmospheric flight. Astronautical engineering or similarly named engineering programs must cover space flight and the means to get to space. Aerospace engineering programs or similarly named engineering programs must include content from both aeronautical engineering and astronautical engineering. The major design experience must include topics content appropriate to the program name.*

The new criteria were specifically designed to admit a greater focus on emerging areas of need for aerospace engineering, including autonomy.

While ABET has accreditation for disciplines such as aerospace engineering, there are not currently accreditation options for areas such as autonomy or data science, so no existing guidance can be leveraged to adapt for an appropriate curriculum of autonomy in aerospace engineering.

4.3 Graduate Program Considerations

Given the historical restrictions in developing new curricular directions at the undergraduate level, many academic programs have focused on building options at the graduate level. Such options include “minors” or “degree options” where students complete a particular set of courses (see e.g., the Data Science option at the University of Washington[18]). Additionally, graduate “micro-degree” certificates, consisting of 3-5 courses are emerging in graduate programs. In the past several years, a range of new master’s-level degrees have been developed in robotics. These programs contain elements of autonomy but often focus more on a computer science perspective than the safety critical verification and validation needed for aerospace certification. Master’s and doctoral students performing autonomy research in aerospace engineering departments (or robotics, or mechanical or electrical engineering departments that also consider autonomy topics like controls) need to take advanced courses in those topics. These students are often in a better position than their undergraduate peers to take advanced autonomy-related courses from other departments because course sizes are often smaller for advanced courses. However, the underlying assumption is that these students have had the opportunity to take foundational courses beforehand, and this can be challenging in some cases, e.g., due to non-research graduate students from the original department taking most slots in those foundational courses.

[17] https://www.abet.org/accreditation/accreditation-criteria/criteria-for-accrediting-engineering-programs-2022-2023/

[18] https://escience.washington.edu/data-science-learning/data-science-at-uw/grad-students/

4.4 Professional Workforce Continuing Education

Continuing education is crucial for professionals in the rapidly evolving field of aerospace autonomy. There's an increasing emphasis on the accreditation of distance learning programs, reflecting the growing trend of aerospace companies, particularly prime contractors, utilizing university online programs for workforce training. Online master's degrees in systems engineering, for instance, have become increasingly popular. While many of these programs are primarily offline, featuring recorded lectures for flexible scheduling, some incorporate in-person components, such as intensive weeklong sessions for team-building or hands-on activities. Alongside these longer programs, shorter online certificate options are gaining traction. These certificates offer a middle ground between brief online courses from platforms like Coursera and multi-year master's degrees, providing a more structured and recognized credential that's often preferred by employers for targeted skill development.

GAPS IN AUTONOMY AND AI FOR AEROSPACE

The AIAA Autonomy Task Force's decision to focus this paper on autonomy "at the edge" for aerospace systems resulted in significant gaps in coverage of other important issues related to autonomy and AI in aerospace. While the committee acknowledged these broader issues, they were deemed too extensive to cover in this document. This section will first address those broader gaps, followed by the two urgent gaps that were deemed to be within the narrowed scope of this paper: safety assurance and workforce development.

5.1 Inadequate Investigation of AI in Aerospace

There is a significant gap[19] in cross-industry collaborations to responsibly apply AI and ML technologies, such as generative AI for design (e.g., coding, documentation), large-scale interaction management (e.g., air or space traffic),[20] improvements to availability and quality of offline training data, data management, and improvement and personalization of K-12 through professional continuing education using AI, among others. Furthermore, a gap exists pertaining to security considerations for autonomous systems, which is amplified by AI. This is due to the skillset required to implement security often being found in a different silo than the organizations deploying autonomy to robotic systems (e.g., autonomous aerospace systems). While this paper was scoped to autonomous aerospace systems that may or may not utilize AI, the broader scope of AI for the aerospace industry is substantial.

5.2 Gaps in Safety Assurance for Autonomous Aerospace Systems

5.2.1 Gaps in Assurance of Autonomous Aerospace Systems with Machine Learning Components

The increasing integration of **machine learning (ML)** into autonomous aerospace systems presents both significant opportunities and challenges. ML enabled capabilities have potential to enhance awareness of a system state in relation to its environment, and provide effective control solutions in highly complex environments, but ensuring their safety and reliability is paramount, particularly as they evolve from advisory to safety critical roles. Strides have been made in developing certificates[21] or other methods of assuring stability, robustness, convergence in autonomous control, but more work is needed. Additionally, similar approaches for assurance of perception are needed. This section identifies related gaps in knowledge representation,

5.2.1.1 Managing complex autonomous aerospace systems is challenging due to the need for ***suitable knowledge representations*** for each function (perception, reasoning, decision making, and action) within each layer of the architecture from the mission level to the inner loop level. Different layers may use different representations, each with its own assumptions, abstractions, architectures, and programming languages. This creates a gap in deliberately defining knowledge representation and associated assumptions, amplified by the increased freedom and complexity of autonomous systems.

[19] Pramod P. Khargonekar. "Five Initiatives to Achieve AI Engineering Dominance", National Defense Magazine, 13 Dec 2024

[20] Karp, Aaron, and Ben Iannotta. "AI at Work: Mastering the Air Space," Aerospace America, Vol 62, No. 11 (2024).

[21] Wang, Li, Aaron D. Ames, and Magnus Egerstedt. "Safety barrier certificates for collisions-free multirobot systems." IEEE Transactions on Robotics 33, no. 3 (2017): 661-674.

5.2.1.2 Composing a system of many autonomous components can lead to **unpredictable emergent behavior**. In some cases, this emergent behavior is desirable; however, mechanisms need to be in place to ensure safety of the composed system.

5.2.1.3 A gap exists in the understanding of how to **modularly design hierarchical levels** in autonomous systems that can be swapped and verified while ensuring or maintaining safety and assurance.

5.2.1.4 The **challenge of verifying machine learning perception** systems necessitates smarter design architectures to ensure safe outcomes.

5.2.1.5 Mechanisms are needed that can detect and respond to failures faster online while autonomous systems are running than human capability.

5.2.1.6 The assurance of safety for a single vehicle is challenging, and these challenges increase significantly when dealing with **multi-agent and human-agent** collaborative systems. Collaboration among multiple autonomous vehicles involves challenges such as joint knowledge representation and understanding, intent understanding/prediction, negotiations, and trust in one another's intent and competence. These challenges are amplified in human-autonomy teaming due to the distinct differences in sensing, perception, reasoning, and decision-making between humans and systems. Significant challenges and gaps exist in representing knowledge to foster shared system state awareness and context, robustness to uncertainty and unknowns, resilience to failures, and continued autonomous agent functionality even without human availability.

5.2.2 Gaps in Verification and Assurance

Verifying, validating, testing, evaluating, and certifying autonomous systems to meet regulatory and process requirements are significant challenges to assuring the safety of autonomous systems.

5.2.2.1 Government acquisition programs and other large bureaucratic processes struggle to adapt to the rapid iteration of technologies like AI due to a reliance on upfront fixed requirements in lengthy programs, resulting in high costs and delays for necessary deviations. This lack of adaptability hinders the effective integration of modern autonomous systems. New methods are needed for fast iteration and development of autonomous systems.

5.2.2.2 A significant gap exists in ensuring the safety of autonomous aerospace systems due to the current reliance on standardized and generalized verification & validation (V&V) and test & evaluation (T&E) requirements. These approaches often demand considerable effort yet may not effectively demonstrate the safety of complex autonomous systems. Furthermore, some V&V and T&E methodologies offer substantially stronger evidence of safety than others, highlighting the need to avoid processes that add rigor without providing commensurate increases in safety assurance. The impact of complexity on V&V and T&E requirements necessitates a more tailored and evidence-based approach.

5.2.2.3 A gap exists in the current approach to certifying autonomy, which primarily focuses on "proving safety" rather than effectively managing risk. While airworthiness processes often aim for safety, the measurable return on safety from satisfying process rigor is often elusive. There is a need for a shift toward risk-acknowledged development, where the focus is on identifying, prioritizing, and mitigating potential threats, rather than solely attempting to prove safety.

5.2.2.4 Additional run-time assurance (RTA) methods and systems are needed, especially for systems with ML components, to complement design-time methods for system assurance. RTA design would need to consider risk and uncertainty to effectively operate within a risk-acknowledged system development process. It should also prevent excessive limiting of the operation of autonomous system features; in other words, RTA systems should not be too conservative.[22]

5.3 Gaps in Workforce Development and Training

There is a growing need for education and training in automation and autonomy at all levels (K-12, undergraduate, graduate, existing workforce, and C-suite). However, autonomy training is often done outside of aerospace engineering departments, and educational offerings are rarely integrated between computer science, robotics, electrical engineering, mechanical engineering, and aerospace engineering. Conversations among members of academia focusing on aerospace engineering report a lack of student access to courses with the desired foundational or advanced skills needed for autonomy, particularly for safety critical systems.

[22] John Schierman, et al. "Runtime Assurance for Autonomous Aerospace Systems," Journal of Guidance, Control, and Dynamics, Vol. 43, No. 12, December 2020, pp. 2205–2217; https://doi.org/10.2514/1.G004862

5.3.1 K-12 STEM Educational Gaps

K-12 STEM education often lacks sufficient focus on engineering and computer science. Although some middle school projects introduce students to the engineering design process using materials like cardboard, Legos, and basic crafts, or advanced kits with microprocessors, servos, and sensors are available, these experiences are not consistently integrated into the curriculum. The Computer Science Teacher Association has developed grade-appropriate standards,[23] but their implementation varies widely. There is no such aerospace K-12 set of standards. In many elementary schools, computer science exposure is limited to basic computer usage or typing skills, with tech lab time often rotated with other activities like physical education, music, and art. Middle and high schools may offer elective courses in computer science, robotics, or related autonomy topics, but these are not universally available and can be dependent on school or district resources. Even where offered, these topics might only be covered briefly within a core class, in a separate STEM lab, or as part of an after-school club, leading to inconsistent learning opportunities. A lack of exposure to engineering and computer science fundamentals in K-12 leaves many students unprepared for the rigors of an engineering education and unaware of the opportunities this field presents.

5.3.2 Undergraduate Aerospace, Aeronautics, and Astronautics Gaps

Faculty members in aerospace engineering departments have identified several challenges or gaps that hinder the application of curricular changes to adapt to new needs, such as those in autonomy, including:

Inertia in processes: The significant bureaucracy and time required to implement necessary curricular changes create inertia to maintain the current system. Achieving consensus on change is often challenging, with departmental groups sometimes protecting their perceived interests or resisting new disciplines that threaten familiar traditional methods.

Collaboration across disciplines: The departmental structure and business models of engineering colleges frequently create challenges for collaboration and disincentivize the creation of cross-departmental courses. As a result, students often must seek out courses in other departments. However, this creates two issues. First, even when enrollment is open to students from other departments, priority is usually given to students within the department offering the course, especially for in-demand courses like artificial intelligence and machine learning. Second, the courses offered in other departments may not align with the needs of aerospace engineering students, with prerequisites that they may not have or coverage that is not relevant or motivating.

Accreditation: While ABET has accreditation for disciplines such as aerospace engineering, there are not currently accreditation options for areas such as autonomy or data science, so no existing guidance can be leveraged to adapt an appropriate curriculum of autonomy in aerospace engineering.

5.3.3 Graduate Program Gaps

The current graduate program structure provides ample flexibility but could be improved in a few areas. These include shifting focus beyond the toy problem (i.e., a simplified or scaled-down version of a more complex issue), incorporating co-mentorship opportunities with industry or government partners, and addressing the scalability of certain academic analysis techniques for assurance. These improvements should be implemented while balancing a continued focus on fundamental discovery.

5.3.4 Professional Workforce Continuing Education Gaps

The field of automation and autonomy, especially that which uses AI and ML, is so active and changing so rapidly that it is exceptionally difficult for working professionals to gain the expertise required to manage and advance autonomous aerospace systems from development through operational use. While professional continuing education courses and certificates are helpful, hands-on collaborative opportunities on the job are lacking.

RECOMMENDATIONS

Recommendations in this section echo gaps identified earlier, within three broad categories.

Recommendation 1: Broader Exploration of AI in Aerospace

Recommendation 1.1 Deeper Exploration of Additional AI for Aerospace Topics

To fully realize AI's potential, the aerospace community should prioritize real-world research and implementation of AI-enabled capabilities, applications to large-scale aerospace challenges, and AI-specific safety and security considerations. Investigation into AI-enabled capabilities includes using large language models (LLMs) to potentially aid in areas such as software development, design, system and test documentation, test case generation, mission design, and end-user anomaly resolution support. Furthermore, it is essential to address research and development requirements for applying AI to large-scale aerospace challenges, such as air and space traffic control and management.

[23] https://csteachers.org/k12standards/interactive/

Responsible and safe AI use in aerospace can potentially accelerate the development and evaluation of new aerospace systems. As autonomy becomes mission-critical and increasingly targeted by adversaries, robust security best practices for autonomous systems become even more crucial.

Recommendation 1.2: Engagement Across Professional Societies

AIAA should collaborate with other professional societies and nonprofit organizations in areas of AI and autonomy. Limiting perspectives to just the AIAA community does not provide a complete picture of R&D needs for the entire lifecycle for autonomous aerospace systems. For example, collaborating with organizations such as the Association for Uncrewed Vehicle Systems International (AUVSI), airline trade associations such as Airlines for America (A4A), and nonprofits such as the Aircraft Owners and Pilots Association (AOPA) may help to capture perspectives and needs of end users, as well as identify new applications of existing technology. Collaboration with standards organizations such as ASTM International and nonprofits like RTCA present a path to rapidly adapt and expand standards to responsibly evaluate novel autonomy technologies. Collaboration with other professional organizations such as the Association for the Advancement of Artificial Intelligence (AAAI), Institute of Electrical and Electronics Engineers (IEEE), or the Human Factors and Ergonomics Society (HFES) will help the AIAA community stay abreast of the most recent developments in AI and autonomy technologies as well as research and best practices in human-autonomy and human-AI teaming and collaboration.

Recommendation 2: Increase Safe Autonomy Research

Recommendation 2.1: Enhancing the Safety and Reliability of Machine Learning in Autonomous Aerospace Systems

This recommendation addresses key aspects of safely using **machine learning (ML)** in autonomous aerospace systems, specifically focusing on exploring voting architectures, implementing monitor and recovery systems, and understanding the importance of clear knowledge representation in both the system architecture and human-autonomy teaming.

First, we recommend further research into **voting architectures** to enhance the reliability and resilience of ML-based autonomy. Contract-based architectures, which rely on formal performance guarantees, can be brittle in the face of unexpected situations and may not always meet the stringent accuracy requirements of safety-critical functions. Voting architectures, on the other hand, integrate machine learning with redundant systems, enabling independent analysis and aggregation of data. This approach allows for a more robust system that can leverage multiple data sources and compensate for the limitations of individual ML components.

Second, the implementation of **monitor and recovery architectures**, alongside control barrier functions, is crucial for both the safe and efficient testing of novel autonomous systems and ensuring safe behavior in unforeseen circumstances. These architectures, already used in aerospace, involve mechanisms such as input/output limiting or filtering and safe reversionary modes. Building upon this, for testing novel autonomous systems, an architecture could include monitors that detect critical ML-based perception or control failures and switch to a backup system. Furthermore, run-time assured perception or control systems can enhance the robustness of operational autonomous systems. This approach reduces the "brittleness" often associated with ML, enabling systems to continue functioning safely even when faced with unexpected inputs or scenarios.

Third, clear and consistent **knowledge representation** is essential for developing and assuring autonomous systems. This includes defining a sufficiently complete and accurate model of the autonomy system, capturing all important connections and contributors, and ensuring that associated assumptions and abstractions are well understood. Different levels of the system architecture, from high-level planning to low-level control, may require drastically different knowledge representations. For example, mode switching might be specified using formal logic, while inner-loop control may rely on physics-based states and differential equations. Understanding these representations and their interactions is vital for ensuring that the system behaves as intended.

Finally, more research is needed to develop **knowledge representation frameworks** that promote shared system state awareness in human-autonomy teams. Trust that an autonomous system will behave in accordance with an operator's intentions is critical to successful teaming. Providing humans with too much or too little information can hinder effective collaboration and erode trust. Determining what specific information about the human teammate should be provided to the autonomous agent is also an important area for investigation. While some aspects of knowledge representation may be application-specific, identifying common elements across systems could significantly improve the performance of human-autonomy teams, especially in uncertain and off-nominal situations.

In summary, enhancing the safety and reliability of ML-based autonomous systems requires a multifaceted approach. By exploring and implementing voting architectures, utilizing monitor and recovery mechanisms, ensuring clear knowledge representation within the system and for human-autonomy teaming, the aerospace community can effectively address the challenges posed by these advanced technologies and move toward safer and more robust autonomous systems.

Recommendation 2.2: "Return on Assurance" Verification Considerations

The verification and validation (V&V) of autonomous aerospace systems should prioritize a "return on assurance," much like a return on investment as well as risk management rather than proof of safety. This approach involves carefully evaluating the value gained from each V&V procedure in relation to its cost. While advanced techniques like digital engineering, digital twins, and formal methods offer significant benefits, they also require substantial resources. A one-size-fits-all strategy, often derived from traditional aviation practices, can lead to excessive and unnecessary rigor in certain areas without a clear increase in safety or mission capability. Therefore, we recommend **a tailored V&V approach**, balancing rigor with the specific system's intended use and acceptable level of risk. This tailoring can be achieved by balancing the use of requirement-based verification and validation, test and evaluation, and/or the use of operating limitations and monitor/recovery methodologies to achieve safe outcomes by design.

While development of novel assessment methods better suited for emerging autonomy technologies is highly encouraged, significant thought is required before just adding these assessments to a growing number of required evaluations for aerospace systems. The idea of applying criteria in a way to achieve proper safety assurance is more important than applying aerospace mission-worthiness processes for the sake of process. Though most aerospace best practices and level of rigor have been designed with safety in mind, the actual measurable return on safety from satisfying the process rigor is elusive. As a result, we recommend the use of pathfinder programs that coordinate with certification authorities to use a new set of criteria tailored to that class of autonomous aerospace systems.

A **risk-acknowledged development approach** is vital for prioritizing assurance efforts. Rather than solely aiming to prove safety, the focus should shift to identifying, managing, and mitigating risks. Drawing inspiration from the National Institute of Standards and Technology's (NIST) Risk Management Framework (RMF), a Risk Acknowledgement and Mitigation Plan could be developed for aerospace autonomy systems. This plan would outline potential risks and the strategies to address them, ensuring that assurance resources are allocated to the most critical areas. Certification criteria, often developed in response to past failures, are important, but should not be the only driver of assurance. Operational risks, in particular, should be managed through operational limitations rather than solely through increased design rigor. A system with lower design assurance can still achieve a high level of operational safety through carefully defined operational boundaries.

Additionally, **segmenting performance controllers and safety controllers** may provide guaranteed safe autonomous behavior in a more cost-effective manner. In some cases, the technology to conduct "safety proofs" on highly complex autonomous systems with large state spaces and neural network controllers may not exist or be mature enough for operational use. However, pairing these autonomous systems, with simpler safety-focused run time assurance bounds or shields, may make certification feasible by placing the burden of proof on the run-time safety assurance mechanisms. These RTA mechanisms would need to be designed considering risk and uncertainty and to prevent enforcement of excessively conservative limitations on the operation of autonomous system features.

In summary, a strategic, risk-informed, and tailored approach to V&V is essential to ensuring the safe and successful deployment of autonomous aerospace systems.

Recommendation 2.3: Aerospace Community, Meet the MVP!

Government acquisition programs and other large bureaucratic processes that emphasize development of requirements at the beginning of many-year programs and punish any deviation with high costs will not scale to the fast iteration of technologies like AI that are powering modern autonomous systems.

In the last 10 years, the concept of a **minimum viable product (MVP)** has taken off within the lean startup and software communities. Lean startup product development processes combine business-driven (or mission-driven) hypothesis experimentation with iterative product releases. A key element is the MVP, and while there is no standard definition, most definitions include some concept of investing the minimum amount possible to quickly prototype a product and get feedback from users or customers before iterating.[24] The MVP concept can also be related to the way technology has traditionally been introduced into aviation, where initial use cases are identified to control risk, and to allow early instantiations of new technology to be introduced safely at an appropriate level of rigor for an initial use case. Then the technology is allowed to mature for more critical use cases based on these initial lessons learned.

Especially for technologies below technology readiness level 7, which is a prototype test in an operational environment (e.g., flight test or in-space test), an MVP approach may enable a team to push a novel autonomous aerospace system through design phases much faster. A minimum viable product can be defined by several factors, such as technical maturity level, expected safety level for its intended use, and level of rigor in the airworthiness process. The MVP concept is beneficial because it enables a quick maturation of a minimum integrated capability through some form of flight test, to find issues and correct assumptions quickly and reduce rework, rather than assuming that the requirements are correct at the beginning of a program (which they never are), and jumping into a complex design right off the bat.

Recommendation 3: Improve Workforce Development and Training

Recommendation 3.1: Develop Content for K-12 Classrooms

Given the inconsistent availability of dedicated autonomy-relevant courses in K-12 education, we recommend that members of the aerospace autonomy community develop accessible training materials that can be readily integrated into core subjects like math or science, used in dedicated STEM labs or computer labs, or offered through after-school clubs. To harness the full talent pool of American students, adding introductions to engineering in the science curriculum in elementary and middle school are critical. Additional emphasis should be placed on the creativity, teamwork, and communication skills needed in engineering. Highlighting the problem-solving and innovation aspects of engineering, as well as the potential career paths and stability of engineering jobs, will further encourage students to pursue this field, maximizing the future talent pool and providing a strong foundation for undergraduate engineering courses.
In addition to the topics mentioned above, it is important to give students foundational skills that can be introduced in K-12 toward the study of autonomous systems include:

- Critical thinking and logical reasoning
- Linear algebra and basic mathematical modeling of systems like free body diagrams in physics
- Physics and basic dynamics
- Computer programming, including computer (boolean) logic and flow diagrams
- Basics of electrical engineering (e.g., electronics, sensors, signals and systems, and information sharing and networking)

Recommendation 3.2: Highly Valued Undergraduate Courses for Students Specializing in Aerospace Autonomy.

Given the challenges discussed earlier that prevent fast and important changes to core undergraduate curricula, an approach based on a portfolio of measures is recommended. These recommendations are primarily intended for students who are looking for which electives are most advantageous, as well as university curriculums looking to add more autonomy-centric electives.

Portfolio of measures

- Probability and statistics: Every engineering undergraduate student in the nation, aerospace engineering or not, should graduate with basic knowledge of probability and statistics. This could be implemented in different ways for different departments, ranging from a dedicated required course to several modules across multiple courses.
- Programming and software engineering skills: Aerospace engineering students should graduate with a basic understanding of data structures and algorithms and the capability proficiency to develop a relatively complex computer program from a set of functional requirements. A single introductory course in freshman year is insufficient in most cases for this purpose. This could again be implemented via a required course or more likely by adding some modules across several courses that include both lectures and programming assignments.

[24] Lenarduzzi, Valentina, and Davide Taibi. "MVP explained: A systematic mapping study on the definitions of minimal viable product." In 2016 42th Euromicro Conference on Software Engineering and Advanced Applications (SEAA), pp. 112-119. IEEE, 2016.

- Linear algebra, statics & dynamics (applied physics), and flight controls courses that cover topics such as Laplace transforms and signals and systems could be necessary depending on application of autonomy. While most of these courses are required in most aerospace engineering departments, linear algebra appears to be a glaring exception in some cases.
- Classical AI/autonomy, machine learning: The bulk of the curriculum in autonomy including classical AI (intelligent agents that search, reason, and plan) and machine learning (supervised, unsupervised, and reinforcement learning) can be delivered as technical electives for students who wish to specialize in these areas. However, some AI concepts can be added to existing core courses. For example, an elective course that covers reinforcement learning, neural networks, genetic algorithms (and other metaheuristic optimization approaches), as well as other AI concepts applied to autonomous system development, would be invaluable
- Autonomy-centric capstones: Finally, an essential ingredient of this portfolio of measures is capstone design. Autonomy-related capstone projects can give students the opportunity to apply the skills and knowledge learned through the elements listed above to a real-world problem. These autonomy capstone projects can also foster fruitful collaborations between industry or government and academia.

Recommendation 3.3: Highly Valued Graduate Courses for Graduate Students Specialization in Aerospace Autonomy.

As noted in the gaps, there is significant flexibility in the autonomy curriculum to gain deep understanding of the foundational topics listed in the undergraduate section in addition to:

- Discrete mathematics, including combinatorics, set theory, relations, finite state machines, and graph theory
- Formal methods including propositional logic, first-order logic, and other logics such as temporal logic or description logics, and their application to automated reasoning, V&V, run-time assurance, etc.
- Probability and statistics, including stochastic processes, simulation, Bayesian and non-Bayesian estimation, hypothesis testing, inference and information theory
- Linear and nonlinear systems, including signals, transfer functions, frequency domain, state-space representation, Kalman filters, model predictive control, adaptive control, control barrier functions, and among others
- Human factors: Even autonomy students who are not doing research in human factors should have at least some awareness of basic human-autonomy teaming concepts, including levels of automation, workload, determining system state awareness, trust, etc.
- Software engineering: Structured development processes, software reliability estimation and reliability improvement techniques, reliable real-time and distributed computing, and cybersecurity considerations
- Classical AI: Search, constraint satisfaction, game theory, logical agents, planning agents, probabilistic reasoning, Markov decision processes, dynamic programming
- Supervised and unsupervised learning, covering the theory and application of basic methods from linear models to deep neural networks
- Reinforcement learning, starting from the foundations (bandits, dynamic programming, Monte Carlo, tabular methods) and getting to modern deep reinforcement learning methods.

Recommendation 3.4: Co-Mentoring of Graduate Students

To promote research beyond simplified academic applications to techniques that scale to real-world applications, as well as improve continuing education of professionals, a co-mentoring approach is recommended. The aerospace community, including government, industry, FFRDC/UARC, and professional organizations like AIAA, establish and support a (formal or informal) co-mentoring program or culture for graduate students. This program would pair graduate students with professionals to provide practical insights and industry context to students, as well as exposure to emerging technologies for professionals. While advisory boards are valuable, active co-mentoring, especially within fellowship or internship programs (e.g., DOD SMART, NDSEG, NASA's NSTGRO, or the Draper's Fellows program), can further enhance student development and address industry's need for job-ready graduates. This initiative should extend beyond AIAA to encompass the broader aerospace community.

ANNEX

Existing Aerospace Standards and Recommended Practices

Several recommendations and regulatory documents exist describing system design and safety standards such as Aerospace Recommended Practice (ARP) from the Society of Automotive Engineers (SAE) International, RTCA guidance, American Society for Testing and Materials (ASTM) standards, the European Union Aviation Safety Agency (EASA) as well as military and federal standards, handbooks, and regulations. Where appropriate, the following recommendations, standards, handbooks, and guidance should be consulted:

- ARP4761: Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment[25]
- ARP4754: Guidelines For Development Of Civil Aircraft and Systems[26]
- RTCA DO-254: Design Assurance Guidance for Airborne Electronic Hardware[27]
- RTCA DO-178: Software Considerations in Airborne Systems and Equipment Certification[28]
- RTCA DO-333: Formal Methods Supplement to DO-178C and DO-278A[29]
- ASTM 3269-21: Standard Practice for Methods to Safely Bound Flight Behavior of Unmanned Aircraft Systems Containing Complex Functions[30]
- MIL-HDBK-516C: Airworthiness Certification Criteria[31]
- MIL-STD-882E: System Safety[32]
- FAA's Standard Airworthiness Certification Regulations[33]
- EASA Concepts of Design Assurance for Neural Networks (CoDANN)[34]
- EASA Concepts of Design Assurance for Neural Networks (CoDANN) II[35]
- NASA NPR 8705.2C Human-Rating Requirements for Space Systems

[25] "Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment," SAE International, Guidance, 1996.

[26] "Guidelines For Development Of Civil Aircraft and Systems," SAE International, Guidance, 2010.

[27] "Design Assurance Guidance for Airborne Electronic Hardware," RTCA Inc., Guidance,2000.

[28] "Software Considerations in Airborne Systems and Equipment Certification," RTCA Inc., Guidance, 2011.

[29] "Formal Methods Supplement to DO-178C and DO-278A," RTCA Inc., Guidance, 2011.

[30] "Standard Practice for Methods to Safely Bound Flight Behavior of Unmanned Aircraft Systems Containing Complex Functions," ASTM International, Standard, 2017.

[31] "MIL-HDBK-516C: Airworthiness Certification Criteria," Department of Defense, Guidance, Dec. 2014.

[32] "MIL-STD-882E: System Safety," Department of Defense, Standard Practice, May 2012.

[33] Federal Aviation Administration (FAA), Standard airworthiness certification regulations: Title 14, code of federal regulations, Website, https://www.faa.gov/aircraft/air_cert/airworthiness_certification/std_awcert/std_awcert_regs/regs/, Oct. 2017.

[34]https://www.easa.europa.eu/en/document-library/general-publications/concepts-design-assurance-neural-networks-codann

[35]https://www.easa.europa.eu/en/document-library/general-publications/concepts-design-assurance-neural-networks-codann-ii

AIAA Groups Working in Autonomy

Autonomy has significant breadth that spans multiple technical committees and task forces, while not necessarily a complete list, the following AIAA Technical Committees hosted autonomy sessions during the 2025 AIAA SciTech Forum:

- Intelligent Systems
- Guidance, Navigation, and Control
- Modeling and Simulation Technologies
- Software
- Uncrewed and Autonomous Systems
- Space Automation and Robotics
- Computer Systems
- Aircraft Operations
- Transformational Flight
- Digital Avionics
- Human-Machine Teaming
- Electrified Aircraft Technology

AIAA Autonomy Documents and Meetings

- 2022 AIAA Summits
  - Space Autonomy Summit Report: https://www.aiaa.org/resources/space-autonomy-summit-report
  - Air Autonomy Summit
- AIAA Intelligent Systems Roadmap: https://aiaa-istc.github.io/IS_roadmap.html