



AFRL-RY-WP-TP-2023-0110

**SAFETY, TRUST, AND ETHICS CONSIDERATIONS FOR
HUMAN-AI TEAMING IN AEROSPACE CONTROL
(Preprint)**

**Kerianne L. Hobbs and Bernard Li
Autonomous Capability Team 3 Branch
Sensors Plans & Advanced Programs Division**

**JUNE 2023
Final Report**

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.

See additional restrictions described on inside pages

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
SENSORS DIRECTORATE
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7320
AIR FORCE MATERIEL COMMAND
UNITED STATES AIR FORCE**

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE June 2023	2. REPORT TYPE Conference Paper Preprint	3. DATES COVERED	
		START DATE 8 June 2023	END DATE 8 June 2023
4. TITLE AND SUBTITLE SAFETY, TRUST, AND ETHICS CONSIDERATIONS FOR HUMAN-AI TEAMING IN AEROSPACE CONTROL (Preprint)			
5a. CONTRACT NUMBER In-House	5b. GRANT NUMBER N/A	5c. PROGRAM ELEMENT NUMBER N/A	
5d. PROJECT NUMBER N/A	5e. TASK NUMBER N/A	5f. WORK UNIT NUMBER N/A	
6. AUTHOR(S) Kerianne L. Hobbs and Bernard Li			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Autonomous Capability Team 3 Branch Sensors Plans & Advanced Programs Division Air Force Research Laboratory, Sensors Directorate Wright-Patterson Air Force Base, OH 45433-7320 Air Force Materiel Command, United States Air Forces			8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory, Sensors Directorate Wright-Patterson Air Force Base, OH 45433-7320 Air Force Materiel Command, United States Air Forces		10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RYZA	11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RY-WP-TP-2023-0110
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.			
13. SUPPLEMENTARY NOTES PAO case number AFRL-2023-1934, Clearance Date 8 June 23. This is a work of the U.S. Government and is not subject to copyright protection in the United States. Report contains color.			
14. ABSTRACT Designing a safe, trusted, and ethical AI may be practically impossible; however, designing AI with safe, trusted, and ethical use in mind is possible and necessary in safety and mission-critical domains like aerospace. Safe, trusted, and ethical use of AI are often used interchangeably; however, a system can be safely used but not trusted or ethical, have a trusted use that is not safe or ethical, and have an ethical use that is not safe or trusted. This manuscript serves as a primer to illuminate the nuanced differences between these concepts, with a specific focus on applications of Human-AI teaming in aerospace system control, where humans may be in, on, or out-of-the-loop of decision-making. What safety, trust, and ethics mean in this context is explored with a focus on Human-AI teams in control of aerospace systems.			
15. SUBJECT TERMS ACT3, AI, AI application, artificial intelligence			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified	
			18. NUMBER OF PAGES 20
19a. NAME OF RESPONSIBLE PERSON Kerianne L. Hobbs			19b. PHONE NUMBER (Include area code) N/A

Safety, Trust, and Ethics Considerations for Human-AI Teaming in Aerospace Control

Kerianne L. Hobbs *

Air Force Research Laboratory, Wright-Patterson AFB, OH, 45433

Bernard Li†

Air Force Research Laboratory, Wright-Patterson AFB, OH, 45433

Designing a safe, trusted, and ethical AI may be practically impossible; however, designing AI with safe, trusted, and ethical use in mind is possible and necessary in safety and mission-critical domains like aerospace. Safe, trusted, and ethical use of AI are often used interchangeably; however, a system can be safely used but not trusted or ethical, have a trusted use that is not safe or ethical, and have an ethical use that is not safe or trusted. This manuscript serves as a primer to illuminate the nuanced differences between these concepts, with a specific focus on applications of Human-AI teaming in aerospace system control, where humans may be in, on, or out-of-the-loop of decision-making. What safety, trust, and ethics mean in this context is explored with a focus on Human-AI teams in control of aerospace systems.

I. Introduction

ON May 11th, 1997, Garry Kasparov became the first World Chess Champion to lose to a computer, IBM's DeepBlue. A year later, Kasparov created a modified version of chess called "advanced chess," where instead of a human playing against a human, a human-computer team would play against another human-computer team [1]. In the first advanced chess tournament hosted in 2005, the person on the winning team was neither a grandmaster nor was the algorithm on the winning team the most advanced. Instead, an amateur who was more trusting of their algorithms to make the right decisions due to his inexperience, while also training with the algorithm to recognize its flaws and weaknesses, took home the trophy with 3 relatively weak laptops that he could manage, each giving separate recommendations [1]. Kasparov, and many others, realized that characteristics such as trust played a large role in human-algorithm team performance. Pairing humans with increasingly intelligent algorithms may result in significant performance gains over the next few years, especially highly capable deep learning models.

Artificial intelligence (AI), is defined as the ability of a computer or robot to perform tasks that normally require human intelligence [2], and *machine learning* (ML), is defined as the use and development of computer systems that are able to learn and adapt without following explicit instructions, by using algorithms and statistical models to analyze and draw inferences from patterns in data [3]. In ML, a *model* is a fixed structure (such as a deep neural network) with modifiable parameters that encodes an input-output relationship [4]. Within ML, there are three broad categories: supervised, unsupervised, and reinforcement learning. In *supervised learning*, labeled input/output data is used to train a model to generate the correct output for any input [4]. *Unsupervised learning* uses unlabeled data to find trends or clusters of output data. In *reinforcement learning* (RL), "learning to take actions," the model is referred to as an *agent* that learns a sequence of actions (outputs) that returns a maximum reward. RL is based on the concept of operant conditioning where desirable behavior is reinforced with rewards (positive reinforcement) and bad behavior is penalized (negative reinforcement), and the agent learns which actions to take to maximize reward through trial and error [5]. Applied to aerospace or robotics, AI and ML applications (in particular, RL) are often embodied as an *autonomous system*, which is a cyber-physical system that is given delegated and bounded authority by a human to make decisions and take actions in the world.

While there are cases in which a model can be trained from optimal lookup table data to decide actions for an aerospace system to take using supervised learning, such as the neural network compression of ACAS-XU [6], RL approaches have gained much more attention, in part due to similarities to control theory. RL and optimal control share the concept of a system *state* (variables that describe the plant / *RL environment*, such as position and velocity), a notion of *partial observability* (where a measurement / *RL observation* can be used to estimate the state), a discrete or

*Safe Autonomy Lead, Autonomy Capability Team 3, 2241 Avionics Circle.

†ML-RCP Intern, Autonomy Capability Team 3, 2241 Avionics Circle

continuous control input / *RL action* is determined by a control law / *RL policy* (a function of the current state or partial observation that is optimized to minimize a cost function in controls or maximize a *reward function* in RL). These similarities are summarized in Table 1. In aviation, RL has the potential to improve air-based wildfire monitoring [8],

Table 1 Comparison of Control Theory and Reinforcement Learning Concepts [7]

Control Theory	Reinforcement Learning
plant	environment
state	state
measurement	observation
controller	agent
control law	policy
control input	action
cost function	reward function

urban air mobility traffic deconfliction [9], air traffic flow [10], landings [11], maintenance scheduling [12], acrobatic maneuvering [13], control system designs [14–19], terminal area operations and planning [20], military engagement scenarios such as the AlphaDogfight [21], and other areas. In the spacecraft domain, RL has been applied to research models for tasks such as hovering near asteroids [22], ground observation missions [23], proximity operations and docking [24–26], autonomous spacecraft maneuvering [27, 28], autonomous imaging and mapping of asteroids [29], planetary landing [30, 31], and other applications.

Safe, trusted, and ethical use of AI are of particular concern because contrary to many applications of AI and ML, aerospace systems feature high safety and reliability requirements. For example, reliability requirements are on the order of 10^{-9} catastrophic fault rate for civil aircraft avionics [32]. The authors argue that anthropomorphizing AI to have qualities such as ethical behavior, places unrealistic expectations on AI. Designing a safe, trusted, or ethical AI system may be practically impossible. However, safe, trusted, and ethical *use* of an AI system from ideation, through development and operational use is feasible. Safety, trust and ethics are related emergent properties of human-AI teams. *Safety* is freedom from harm during operations, *trust* is a willingness to accept vulnerability in situations characterized by uncertainty [33], *ethics* in this context are rules created by societies and cultures governing moral and just usage of technology, and *emergent properties* are characteristics that an entity gains when it becomes part of a larger system. This publication explores safety, trust, and ethics considerations that should be taken into consideration for an AI that makes decisions and controls behavior of an aerospace system in a team with a human.

II. Humans In, On, and Out-of-the-Loop

Control systems and AI models that take actions based on observations of the state, make decisions in a discrete loop. The relationship between humans and that decision loop, is often described as in, on, or out-of-the-loop, and is depicted in Figure 1. For *human-in-the-loop* systems, no action is taken without the approval of a human. Actions are taken at a slow enough rate that both the human and the autonomous agent know the state of the system, a human provides guidance to an autonomous agent, and the human has time to evaluate the recommended action from the autonomous agent, and approve the action that is sent to the system to execute. Human-in-the-loop teams are more appropriate for discrete actions where long periods of time are available between decisions, such as planning a trajectory. *Human-on-the-loop* systems feature a human that gets status updates from the autonomous agent and the state of the state of the system under control, and provides guidance to the autonomous agent as it takes independent actions. This category is appropriate in situations in which human decision making may not be fast enough, human approval of every decision is not required, and a human supervisory role makes more sense. *Human-out-of-the-loop* teams feature an autonomous agent that observes the state and takes actions, perhaps with some initial guidance from a human operator, but otherwise independently. This approach may be appropriate in situations where regular communication with a human operator is not possible. More research is needed to explore differences in the safety, trust, and ethical principles applied to designing models that make decisions and take actions. The rest of this manuscript considers decisions made at a variety of levels of human involvement.

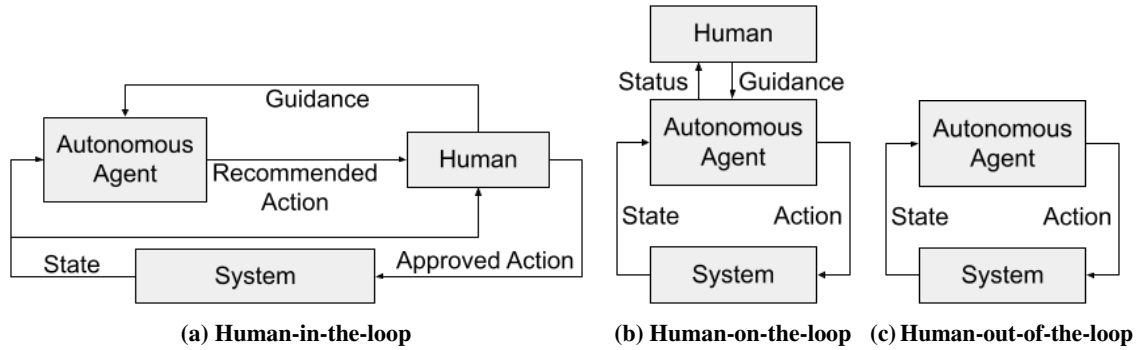


Fig. 1 Comparison of human-AI teaming for control in which the human is in, on, or out-of-the-loop.

III. Safety

Although trusted and ethical use of AI in a human-AI team, may be difficult to define formally, safety has been a huge part of aerospace for more than a century, and can be formally evaluated with respect to certification criteria, risk definitions, and design principles. The Center for Security and Emerging Technology defines AI safety as a "an area of machine learning research that aims to identify causes of unintended behavior in machine learning systems and develop tools to ensure these systems work safely and reliably," and defines safety AI issues in categories of specification, assurance, and robustness [34]. Specification describes the challenge of describing the intended system behavior, sometimes in the form of a reward function during training, and the opportunity for misspecification that leads to unintended behavior. Assurance is analogous to interpretability of the AI in the definition, and robustness refers to the ability of the AI to behave reliably in a wide range of decisions. While these issues are important to human-AI teams in aerospace, they are not sufficient. This section describes several other important considerations for human-AI teams in aerospace control.

A. Standards for Aerospace Safety

Several standards provide guidance for assuring safety in the design of aerospace systems, including Aerospace Recommended Practice (ARP) from the Society of Automotive Engineers (SAE) International, Radio Technical Commission for Aeronautics (RTCA) guidance, American Society for Testing and Materials (ASTM) standards, as well as military and federal standards, handbooks, and regulations. Where appropriate, the following standards, handbooks, and guidance should be consulted:

- ARP4761: Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment [35]
- ARP4754: Guidelines For Development Of Civil Aircraft and Systems [36]
- RTCA DO-254: Design Assurance Guidance for Airborne Electronic Hardware [37]
- RTCA DO-178: Software Considerations in Airborne Systems and Equipment Certification [38]
- RTCA DO-333: Formal Methods Supplement to DO-178C and DO-278A [39]
- ASTM3269-21 (update from ASTM 3269-17): Standard Practice for Methods to Safely Bound Flight Behavior of Unmanned Aircraft Systems Containing Complex Functions [40, 41]
- MIL-HDBK-516C: Airworthiness Certification Criteria [42]
- MIL-STD-882E: System Safety [43]
- FAA's Standard Airworthiness Certification Regulations [44]

Research groups and the public have raised concerns about developing AI based on big data by including equitable samples across ages, race, gender, and other variables. Similar considerations are required for developing AI that will be placed in control of robotic systems such as cars or aircraft. Even in aerospace, roadmaps such as Concepts of Design Assurance for Neural Networks (CoDANN) I and II [45, 46] have primarily focused on supervised and unsupervised learning applications such as machine perception.

B. Safety by Design

The Automatic Ground Collision Avoidance System (Auto GCAS) and the Automatic Air Collision Avoidance System (Auto ACAS) [47–53] are interesting case studies here in two separate contexts. First, it is an example of a run

time assurance system, and second it is arguably an example of the successful use of Human-autonomy teaming (there is no machine learning or neural network in the example though). *Run time assurance* systems monitor the output of an advanced, primary controller and intervene when necessary to assure safety [54], as depicted in Figure 2. Until technology advances enough to prove properties of neural network control systems, run time assurance is the most likely near term path to certify advanced control systems. The ASTM3269-21 [41], and its predecessor are the first examples of run time assurance certification standards.




Fig. 2 A generalization of a Run Time Assurance (RTA) architecture

A generalization of a Run Time Assurance (RTA) architecture is shown in 2, in which a black-box complex function, such as a trained neural network, produces an output based on an input from another gray-box system or user, a white-box RTA system that may modify the output based on its observation of the input to produce a satisfactory output (if the original complex system output satisfies desired properties, it will be passed through as the satisfactory output with no modification). Here, a black-box system can only be viewed in terms of inputs and outputs while the internal mechanisms are opaque, white-box system inner workings are fully known, and gray-boxes internal mechanisms are partially known. Note that while the weights, biases, activation functions, connections, and other features of neural networks can be known, it is generally argued that their combination and interactions are so complex as to effectively be a black box to a human user.

In Auto GCAS and Auto ACAS a human provides the primary performance control to complete a mission, while an autonomous system monitors the state of the aircraft relative to the ground or other aircraft, and intervenes when necessary to avoid a collision. Auto GCAS was transitioned to the F-16 fleet with digital flight control computers in fall of 2014 and has since been credited with saving the lives of 13 pilots in 12 F-16s. It is also a unique aerospace system that takes control from a pilot in life or death situations, is the result of more than three decades of research, and has been studied extensively by psychologists as a human-automation teaming example [51, 55].

While there is debate whether trust [56] and ethics [57] can be built into a system, safety is agreed upon to be a necessary design criteria of any autonomous robotic system interfacing in real-world scenarios, especially in proximity to humans. Auto GCAS [58] which has several critical design principles that contribute to safety, as follows.

- understandable response: the human teammate should have an acceptable understanding of what the autonomous agent is doing, e.g. in Auto GCAS the response is a simple roll to wings level and 5g pull maneuver.
- appropriate transparency: visual or auditory explanations of the AI system perceptions should be balanced, so they do not result in overtrust of the system or create a distraction or ignored nuisance to the operator.
- nuisance freedom: the responses of the AI agent should not be so excessive that they compromise performance of the team, e.g. in Auto GCAS, the system would only maneuver at a point in time beyond when a human pilot would maneuver to reduce possible mission impacts.
- high level requirements precedence: the human-AI team should agree on the order of high level requirements, e.g. on Auto GCAS the requirements were in the order of "do no harm, do not interfere, prevent collisions," in which the primary task of the AI was placed at a lower priority than not interfering with the mission or causing harm to the human or system under control.
- system wide integrity management - one of the biggest keys to ensuring safety of the system is recognizing and accommodating failures. Most aircraft are triple or quad redundant and use voting schemes to identify failures. That said, AI, depending on its form, can be brittle and fail catastrophically in non-traditional ways, so redundancy may not be the best antidote. Auto GCAS uses a hardware and software checks to ensure safety of maneuver requests.
- variable risk tolerance - different missions may have different definitions of what constitutes "nuisance" AI, so an ability to modify parameters, such as those related to risk are necessary.

C. Risk

The working definition of trust includes a willingness to be vulnerable, which is related to safety and risk tolerance. Risk is defined by government organizations like the military and NASA in terms of likelihood and severity of particular risks. Definitions of probability and severity in these contexts are provided in Tables 2-3. These probabilities and

severities are then used to assign a risk level via a risk matrix such as Figure 3a for the military and Figure 3b for NASA. Also of relevance to AI and autonomy software is the U.S. Military definition of software safety criticality levels: no safety impact, influential, redundant fault-tolerant, semi-autonomous, and autonomous control.

Table 2 Probability Categories from MIL-STD-882E [43] and NASA S3001 [59]

882E Level	882E Criteria	S3001 Level	S3001 Criteria	S3001 Probability
A	Likely to occur often in the life of an item.	5	Near certainty	$p > 80\%$
B	Will occur several times in the life of an item.	4	Highly Likely	$80\% > p > 60\%$
C	Likely to occur sometime in the life of an item.	3	Likely	$60\% > p > 40\%$
D	Unlikely, but possible to occur in the life of an item.	2	Low likelihood	$40\% > p > 20\%$
E	So unlikely, it can be assumed occurrence may not be experienced in the life of an item.	1	Not likely	$20\% > p > 0\%$
F	Incapable of occurrence. This level is used when potential hazards are identified and later eliminated.			

Table 3 Severity categories used to assign risk under MIL-STD-882E [43]

Description	Severity	Mishap Result Criteria
Catastrophic	1	Death, permanent total disability, irreversible significant environmental impact, or monetary loss \geq \$10M.
Critical	2	Permanent partial disability, injuries or occupational illness that may result in hospitalization of at least three personnel, reversible significant environmental impact, or monetary \geq \$1M and $<$ \$10M.
Marginal	3	Injury or occupational illness resulting in one or more lost work day(s), reversible moderate environmental impact, or monetary loss \geq \$100K and $<$ \$1M.
Negligible	4	Injury or occupational illness not resulting in a lost work day, minimal environmental impact, or monetary $<$ \$100K.

D. Cross-domain Human-Automation Interaction Safety Examples

A critical section of designing safety with autonomous systems that cannot be ignored is the interaction of the algorithm with the human. In particular, automation complacency (or appropriate reliance) and poor transparency of the system operation to human users have contributed to major accidents across safety critical industries.

1. Automation Complacency

The NTSB Highway Accident Report on the Collision between Vehicle Controlled by Automated Driving System and Pedestrian [60], which was a landmark for being the first ever investigation of a human death by self-driving vehicle, discovered one of the major factors leading to the crash was *automation complacency*. Automation complacency is first

RISK ASSESSMENT MATRIX				
SEVERITY PROBABILITY	Catastrophic (1)	Critical (2)	Marginal (3)	Negligible (4)
Frequent (A)	High	High	Serious	Medium
Probable (B)	High	High	Serious	Medium
Occasional (C)	High	Serious	Medium	Low
Remote (D)	Serious	Medium	Medium	Low
Improbable (E)	Medium	Medium	Medium	Low
Eliminated (F)	Eliminated			

(a) Military Risk Assessment Matrix [43]

L I K E L I H O O D	5	7	16	20	23	25
	4	6	13	18	22	24
	3	4	10	15	19	21
	2	2	8	11	14	17
	1	1	3	5	9	12
		1	2	3	4	5
		CONSEQUENCE				

(b) NASA Risk Assessment Matrix [59]

Fig. 3 NASA and Military risk assessment matrices.

described Parasuraman paper [61], and is a form of overtrust under the Lee and See category of appropriate reliance [62]. *Appropriate reliance* is related to *calibrated trust*. When a user overtrusts an AI or Automation solution, they can become complacency and use the AI in a context that it was not designed for. In cases where a human undertrusts AI, it may lead to a situation where the human stops using the AI in a context where it is better suited. Both under and overtrust can contribute to safety violations. The NTSB concluded that autonomous vehicles had high levels of automation complacency, and so began recommending all automated vehicle manufacturers monitor their drivers for attentiveness [60]. In response car manufacturers, notably Tesla, Ford and GM, began to implement torque sensors, which required a human to hold the steering wheel, and are testing driver-facing cameras to detect if drivers become inattentive [63]. The debate over whether a system should monitor its users comes down to the culture of the operator and the kind of trust between the operator, which can be explored in future research.

2. Appropriate Transparency

Designing appropriate transparency that communicates system information and alarms in human-autonomy teams can be challenging. On March 28th, 1979, at 4:00AM, the Three Mile Island nuclear reactor began a series of events which led to America's worst nuclear meltdown. The US Presidential Commission on the Accident at 3 Mile Island found 3 major factors that caused the meltdown, one of which was not a design flaw in the reactor, but in the operator console [64]. The Three Mile Island plant was one of the most complex computer controlled systems of its time, but the investigation found that the system did not give clear warnings of incidents, or explain the system was doing in response [64]. Another example of the need for transparency came from the unfortunate set of crashes in 2019 related to the Boeing 737 Max 8. The NTSB noted in their Safety Recommendation Report [65] that one of the primary failures leading to the crashes was improper training of flight crew and withholding of information to flight personnel [65]. Clear and appropriate transparency, combined with good operator training, is a requirement in automated and complex situations.

IV. Trust

Within the trust in automation community, trust is defined as "the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability" [33]. Trust is inherently an evaluation of whether the benefits outweigh the risk. In teaming scenarios, trust may occur between humans, between AIs or between humans and an AI. One of the biggest challenges is ensuring that trust between two agents (human or AI) is calibrated. Overtrust may result in relying on an agent in a scenario they are ill-suited for. Undertrust may result in lower system performance when the agent is not used in a situation they are better suited to handle. Each agent, whether human or automation must have a good representation of the strengths and weaknesses of the partner, which may be gained through experience.

A. Trust Dynamics

Trust is a dynamic rather than a static process. When it comes to humans and technology, the general default assumption is that the automation is perfect, initially resulting in overtrust. However, any mistake by the automation quickly erodes that trust [66]. Additionally, high false alarm rates or nuisance alerts or activations of AI can result in the system being turned off [55, 67]. In one study of pilot trust in an automatic ground collision avoidance system, a pilot indicated that one mistake by the automation would likely cause pilots to turn the system off [55].

B. Transparency

To facilitate trust between human-human, human-AI, and AI-AI systems, each agent needs some model of the cooperating agent. Transparency could include an explanation of the agent's rationale, facilitate understanding of the agent's behavior, and highlight shared understanding of interaction [68, 69]. As operators become familiar with a system and see its operation is consistent with how its function was explained to them, they gain positive impressions [51].

C. Human-Automation Trust Variability

Trust in automation varies from human to human. One organization categorizes trust is to categorize it as situational, learned, or dispositional trust [70]. *Situational* trust describes an individual's level of trust for a specific context. For example, an operator may rely on automation more when overwhelmed with a high workload [71]. *Learned* trust is based on experiences with the specific automated system [70], and changes over time [72]. *Dispositional* trust is an individual's propensity to trust automation may in some cases be cultural or personality driven. For example, studies have found that citizens of Mexico are more likely to trust automated decision aids than citizens of the United States [73], older adults are more likely to have better calibrated trust of decision aids [74, 75] and are less likely to be swayed by a picture of a human expert in the automation interface [76], women and men respond different to flattery and communication styles from automation [77–79]. Trust can also be influenced by different personality categorizations. People are more likely to trust if they are extroverted [80] or intuitive [81], and less likely if they have neurotic [82] or sensing [81] personalities.

D. Anthropomorphizing Autonomy

Trust may also be influenced by a tendency to anthropomorphize AI and robots. While science fiction has long been concerned with anthropomorphizing robots [83], real world examples of humans anthropomorphizing robots grew in the early 2000s. The first documented case of this is with iRobot's Roomba platform, a robot vacuum which user gave positive traits and characteristics due to its function and proximity to users [84]. Not only did users want to personalize their unique Roomba, operators assigned Roombas specific names and considered them similar to house pets, and became more willing to tolerate more extraneous behavior, compared to other cleaning robots [84] due to their nonthreatening nature and positive contribution to the household. The Army Research Laboratory has expanded this concept further, finding similarities in Human-Autonomy Teaming and Human-Animal Trust [85]. If the robot interacted with humans in a friendly way on occasion, if the errors it made were not severe, if the operator underwent training with the robot, if it was "cute" in either behavior or design, and if the environment the human and robot operated in was non-threatening, operators were more likely to accept unpredictable behavior [85].

E. Stakeholder Relationships

Multiple studies have identified that mutual understanding and respect between stakeholders of a technology impact trust. Stakeholders may include researchers, regulators, program managers, design engineers, system acquirer, insurers, commander or supervisor, operator, patron, the community, and in some cases the science team and science community [55, 86]. Each group features risks such as personal safety, property and finances, mission effectiveness, reputation, job security, and public trust. During the design process, researchers, program managers, and design engineers are accountable for designing the system to perform well, are vulnerable to customer backlash including loss of reputation and future resources, and expect the system to be used as intended. Regulators and acquirers are responsible for accomplishing mission goals, vulnerable to a loss of required resources to complete design and evaluation, and expect the design to meet operator needs. Operators, patrons, the community, and in some cases, the scientific community, are accountable for using tools as intended to successfully complete the mission, vulnerable to failure, which could include loss of priceless science data, property, finances, reputation or life, and expect the AI system to perform acceptably and free from errors or biases that could call analysis of mission data into question [55, 87].

V. Ethics

In this section, we summarize the ethical viewpoints of philosophy, research/engineering, and law on the use of AI through published documentation and literature. We expand on the primary ethical challenges researchers and engineers face in development, deployment, and use in machine learning, and create recommendations for ethical deployment. The authors of this manuscript argue that ethical use of AI from design through operations is more appropriate and achievable than developing AI that has "ethics." This section explores some of the prominent AI ethics historical examples and

A. Philosophical Machine Learning Ethics

Ethics in philosophy concerns the morality of decisions, of right and wrong, good and evil. For the last decade, philosophers have become concerned with the rise of AI; in particular, how to apply concepts of right and wrong to artificial intelligence. In general, philosophers tackling ethics subscribe to one of the three frameworks [57], shown in Table 4. These philosophical considerations help build the foundation of trust.

Table 4 Types of Ethics in Philosophy [57]

Ethical Framework	Description
Deontology	Right and wrong are set by a predetermined set of rules, and ethics works to codify and set those sets of rules.
Virtue Ethics	A person or groups character, in particular if they have good traits or characteristics, affect their morality and decisions.
Consequentialism	The way a person or group forms their ethical belief is due to feedback from society or others on if an action was right or wrong.

1. *The Golem and Trust: Believing Humans Can Anticipate all Scenarios is Moral Hubris*

The first story of AIs when it comes to ethics is arguably the 16th century Jewish myth, "The Golem of Prague" [88]. The most common version is that a golem was created to defend the city of Prague, but the golem is sealed for various reasons. The most relevant reason to AI is that the golem misinterprets its programming and performs an action which is unexpected [89]. This breaking of trust between the golem and its users is supposed to be a moral of hubris [90], but also demonstrates how ethical behavior factors into trust.

2. *The 21st Century Trolley Problem: Not all Cultures Share the Same Ethics*

The trolley problem is a well known moral psychology question which was first introduced in 1967 asking the reader to suppose a man "is the driver of a runaway tram which he can only steer from one narrow track onto another; five men are working on one track and one man on the other; anyone on the track he enters is bound to be killed" [91]. The trolley problem was later formalized and expanded on with several variations [92]. While initially just a starter for ethics debates, modern scientists use data driven psychology in order to understand the trolley dilemma, specifically by polling large groups of people on the situation in order to break down morality by location, gender, income gap, ethnicity, and other factors [93]. MIT has put together their own version, asking users what kind of decision an autonomous car should make given varying factors [94]. MIT was able to break down how the public would see a autonomous car crash by ethnicity, gender, religion, and affiliation in their paper *The MIT Moral Machine* [94]. MIT and other ethics researchers have shown that different societal groups are likely to have different ethical interpretations.

3. *Asimov and the Three Laws of Robotics*

The most well known ethical boundary in robotics and automation was created by Issac Asimov in 1942, known as the three laws of robotics which feature primarily in two books, *I, Robot* [83] and *Foundations* [95]. In his stories, Asimov proposes the following three laws be programmed as the foundational starting seed for robots to behave ethically [95]. These rules are:

- 1) A robot may not injure a human being or, through inaction, allow a human being to come to harm.
- 2) A robot must obey the orders given it by human beings, except where such orders would conflict with the First Law.

3) A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

Much of *I, Robot* [83], explores how modifying rule strength, or reordering rules leads to unintended consequences. The three laws are just as much of a societal expectation as they are a programmed rule, as Asimov mentions in *Lost Little Robots* [83] and *The Caves of Steel* [83]. He explores this in greater detail in *Runaround* [83], and then how those rules are manipulated by differing groups with different ethics in *The Naked Sun* [83][95]. Societies, even when they follow the same ethical guidelines, will have different interpretations [95] of ethics and use those interpretations in unforeseen methods[83].

4. ML Ethics in Philosophy Summary

The Golem, the MIT Moral Machine, and Asimov’s writings demonstrate a modern philosophical approach to ethics autonomous decision making agents. What one individual considers a rational or ethical decision can be up for interpretation, but as a whole societies and organizations share differences in what they define as moral and their ethical boundaries. Therefore, any algorithm attempting to resolve ethical dilemmas should share the ethical boundaries of its users. Designing for the ethics of users than designers may be considered controversial in cases where algorithms are applied outside of their country of origin, but the alternative is a lack of trust between the user and the algorithm. If practical, users should be polled during the development of algorithms to factor in their ethical considerations before wide scale deployment.

B. ML Ethics in Research and Engineering

While philosophy considers the ethical and moral boundaries of AI, ethics in research and engineering place more emphasis on ethical usage. Most engineering literature define ethics as rules of right and wrong, following the philosophical principle of Deontology [96]. In this subsection we introduce the challenge of ethical design and usage of AI systems.

1. Ethical Behavior in ML

Designing ethical behavior into algorithms during development and deployment is an active area of research. The primary focus as of publishing has been on three subjects: 1) whether algorithms can be programmed with ethical behavior, 2) how to remove bias in algorithm data, and 3) how organizational ethics, particularly around privacy, affect the Manned-AI Team decisions.

Ethics in Human-AI Teaming [57] suggests that ML algorithms can and should be programmed with an ethics decision algorithm. However, others argue that since it is not really possible to code those principles to ML [97]. David M. Tate of the Institute for Defense Analysis [56] also postulates that ethics is more about trust and transparency. In all cases, it matters more that there is trust between the human and the ML to not violate the human’s ethical boundaries[98].

Table 5 Types of Bias in AI [99]

Type of Bias	Description
Systemic Bias	Systemic bias comes from information sourced from history, society, or institutions. This mostly manifests in the data that arrives from these sources and the bias this data was collected with.
Human Bias	Human bias comes from the user and the biases the user has towards their interaction, as the bias of the user or organization will manifest in the AI algorithm.
Statistical Bias	Statistical bias comes from the development of the AI, particularly in the use of the system, the sampling and filtering of the algorithm, and the AI model selection.

Another important problem in ethical design of ML models is removing bias. The National Institute for Standards and Technology (NIST) has put together a table of possible bias, [99] shown in Table 5. When not managed, bias can be amplified into bad behavior in ML [100] [101]. To mitigate bad behavior, one approach is to use feedback generated by the user and stakeholders to evaluate ethical behavior [102] [101], while another approach [57] is to build in a series of

ethical safety checks during the training process, as demonstrated in Figure 4.

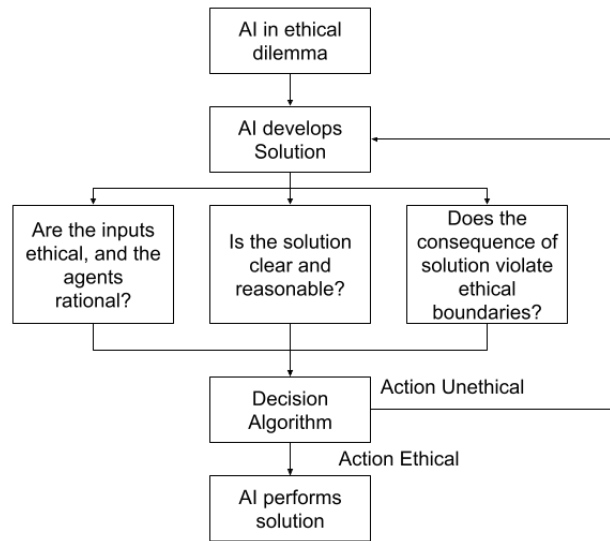


Fig. 4 Flow chart of modified of Agent, Deed, Consequence model

2. Responsible Use of ML

As ML and AI become more prolific, the debate over using AI has begun to shift away from whether AI can be ethical in general and into whether AI is used responsibly in particular applications. For example, there are ethical concerns surrounding the the development and deployment of Chatbots [102–104] Another controversy has arose at Alphabet Inc. when an engineer claimed their model, LaMDA, is sentient [105]. Unfortunately, examples of chatbots behaving unethically have become more prolific due to bias [106], or unethical applications [107]. Facial recognition is another controversial application considering it’s potential use cases [108] and biases [109]. The primary concern, however, is focused on privacy and ethics of machines in defense and combat roles [110]. Responsible design of robots in defense focuses primarily on the level of human input in the system (on, in, or out-of-the-loop, as discussed in Section II) when the ML is in an ethical dilemma leading to international recommendations for human in-the-loop systems [111]. The U.S. Department of defense has defined tactical autonomy as a system acting with delegated and bounded authority of humans [112].

C. Differences in Ethical Perceptions Across Cultures

Ethical use of AI is not universally defined, and two particularly different approaches depend on Eastern (in particular, China) vs. Western (United States and European Union) philosophy. While there are many differences, this section will discuss differences in cultural philosophies, attitudes towards AI, trust in institutions and concerns over surveillance and privacy.

1. Trust in Institutions using Artificial Intelligence

Another large cultural difference is trust in institutions and their ability to ethically use AI. For example, according to a 2018 Pew Research study the percent of U.S. adults that believe many institutions behave unethically all or most of the time is high (e.g. members of congress (81%), journalists (66%), leaders of technology companies (77%), religious leaders (69%), police officers (61%), military leaders (50%), and local elected officials (66%)) [113]. There are also significant differences in Eastern and Western countries attitudes towards AI. According to a 2020 Pew Research Study, the public of Eastern countries have a more favorable view of AI (e.g. Singapore (72%), South Korea (69%), India (67%), Taiwan (66%) and Japan (65%)), while Western countries are more split with lower favorable views (e.g. U.S. (47%), UK (46%), Germany (47%), and France (37%)) [114]. American attitudes towards AI, in particular autonomous weapons, are influenced by the films that depict armed AI systems [115]. Americans (82%) also believe that robots and/or AI should be carefully managed (similar to EU survey results), and that the highest importance issues were [116]:

- preventing AI-assisted surveillance from violating privacy and civil liberties,
- preventing AI from being used to spread fake and harmful content online,
- preventing AI cyberattacks against governments, companies, organizations, and individuals, and
- protecting data privacy.

2. Chinese AI Investment Strategy

When Google Deepmind’s AlphaGo defeated a professional Go player in 2016 [117], it was watched live by more than 280 million people in China and described by some as a "Sputnik moment," triggering new national strategies around AI, while the event was not nearly as significant to Western cultures. [118]. The 2017 “New Generation Artificial Intelligence Development Plan” Strategy [119] aimed to make China the AI leader and with a trillion yuan industry by 2030, has an incentives list of technologies, described broad usage in sectors ranging from defense to social welfare, and became a driving force in ethical norms for AI in China [118]. Concern over "moral decline" in China outweighs the concerns for privacy culturally, providing the impetus to use AI for moral governance through the Social Credit System, which regulates citizen behavior with regard to facets such as finances, taxes, academic dishonesty, and behavior on public transit [118]. Chinese companies have cooperated more on creating and applying surveillance technology due to the passing of the National Intelligence Law of the People’s Republic [120], which broadly mandates all Chinese companies, citizens, and companies with assets in China to collaborate with a Chinese intelligence agency [121].

3. Westerns Concerns over Military use of AI

It’s worth noting that western culture is concerned with the collaboration between defense and civilian technology. An example of this concern in the US is the public and corporate reaction to the DoD when it asked Google to work on Project Maven in 2017. In order to analyze the massive amount of video and other data generated from the growth in unmanned aerial vehicles [122], US Special Operations Command created a special joint task force that opened to industry to help solve the problem. Teams within Google, Microsoft, and Amazon Web Services were contracted to support the project [123]. However, major protests occurred within each company over worker concerns that their work may contribute to automated warfare [124]. Protests drove Google to withdraw from the contract [125]. Germany and other EU nations faced similar worries over military adoption of AI platforms. In 2020 a German parliamentary spokesperson resigned alongside protests against the adoption of armed UAVs in the Bundeswehr [126]. Debate over whether the German military should use armed autonomous vehicles polarize national political debate in the EU, and show the level of public concerns over military use of autonomous and AI platforms[127].

D. AI Ethics in Application

While the concerns of philosophy and ethical design/use are to be considered, an organization’s ethical framework and frameworks from law, affect the ethical use case and responsibility of operating ML agents in application. In international law, the International Committee of the Red Cross has taken the opinion that autonomous platforms follow the Martens Clause of the 1907 Hague Convention [111], proposing that the operator is responsible for an AI’s unethical behavior. The United States Department of Defense’s Innovation Board released a series of 5 guidelines of ethical principles that were officially adopted in 2020 [128] and align with civilian organizations, as shown in Table 6.

VI. Conclusion

This publication serves as a primer on considerations for safe, trusted, and ethical use of AI in aerospace control systems. While these concepts are closely related, and presence of each strengthens the others, there are some nuanced differences in considerations. A brief introduction was provided on AI techniques such as reinforcement learning, with examples of how it is being applied in aerospace research. There was also discussion of what it means for humans in a human-AI team to be in, on, or out-of-the-loop. Next, specific considerations and references for safety, trust and ethics were presented. For safety, existing aerospace standards, design practices, relationship to risk, and examples of lessons learned across domains is presented. For trust, a discussion of the dynamic nature of trust, need for AI transparency, variability in automation trust between humans for a variety of reasons, the risks associated with anthropomorphizing AI, and trust built between AI stakeholders. For ethics, different ethical frameworks and famous examples are introduced, as well as types of bias, methods of analyzing consequence, cultural differences, and ethical principles. AI and its applications to domains such as aerospace are evolving rapidly. Public access to generative AI tools like ChatGPT, DALL-E, Midjourney, Stable Diffusion, Google’s Bard, and others in the last couple of years has led to a corresponding

Table 6 DoD and NASA Ethical Principles [128, 129]

Department of Defense Innovation Board	NASA Guidelines
<i>Responsible:</i> DoD personnel will exercise appropriate levels of judgment and care, while remaining responsible for the development, deployment, and use of AI capabilities.	<i>Accountable:</i> Organizations and individuals must be accountable for the systems they create, and organizations must implement AI governance structures to provide oversight. AI developers should consider potential misuse or misinterpretation of AI-derived results (intentional or otherwise) and take steps to mitigate negative impact.
<i>Equitable:</i> The Department will take deliberate steps to minimize unintended bias in AI capabilities.	<i>Fair:</i> AI systems must include considerations regarding how to treat people, including refining solutions to mitigate discrimination and bias, preventing covert manipulation, and supporting diversity and inclusion.
<i>Traceable:</i> The Department’s AI capabilities will be developed and deployed such that relevant personnel possess an appropriate understanding of the technology, development processes, and operational methods applicable to AI capabilities, including with transparent and auditable methodologies, data sources, and design procedure and documentation.	<i>Explainable and Transparent:</i> Solutions must clearly state if, when, and how an AI system is involved, and AI logic and decisions must be explainable. AI solutions must protect intellectual property and include risk management in their construction and use. AI systems must be documented.
<i>Governable:</i> The Department’s AI capabilities will have explicit, well-defined uses, and the safety, security, and effectiveness of such capabilities will be subject to testing and assurance within those defined uses across their entire life-cycles.	<i>Human-Centric:</i> AI systems must obey human legal systems and must provide benefits to society. At the current state of AI, humans must remain in charge, though future advancements may cause reconsideration of this requirement.
<i>Reliable:</i> The Department’s AI capabilities will have explicit, well-defined uses, and the safety, security, and effectiveness of such capabilities will be subject to testing and assurance within those defined uses across their entire life-cycles.	<i>Secure and Safe:</i> AI systems must respect privacy and do no harm. Humans must monitor and guide machine learning processes. AI system risk tradeoffs must be considered when determining benefit of use.
	<i>Scientifically and Technically Robust:</i> AI systems must adhere to the scientific method NASA applies to all problems, be informed by scientific theory and data, robustly tested in implementation, well-documented, and peer reviewed in the scientific community.

growth in public interest and discourse in the use of AI. Safe, ethical, and trusted use of AI in aerospace will likely also rapidly evolve as AI technologies continue to mature.

Acknowledgments

The views expressed are those of the authors and do not reflect the official guidance or position of the United States Government, the Department of Defense, or the United States Air Force.

References

- [1] Kasparov, G., “The Chess Master and the Computer,” *Chess Metaphors; Artificial Intelligence and the Human Mind*, Vol. 205, 2010. URL <http://web.mit.edu/6.034/wwwbob/kasparov-article.pdf>.
- [2] Oxford Languages, “Artificial Intelligence,” , 2023.
- [3] Oxford Languages, “Machine Learning,” , 2023.
- [4] Alpaydin, E., *Machine Learning*, MIT Press, 2016.

- [5] Sutton, R. S., and Barto, A. G., *Reinforcement Learning: An Introduction*, MIT press, 2018.
- [6] Julian, K. D., Kochenderfer, M. J., and Owen, M. P., “Deep Neural Network Compression for Aircraft Collision Avoidance Systems,” *Journal of Guidance, Control, and Dynamics*, Vol. 42, No. 3, 2019, pp. 598–608.
- [7] Ravaoli, U. J., Cunningham, J., McCarroll, J., Gangal, V., Dunlap, K., and Hobbs, K. L., “Safe reinforcement learning benchmark environments for aerospace control systems,” *2022 IEEE Aerospace Conference (AERO)*, IEEE, 2022, pp. 1–20.
- [8] Julian, K. D., and Kochenderfer, M. J., “Distributed Wildfire Surveillance with Autonomous Aircraft using Deep Reinforcement Learning,” *Journal of Guidance, Control, and Dynamics*, Vol. 42, No. 8, 2019, pp. 1768–1778.
- [9] Deniz, S., and Wang, Z., “A Multi-Agent Reinforcement Learning Approach to Traffic Control at Future Urban Air Mobility Intersections,” *AIAA SCITECH 2022 Forum*, 2022, pp. 1–12.
- [10] Crespo, A. M. F., Weigang, L., and de Barros, A. G., “Reinforcement Learning Agents to Tactical Air Traffic Flow Management,” *International Journal of Aviation Management*, Vol. 1, No. 3, 2012, pp. 145–161.
- [11] Tang, C., and Lai, Y.-C., “Deep Reinforcement Learning Automatic Landing Control of Fixed-Wing Aircraft Using Deep Deterministic Policy Gradient,” *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, IEEE, 2020, pp. 1–9.
- [12] Andrade, P., Silva, C., Ribeiro, B., and Santos, B. F., “Aircraft Maintenance Check Scheduling using Reinforcement Learning,” *Aerospace*, Vol. 8, No. 4, 2021.
- [13] Clarke, S. G., and Hwang, I., “Deep Reinforcement Learning Control for Aerobatic Maneuvering of Agile Fixed-wing Aircraft,” *AIAA Scitech 2020 Forum*, 2020, p. 0136.
- [14] Konatala, R., Van Kampen, E.-J., and Looye, G., “Reinforcement Learning based Online Adaptive Flight Control for the Cessna Citation II (PH-LAB) Aircraft,” *AIAA Scitech 2021 Forum*, 2021, p. 0883.
- [15] Waslander, S. L., Hoffmann, G. M., Jang, J. S., and Tomlin, C. J., “Multi-agent Quadrotor Testbed Control Design: Integral Sliding Mode vs. Reinforcement Learning,” *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2005, pp. 3712–3717.
- [16] Bøhn, E., Coates, E. M., Moe, S., and Johansen, T. A., “Deep Reinforcement Learning Attitude Control of Fixed-wing UAVs using Proximal Policy Optimization,” *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*, IEEE, 2019, pp. 523–533.
- [17] Tandale, M., Rong, J., and Valasek, J., “Preliminary Results of Adaptive Reinforcement Learning Control for Morphing Aircraft,” *AIAA Guidance, Navigation, and Control Conference*, 2004, p. 5358.
- [18] De Vries, P. S., and Van Kampen, E.-J., “Reinforcement Learning-based Control Allocation for the Innovative Control Effectors Aircraft,” *AIAA Scitech 2019 Forum*, 2019, pp. 1–22.
- [19] Woodbury, T. D., Dunn, C., and Valasek, J., “Autonomous Soaring using Reinforcement Learning for Trajectory Generation,” *52nd Aerospace Sciences Meeting*, 2014, p. 0990.
- [20] Balakrishna, P., Ganesan, R., and Sherry, L., “Accuracy of Reinforcement Learning Algorithms for Predicting Aircraft Taxi-out Times: A Case-study of Tampa Bay Departures,” *Transportation Research Part C: Emerging Technologies*, Vol. 18, No. 6, 2010, pp. 950–962.
- [21] Pope, A. P., Ide, J. S., Mićović, D., Diaz, H., Rosenbluth, D., Ritholtz, L., Twedt, J. C., Walker, T. T., Alcedo, K., and Javorsek, D., “Hierarchical Reinforcement Learning for Air-to-Air Combat,” *2021 International Conference on Unmanned Aircraft Systems (ICUAS)*, IEEE, 2021, pp. 275–284.
- [22] Gaudet, B., and Furfaro, R., “Robust Spacecraft Hovering Near Small Bodies in Environments with Unknown Dynamics Using Reinforcement Learning,” *AIAA/AAS Astrodynamics Specialist Conference*, 2012, p. 5072.
- [23] Harris, A. T., and Schaub, H., “Spacecraft Command and Control with Safety Guarantees using Shielded Deep Reinforcement Learning,” *AIAA Scitech Forum*, 2020, p. 0386.
- [24] Scorsoglio, A., Furfaro, R., Linares, R., Massari, M., et al., “Actor-critic Reinforcement Learning Approach to Relative Motion Guidance in Near-Rectilinear Orbit,” *Advances in the Astronautical Sciences*, 2019, pp. 1–20.
- [25] Broida, J., and Linares, R., “Spacecraft Rendezvous Guidance in Cluttered Environments via Reinforcement Learning,” *29th AAS/AIAA Space Flight Mechanics Meeting*, 2019, pp. 1–15.

- [26] Dunlap, K., Mote, M., Delsing, K., and Hobbs, K. L., "Run time assured reinforcement learning for safe satellite docking," *Journal of Aerospace Information Systems*, Vol. 20, No. 1, 2023, pp. 25–36.
- [27] Oestreich, C. E., Linares, R., and Gondhalekar, R., "Autonomous Six-Degree-of-Freedom Spacecraft Docking Maneuvers via Reinforcement Learning," *2020 AAS/AIAA Astrodynamics Specialist Virtual Lake Tahoe*, 2020.
- [28] Hovell, K., and Ulrich, S., "On Deep Reinforcement Learning for Spacecraft Guidance," *AIAA Scitech Forum*, 2020, p. 1600.
- [29] Chan, D. M., and Agha-mohammadi, A.-a., "Autonomous Imaging and Mapping of Small Bodies Using Deep Reinforcement Learning," *2019 IEEE Aerospace Conference*, IEEE, 2019, pp. 1–12.
- [30] Gaudet, B., Linares, R., and Furfaro, R., "Deep Reinforcement Learning for Six Degree-of-freedom Planetary Landing," *Advances in Space Research*, Vol. 65, No. 7, 2020, pp. 1723–1741.
- [31] Gaudet, B., Linares, R., and Furfaro, R., "Adaptive Guidance and Integrated Navigation with Reinforcement Meta-learning," *Acta Astronautica*, Vol. 169, 2020, pp. 180–190.
- [32] Rushby, J., "Software Verification and System Assurance," *2009 Seventh IEEE International Conference on Software Engineering and Formal Methods*, IEEE, 2009, pp. 3–10.
- [33] Lee, J. D., and See, K. A., "Trust in automation: Designing for appropriate reliance," *Human Factors*, Vol. 46, No. 1, 2004, pp. 50–80.
- [34] Rudner, T. G., and Toner, H., "Key concepts in AI safety: an overview," , 2021. URL <https://cset.georgetown.edu/publication/key-concepts-in-ai-safety-an-overview/>.
- [35] ARP4761, "Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment," Guidance, SAE International, 1996.
- [36] ARP4754A, "Guidelines For Development Of Civil Aircraft and Systems," Guidance, SAE International, 2010.
- [37] RTCA DO-254, "Design Assurance Guidance for Airborne Electronic Hardware," Guidance, RTCA Inc., 2000.
- [38] RTCA DO-178C, "Software Considerations in Airborne Systems and Equipment Certification," Guidance, RTCA Inc., 2011.
- [39] RTCA DO-333, "Formal Methods Supplement to DO-178C and DO-278A," Guidance, RTCA Inc., 2011.
- [40] ASTM3269-17, "Standard Practice for Methods to Safely Bound Flight Behavior of Unmanned Aircraft Systems Containing Complex Functions," Standard, ASTM International, 2017.
- [41] ASTM3269-21, "Standard Practice for Methods to Safely Bound Flight Behavior of Unmanned Aircraft Systems Containing Complex Functions," Standard, ASTM International, 2021.
- [42] MIL-HDBK-516C, "MIL-HDBK-516C: Airworthiness Certification Criteria," Guidance, Department of Defense, 12 2014.
- [43] MIL-STD-882E, "MIL-STD-882E: System Safety," Standard Practice, Department of Defense, May 2012.
- [44] Federal Aviation Administration (FAA), "Standard Airworthiness Certification Regulations: Title 14, Code of Federal Regulations," Website, 10 2017. https://www.faa.gov/aircraft/air_cert/airworthiness_certification/std_awcert/std_awcert_regs/regs/.
- [45] CoDANNI, "Concepts of Design Assurance for Neural Networks (CoDANN)," Technical Report, European Union Aviation Safety Agency, 2020.
- [46] CoDANNI, "Concepts of Design Assurance for Neural Networks (CoDANN) II," Technical Report, European Union Aviation Safety Agency, May 2021.
- [47] Burns, A., Harper, D., Barfield, A. F., Whitcomb, S., and Jurusik, B., "Auto GCAS for analog flight control system," *2011 IEEE/AIAA 30th Digital Avionics Systems Conference*, IEEE, 2011, pp. 8C5–1.
- [48] Turner, R., Lehmann, R., Wadley, J., Kidd, D., Swihart, D., Bier, J., and Hobbs, K., "Automatic Aircraft Collision Avoidance algorithm design for fighter aircraft," *Asia-Pacific International Symposium on Aerospace Technology*, 2012.
- [49] Swihart, D., Barfield, A., Griffin, E., Lehmann, R., Whitcomb, S., Skoog, M., Flynn, B., and Prosser, K., "Design, integration and flight test of an autonomous ground collision avoidance system," *Gyroscope and Navigation*, Vol. 2, No. 2, 2011, pp. 84–91.

- [50] Swihart, D. E., Barfield, A. F., Griffin, E. M., Lehmann, R. C., Whitcomb, S. C., Flynn, B., Skoog, M. A., and Prosser, K. E., "Automatic Ground Collision Avoidance System Design, Integration, & Flight Test," *IEEE Aerospace and Electronic Systems Magazine*, Vol. 26, No. 5, 2011, pp. 4–11.
- [51] Lyons, J. B., Ho, N. T., Ferguson, W. E., Sadler, G. G., Cals, S. D., Richardson, C. E., and Wilkins, M. A., "Trust of an automatic ground collision avoidance technology: A fighter pilot perspective," *Military Psychology*, Vol. 28, No. 4, 2016, pp. 271–277.
- [52] Wadley, J., Jones, S., Stoner, D., Griffin, E., Swihart, D., Hobbs, K., Burns, A., and Bier, J., "Development of an Automatic Aircraft Collision Avoidance System for Fighter Aircraft," *AIAA Infotech@ Aerospace Conference*, 2013, p. 4727.
- [53] Sorokowski, P., Skoog, M., Burrows, S., and Thomas, S., "Small UAV Automatic Ground Collision Avoidance System Design Considerations and Flight Test Results," 2015.
- [54] Hobbs, K., Mote, M., Abate, M., Coogan, S., and Feron, E., "Run time assurance for safety-critical systems: An introduction to safety filtering approaches for complex control systems," *Control Systems Magazine*, Vol. 43, No. 2, 2023, p. TBD.
- [55] Ho, N. T., Sadler, G. G., Hoffmann, L. C., Lyons, J. B., and Johnson, W. W., "Trust of a Military Automated System in an Operational Context," *Military Psychology*, Vol. 29, No. 6, 2017, pp. 524–541.
- [56] Tate, D. M., "Trust, Trustworthiness, and Assurance of AI and Autonomy," 2021. URL <https://www.ida.org/research-and-publications/publications/all/t/tr/trust-trustworthiness-and-assurance-of-ai-and-autonomy>.
- [57] Pflanzner, M., "Ethics in Human-AI Teaming: Principles and Perspectives," 2022. Unpublished Manuscript.
- [58] Hobbs, K. L., "Elicitation and Formal Specification of Run Time Assurance Requirements for Aerospace Collision Avoidance Systems," Ph.D. thesis, Georgia Institute of Technology, 2020.
- [59] S3001, "S3001: Guidelines for Risk Management," Standard Practice, NASA, October 2017.
- [60] "Highway Accident Report: Collision Between Vehicle Controlled by Developmental Automated Driving System and Pedestrian," Tech. rep., 2018. URL <https://www.nts.gov/investigations/accidentreports/reports/har1903.pdf>.
- [61] Parasuraman, R., "Complacency and Bias in Human Use of Automation: An Attentional Integration," 2010. URL <http://dx.doi.org/10.1177/0018720810376055>.
- [62] Lee, J. D., "Trust in automation: Designing for appropriate reliance," 2004. URL https://journals.sagepub.com/doi/10.1518/hfes.46.1.50_30392.
- [63] Barry, K., and Bartlett, J. S., "Tesla's Camera-Based Driver Monitoring Still Fails to Keep Driver Attention on the Road, CR Tests Show," 2022. URL <https://www.consumerreports.org/car-safety/tesla-driver-monitoring-fails-to-keep-driver-focus-on-road-a3964813328/>.
- [64] Kenny, J. G., "The Need for Change, the legacy of Three Mile Island," 1979. URL https://www.ans.org/file/6651/1/Kemeny_20Commission_20Report.pdf.
- [65] "Assumptions Used in the Safety Assessment Process and the Effects of Multiple Alerts and Indications on Pilot Performance," Tech. rep., 2019. URL <https://www.nts.gov/investigations/accidentreports/reports/asr1901.pdf>.
- [66] Madhavan, P., and Wiegmann, D. A., "Similarities and differences between human–human and human–automation trust: an integrative review," *Theoretical Issues in Ergonomics Science*, Vol. 8, No. 4, 2007, pp. 277–301.
- [67] Sorkin, R. D., "Why are People Turning off our Alarms?" *The Journal of the Acoustical Society of America*, Vol. 84, No. 3, 1988, pp. 1107–1108.
- [68] Chen, J. Y., Flemisch, F. O., Lyons, J. B., and Neerincx, M. A., "Guest editorial: Agent and system transparency," *IEEE Transactions on Human-Machine Systems*, Vol. 50, No. 3, 2020, pp. 189–193.
- [69] Mercado, J. E., Rupp, M. A., Chen, J. Y., Barnes, M. J., Barber, D., and Procci, K., "Intelligent agent transparency in human–agent teaming for Multi-UxV management," *Human factors*, Vol. 58, No. 3, 2016, pp. 401–415.
- [70] Marsh, S., and Dibben, M. R., "The role of trust in information science and technology," *Annual Review of Information Science and Technology (ARIST)*, Vol. 37, 2003, pp. 465–98.

- [71] Biro, D. P., Daly, M., and Gunsch, G., "The influence of task load and automation trust on deception detection," *Group Decision and Negotiation*, Vol. 13, No. 2, 2004, pp. 173–189.
- [72] Hoff, K. A., and Bashir, M., "Trust in automation: Integrating empirical evidence on factors that influence trust," *Human Factors*, Vol. 57, No. 3, 2015, pp. 407–434.
- [73] Huerta, E., Glandon, T., and Petrides, Y., "Framing, decision-aid systems, and culture: Exploring influences on fraud investigations," *International Journal of Accounting Information Systems*, Vol. 13, No. 4, 2012, pp. 316–333.
- [74] Ho, G., Wheatley, D., and Scialfa, C. T., "Age differences in trust and reliance of a medication management system," *Interacting with Computers*, Vol. 17, No. 6, 2005, pp. 690–710.
- [75] Sanchez, J., Fisk, A. D., and Rogers, W. A., "Reliability and age-related effects on trust and reliance of a decision support aid," *proceedings of the human factors and ergonomics society annual meeting*, Vol. 48, Sage Publications Sage CA: Los Angeles, CA, 2004, pp. 586–589.
- [76] Pak, R., Fink, N., Price, M., Bass, B., and Sturre, L., "Decision support aids with anthropomorphic characteristics influence trust and performance in younger and older adults," *Ergonomics*, Vol. 55, No. 9, 2012, pp. 1059–1072.
- [77] Lee, E.-J., "Flattery may get computers somewhere, sometimes: The moderating role of output modality, computer gender, and user gender," *International Journal of Human-Computer Studies*, Vol. 66, No. 11, 2008, pp. 789–800.
- [78] Nomura, T., Kanda, T., Suzuki, T., and Kato, K., "Prediction of human behavior in human–robot interaction using psychological scales for anxiety and negative attitudes toward robots," *IEEE transactions on robotics*, Vol. 24, No. 2, 2008, pp. 442–451.
- [79] Tung, F.-W., "Influence of gender and age on the attitudes of children towards humanoid robots," *International Conference on Human-Computer Interaction*, Springer, 2011, pp. 637–646.
- [80] Merritt, S. M., and Ilgen, D. R., "Not all trust is created equal: Dispositional and history-based trust in human-automation interactions," *Human factors*, Vol. 50, No. 2, 2008, pp. 194–210.
- [81] McBride, M., Carter, L., and Ntuen, C., "The impact of personality on nurses' bias towards automated decision aid acceptance," *International Journal of Information Systems and Change Management*, Vol. 6, No. 2, 2012, pp. 132–146.
- [82] Szalma, J. L., and Taylor, G. S., "Individual differences in response to automation: the five factor model of personality," *Journal of Experimental Psychology: Applied*, Vol. 17, No. 2, 2011, p. 71.
- [83] Asimov, I., *The Complete Robot*, 1982.
- [84] Sung, J.-Y., "'My Roomba Is Rambo': Intimate Home Appliances," 2007. URL <https://faculty.cc.gatech.edu/~beki/c35.pdf>.
- [85] Billings, D. R., "Human-Animal Trust as an Analog for Human-Robot Trust: A Review of Current Evidence," Tech. rep., 2012.
- [86] Lacher, A., Grabowski, R., and Cook, S., "Autonomy, trust, and transportation," *2014 AAAI Spring Symposium Series*, 2014.
- [87] Mandrake, L., Doran, G., Goel, A., Ono, H., Amini, R., Feather, M. S., Fesq, L., Slingerland, P., Perry, L., Bycroft, B., and Kaufman, J., "Space Applications of a Trusted AI Framework: Experiences and Lessons Learned," *2022 IEEE Aerospace Conference*, IEEE, 2022.
- [88] Auerbach, B., *Spinoza*, 1837.
- [89] Lewis, J., "The Golem of Prague: Man or Monster?" *Hohonu Academic Journal*, Vol. 205, 2019, pp. 23–27. URL <https://hilo.hawaii.edu/campuscenter/hohonu/volumes/documents/TheGolemofPragueManorMonster.pdf>.
- [90] Gelbin, C. S., "The Golem Returns: From German Romanticism Literature to Global Jewish Literature," 2011. URL <https://www.press.umich.edu/pdf/9780472117598-intro.pdf>.
- [91] Foot, P., "The problem of abortion and the doctrine of the double effect," *Oxford Review*, , No. 5, 1967. URL <https://philpapers.org/archive/footpo-2.pdf>.
- [92] Thomson, J. J., "The trolley problem," *Yale LJ*, Vol. 94, 1984, p. 1395.
- [93] Navarrete, C. D., "Virtual Morality: Emotion and Action in a Simulated Three-Dimensional 'Trolley Problem,'" *Emotion*, Vol. 12, 2012, pp. 364–370. URL <https://doi.apa.org/doiLanding?doi=10.1037/a0025561>.

- [94] Awad, E., “The Moral Machine experiment,” *Nature*, Vol. 563, 2018, pp. 59–64. URL <https://doi.org/10.1038/s41586-018-0637-6>.
- [95] Asimov, I., *Foundation*, 1951.
- [96] “ASME CODE OF ETHICS OF ENGINEERS,” Tech. rep., 1976. URL <https://www.asme.org/wwwasmeorg/media/resourcefiles/aboutasme/get%20involved/advocacy/policy-publications/p-15-7-ethics.pdf>.
- [97] Chaudhury, D., “Why The Ethics Of AI Are Complicated,” 2021. URL <https://www.forbes.com/sites/forbestechcouncil/2021/08/30/why-the-ethics-of-ai-are-complicated/?sh=3d98f49b43a7>.
- [98] Ho, N. T., “Trust of a Military Automated System in an Operational Context,” Vol. 29, 2017, pp. 524–541. URL <https://doi.org/10.1037/mil0000189>.
- [99] Schwartz, R., “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence,” Tech. rep., 2022. URL <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf>.
- [100] Allen, G., “Can AI be ethical?” 2022. URL <https://www.youtube.com/watch?v=6sTGXUQysaQ>.
- [101] O’Neil, C., “Near-Term Artificial Intelligence and the Ethical Matrix,” Vol. 8, 2020, pp. 237–269. URL <http://www.eliassi.org/Near-Term%20AI%20and%20the%20Ethical%20Matrix.pdf>.
- [102] Google, “LaMDA: Language Models for Dialog Applications,” Tech. rep., 2022. URL <https://arxiv.org/pdf/2201.08239.pdf>.
- [103] Press Association, “Computer simulating 13-year-old boy becomes first to pass Turing test,” 2014. URL <https://www.theguardian.com/technology/2014/jun/08/super-computer-simulates-13-year-old-boy-passes-turing-test>.
- [104] Press Association, “Eugene the Turing test-beating ‘human computer’ – in ‘his’ own words,” 2014. URL <https://www.theguardian.com/technology/2014/jun/09/eugene-person-human-computer-robot-chat-turing-test>.
- [105] Lemoine, B., “Is LaMDA Sentient? — an Interview,” 2022. URL <https://cajundiscordian.medium.com/is-lambda-sentient-an-interview-ea64d916d917>.
- [106] Neff, G., “Talking to Bots: Symbiotic Agency and the Case of Tay,” 2016. URL https://www.researchgate.net/publication/309101946_Talking_to_Bots_Symbiotic_Agency_and_the_Case_of_Tay.
- [107] Gault, M., “AI Trained on 4Chan Becomes ‘Hate Speech Machine,’” 2022. URL <https://www.vice.com/en/article/7k8zwx/ai-trained-on-4chan-becomes-hate-speech-machine>.
- [108] Oliver, J., “Facial Recognition: Last Week Tonight with John Oliver (HBO),” 2020. URL <https://www.youtube.com/watch?v=jZjmLJPJgug>.
- [109] Buolamwini, J., “Gender Shades,” Tech. rep., 2017. URL <https://www.media.mit.edu/publications/full-gender-shades-thesis-17/>.
- [110] PROJECT, P., “STATEMENT OF RESIGNING AXON AI ETHICS BOARD MEMBERS,” 2022. URL <https://www.policingproject.org/statement-of-resigning-axon-ai-ethics-board-members>.
- [111] “Ethics and autonomous weapon systems: An ethical basis for human control?” 2018. URL <https://blogs.icrc.org/law-and-policy/2017/11/14/ethics-source-law-martens-clause-autonomous-weapons/>.
- [112] Zacharias, G., “Autonomous Horizons, the Way Forward ,” , Office of the Chief Scientist, USAF, March 2019.
- [113] Gecewicz, C., and Rainie, L., *Why Americans Don’t Fully Trust Many who Hold Positions of Power and Responsibility: Members of Congress and Technology Leaders are Rated Lower in Empathy, Transparency and Ethics; Public Gives Higher Scores to Military Leaders, Public School Principals and Police Officers*, Pew Research Center., 2019. URL <https://www.pewresearch.org/politics/2019/09/19/why-americans-dont-fully-trust-many-who-hold-positions-of-power-and-responsibility/>.
- [114] Johnson, C., and Tyson, A., “People globally offer mixed views of the impact of artificial intelligence, job automation on society,” 2020. URL <https://www.pewresearch.org/fact-tank/2020/12/15/people-globally-offer-mixed-views-of-the-impact-of-artificial-intelligence-job-automation-on-society/>.

- [115] Young, K. L., and Carpenter, C., “Does science fiction affect political fact? Yes and no: A survey experiment on “Killer Robots”,” *International Studies Quarterly*, Vol. 62, No. 3, 2018, pp. 562–576.
- [116] Zhang, B., and Dafoe, A., “Artificial intelligence: American attitudes and trends,” *Center for the Governance of AI, Future of Humanity Institute, University of Oxford*, 2019. URL https://isps.yale.edu/sites/default/files/files/Zhang_us_public_opinion_report_jan_2019.pdf.
- [117] Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al., “Mastering the game of go without human knowledge,” *nature*, Vol. 550, No. 7676, 2017, pp. 354–359.
- [118] Roberts, H., Cowls, J., Morley, J., Taddeo, M., Wang, V., and Floridi, L., “The Chinese approach to artificial intelligence: an analysis of policy, ethics, and regulation,” *AI & society*, Vol. 36, No. 1, 2021, pp. 59–77.
- [119] “China’s ‘New Generation Artificial Intelligence Development Plan’ Translated by Stanford University,” , 2017. URL <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.
- [120] “National Intelligence Law of the People’s Republic Translated by Brown University,” , 2017. URL https://cs.brown.edu/courses/csci1800/sources/2017_PRC_NationalIntelligenceLaw.pdf.
- [121] Government of Canada, *Rethinking Security: China and the Age of Strategic Rivalry*, Government of Canada, 2018. URL <https://www.canada.ca/en/security-intelligence-service/corporate/publications/china-and-the-age-of-strategic-rivalry.html>.
- [122] Shultz, R., and Clarke, G. R. D., “BIG DATA AT WAR: SPECIAL OPERATIONS FORCES, PROJECT MAVEN, AND TWENTY-FIRST-CENTURY WARFARE,” 2020. URL <https://mwi.usma.edu/big-data-at-war-special-operations-forces-project-maven-and-twenty-first-century-warfare/>.
- [123] Brewster, T., “Project Maven: Amazon And Microsoft Scored 50 Million In Pentagon Surveillance Contracts After Google Quit,” 2021. URL <https://www.forbes.com/sites/thomasbrewster/2021/09/08/project-maven-amazon-and-microsoft-get-50-million-in-pentagon-drone-surveillance-contracts-after-google/?sh=6e34580c6f1e>.
- [124] Schneider, A., and Sydell, L., “Microsoft Workers Protest Army Contract With Tech ‘Designed To Help People Kill’,” 2019. URL <https://www.npr.org/2019/02/22/697110641/microsoft-workers-protest-army-contract-with-tech-designed-to-help-people-kill>.
- [125] Conger, K., “Google Plans Not to Renew Its Contract for Project Maven, a Controversial Pentagon Drone AI Imaging Program,” 2018. URL https://gizmodo.com/google-plans-not-to-renew-its-contract-for-project-mave-1826488620?rev=1527878336532&utm_campaign=socialflow_gizmodo_twitter&utm_source=gizmodo_twitter&utm_medium=socialflow.
- [126] Gady, F.-S., “Germany to operate armed UAVs: too little too late?” Tech. rep., 2021. URL <https://www.iiss.org/online-analysis/online-analysis//2021/11/germany-to-operate-armed-uavs-too-little-too-late>.
- [127] Wachs, L., “The Latecomer: Germany’s Debate on Armed Drones,” Tech. rep., 2020. URL <https://www.rusi.org/explore-our-research/publications/commentary/latecomer-germanys-debate-armed-drones>.
- [128] “DOD Adopts Ethical Principles for Artificial Intelligence,” 2020. URL <https://www.defense.gov/News/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/>.
- [129] McLarney, E., “NASA Framework for the Ethical Use of Artificial Intelligence (AI),” Tech. rep., 2021. URL <https://ntrs.nasa.gov/api/citations/20210012886/downloads/NASA-TM-20210012886.pdf>.