

AUTONOMOUS HORIZONS

System Autonomy in the Air Force – A Path to the Future

Volume I: Human-Autonomy Teaming



United States Air Force
Office of the Chief Scientist

AF/ST TR 15-01
June 2015

The views expressed in this document are those of the Air Force Office of the Chief Scientist and do not necessarily reflect the official policy or position of the Air Force, the Department of Defense, or the U.S. Government. The Office of the Chief Scientist serves an advisory role within the Air Force and does not make decisions regarding official Air Force policy or positions.



Office of the
USAF Chief Scientist

DEPARTMENT OF THE AIR FORCE
HEADQUARTERS OF THE AIR FORCE
WASHINGTON DC

1 June 2015

America's Air Force: A Call to the Future provides a compelling 30-year vision for ensuring that the U.S. Air Force can continue to provide responsive and effective *Global Vigilance—Global Reach—Global Power*. Creating innovative game-changing technologies that are agile and able to amplify many of the enduring attributes of airpower – speed, range, flexibility, and precision – form a critical portion of this vision.

Unmanned systems and autonomous software offer significant potential advantages for meeting the challenges of a newly forming adversarial environment. Speed of light cyber-attacks, anti-access/area-denial (A2SD) actions that keep our forces operating at a distance, and potential attacks on our space-based assets all require innovative solutions for maintaining mission-effective air, space and cyber operations in the face of these new challenges.

Autonomous Horizons depicts a path to the future for system autonomy in the Air Force. It describes an evolutionary progression that obtains the best benefits of autonomous software working synergistically with the innovation of empowered airmen. This vision is both obtainable and sustainable – it leaves the authority and responsibility for warfare in the hands of airmen while creating tools that enhance their situation awareness and decision-making, speed effective actions, and bring needed extensions to their capabilities. Rather than attempting to design the airman out of the equation, the Air Force embraces the agility, intelligence and innovation that airmen provide, along with the advanced capabilities of autonomy, to create effective teams in which activities can be accomplished smoothly, simply and seamlessly.

In order to thrive in the future, we must pursue strategic agility – in our people and in our technology. New paradigms of operation are needed to maintain the strategic advantage for our Airmen, our Air Force, and our Nation. *Autonomous Horizons* provides a pathway to the future for Air Force systems to ensure our ability to fly, fight, and win in air, space, and cyberspace.

A handwritten signature in black ink, reading "Mica R. Endsley".

Mica R. Endsley
Chief Scientist
United States Air Force

Executive Summary

Autonomous systems provide a considerable opportunity to enhance future Air Force operations by potentially reducing unnecessary manning costs, increasing the range of operations, enhancing capabilities, providing new approaches to air power, reducing the time required for critical operations, and providing increased levels of operational reliability, persistence and resilience. Increased levels of autonomy can be brought to bear to enhance operations in both manned and unmanned aircraft, and in operations in space, cyber, command and control, intelligence, surveillance and reconnaissance, readiness, and sustainment across the Air Force.

Autonomous Horizons serves to provide direction and guidance on the opportunities and challenges for the development of autonomous systems for Air Force operations. This goal is to be addressed in three volumes. Volume I describes a vision for autonomous systems that will work synergistically with our airmen as a part of an effective human-autonomy team where functions and situation awareness flow smoothly, simply, and seamlessly between them. Volume II details the many technical issues involved in creating machine intelligence that can deal effectively with the challenges of uncertainty and variability in operational environments. Volume III will address key issues associated with cyber security and reliability, communications links, and command and control systems to support autonomous vehicles, as well as issues in the development of autonomous systems, including verification and validation of autonomy software and hardware.

In this first volume, a summary of the challenges of automation and autonomy for airman interaction are presented, based on some four decades of experience and research on this issue. These include (1) difficulties in creating autonomy software that is robust enough to function without human intervention and oversight, (2) the lowering of human situation awareness that occurs when using automation leading to out-of-the-loop performance decrements, (3) increases in cognitive workload required to interact with the greater complexity associated with automation, (4) increased time to make decisions when decision aids are provided, often without the desired increase in decision accuracy, and (5) challenges with developing a level of trust that is appropriately calibrated to the reliability and functionality of the system in various circumstances. Given that it is unlikely that autonomy in the foreseeable future will work perfectly for all functions and operations, and that airman interaction with autonomy will continue to be needed at some level, each of these factors works to create the need for a new approach to the design of autonomous systems that will allow them to serve as an effective teammate with the airmen who depend on them to do their jobs.

In this vision of the future, autonomous systems will be designed to serve as a part of a collaborative team with airmen. Flexible autonomy will allow the control of tasks, functions, sub-systems, and even entire vehicles to pass back and forth over time between the airman and the autonomous system, as needed to succeed under changing circumstances. Many functions will be supported at varying levels of autonomy, from fully manual, to recommendations for decision aiding, to human-on-the-loop supervisory control of an autonomous system, to one that operates fully autonomously with no human intervention at all. The airman will be able to make informed choices about where and when to invoke autonomy based on considerations of trust,

the ability to verify its operations, the level of risk and risk mitigation available for a particular operation, the operational need for the autonomy, and the degree to which the system supports the needed partnership with the airman. In certain limited cases the system may allow the autonomy to take over automatically from the airman, when timelines are very short for example, or when loss of lives are imminent. However, human decision making for the exercise of force with weapon systems is a fundamental requirement, in keeping with Department of Defense directives.

The development of autonomy that provides sufficient robustness, span of control, ease of interaction, and automation transparency is critical to achieving this vision. In addition, a high level of shared situation awareness between the airman and the autonomy will be critical. Shared situation awareness is needed to ensure that the autonomy and the airman are able to align their goals, track function allocation and re-allocation over time, communicate decisions and courses of action, and align their respective tasks to achieve coordinated actions. Critical situation awareness requirements that communicate not just status information, but also comprehension and projections associated with the situation (the higher levels of situation awareness), must be built into future two-way communications between the airman and the autonomy.

Developing future autonomous systems that achieve this vision will require addressing many key technical challenges. Future autonomy will need to be able to more effectively process sensor data and airman inputs to create its own internal situation model to direct its decision making. By drawing on research on human situation awareness, a cognitively inspired architecture for autonomy situation models can provide significant gains for creating effective and robust autonomous systems.

As many approaches to autonomy are based on adaptive technologies and learning techniques, many new challenges will also be created, including new problems with supporting understandability of the autonomous system, the need to manage standardization among potentially varied systems that have learned different lessons, and successful methods for verification and validation of the autonomy. In addition, methods for creating resilience to cyber-attacks must be carefully considered throughout system development.

Many Air Force systems will experience an evolution towards increasing levels of autonomy over the next several decades. These advances will only be successful in achieving their goals of increased range and speed of operations, increased mission capabilities, increased reliability, persistence and resilience, or reduced manning loads if they take careful consideration of the need for effective airman-autonomy teaming. Past paradigms that created brittle automation, with limited capabilities and limited consideration of human operators, will be replaced by an explicit focus on synergistic airman-autonomy teams. This new paradigm will directly support high levels of shared situation awareness between the airman and the autonomy, creating situationally relevant informed trust, ease of interaction and control, and the manageable workload levels needed for mission success. By focusing on airman-autonomy teaming, the Air Force will create successful systems that get the best benefits of autonomous software along with the innovation of empowered airmen.

Table of Contents

1.0 The Promise of Autonomy in Air Force Operations.....	1
1.1 Air.....	1
1.2 Space.....	2
1.3 Cyber.....	2
1.4 Command, Control and ISR.....	2
1.5 Readiness and Sustainment.....	3
2.0 Definitions.....	3
2.1 Automation.....	3
2.2 Autonomy.....	3
2.3 Remotely Controlled Vehicles.....	4
3.0 Challenges for Use of Autonomous Systems.....	4
3.1 System Capabilities.....	5
3.2 Situation Awareness and Out-of-the Loop Performance Problems.....	6
3.3 Optimal Workload Levels.....	7
3.4 Integrating Human & Autonomy Decision Making.....	7
3.5 Informed, Situational Trust in Autonomy.....	8
4.0 Towards Symbiotic Human-Autonomy Systems.....	9
4.1 Flexible Autonomy.....	9
4.1.1 Levels of Autonomy.....	9
4.1.2 Dynamic Autonomy Selection.....	11
4.1.3 Limitations on Autonomy for Weapons Systems.....	12
4.1.4 Human–Autonomy Interaction.....	13
4.2 Building Shared Situation Awareness to Support Airman-Autonomy Teams..	14
4.2.1 Human-Interaction Guidelines.....	14
4.2.2 Shared Situation Awareness.....	15
4.2.3 Supporting Unmanned Operations.....	18
4.2.4 Trustworthy Autonomy.....	19
5.0 Challenges for the Development of Autonomous Systems.....	20
5.1 Situation Models for Autonomy.....	21
5.2 Learning Systems.....	22
5.3 Verification & Validation.....	23
5.4 Cyber.....	23
6.0 Conclusions.....	24
7.0 References.....	24

1.0 The Promise of Autonomy in Air Force Operations

In laying out a vision for the next 30 years, the U.S. Air Force strategy provides an emphasis on the development of system autonomy as a means of achieving strategic advantage in future operations [1]. This includes an increased use of automation software, as well as more advanced algorithms that enable systems to act autonomously or “*react to their environment and perform more situational-dependent tasks as well as synchronized and integrated functions with other autonomous systems*”. Increased levels of system autonomy promise to:

- Reduce unnecessary manual labor and lower system manning costs,
- Increase the range of operations and extend manned capabilities,
- Reduce the time required to conduct time-critical operations, and
- Provide increased levels of operational reliability, persistence and resilience.

Increased levels of system autonomy are being envisioned for a number of Air Force operations.

1.1 Air

Manned aircraft have experienced increasing levels of automation over the past three decades. The F-35 contains over eight million lines of code, including advanced automation for sensor fusion, voice recognition, and missile and threat management systems [2]. This trend is likely to continue into the future, as autonomy is applied to a much wider variety of tasks and functions for manned aircraft, both on-board and in supporting functions, as shown in Figure 1.

Similarly, increased automation and system autonomy may enhance operation of remotely piloted aircraft (RPA). Currently most RPAs in the Air Force inventory are primarily controlled manually, waypoint-to-waypoint or through flight management computers, by pilots who are located external to the aircraft. Future RPAs will be capable of performing many functions autonomously, allowing them to be employed in areas where: (1) people would be at high levels of risk (e.g., near to hostilities), 2) communications links for direct control are unreliable due

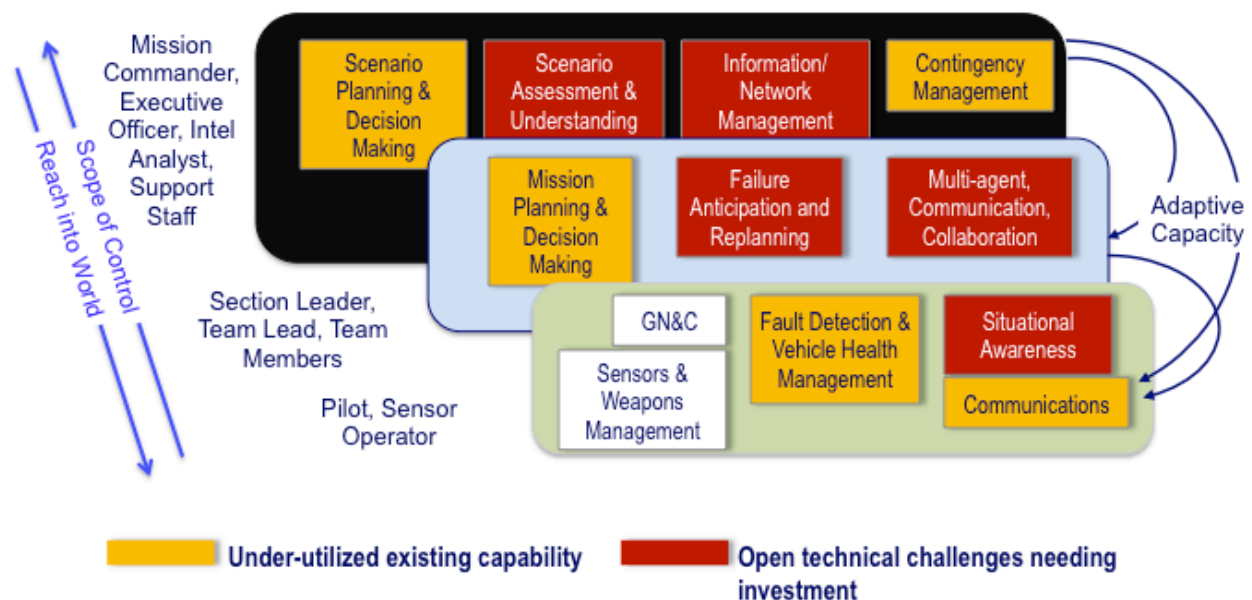


Figure 1. Areas for insertion of autonomy in aircraft systems [3]

to jamming or other interference effects, (3) where speed of operations is useful (e.g., re-tasking sensors based on observed target features), or (4) to undertake new forms of warfare that may be enabled by intelligent, but expendable, systems, or closely coordinated flights of RPAs (e.g., swarms)[4]. In addition, close teaming of manned and intelligent unmanned systems will allow manned aircraft to offload certain functions or extend their payloads significantly.

1.2 Space

U.S. military operations depend significantly on space assets for communication, navigation, and intelligence, surveillance, and reconnaissance (ISR) capabilities. The current network of satellites operated by the Air Force provides this much needed functionality across Department of Defense (DOD) operations. These assets may be at risk, however, if future adversaries work to degrade, deny, or disrupt our ability to operate in space. Autonomy provides a means to build resilient networks that can reconfigure themselves in the face of such attacks, preserving essential functions under duress. It also provides a mechanism for significantly reducing the extensive manpower requirements for manual control of satellites and generation of space situation awareness through real-time surveillance and analysis of the enormous number of objects in orbit around the Earth

1.3 Cyber

Air Force operations in air and space, as well as systems for their command and control, are highly dependent on software and electronic systems that are vulnerable to cyber-attack. Due to the rapidity of cyber-attacks, and the sheer volume of attacks that could potentially occur, there is a need for autonomy that can react in milliseconds to protect critical systems and mission components. As these speeds are far faster than human operators can perform, system autonomy will form a critical aspect of cyber defense. (Although normally considered distinct from cyber-warfare, electronic warfare (EW) is often subject to similar time demands, and would similarly benefit from autonomy.) In addition, the ever-increasing volume of novel cyber threats creates a need for autonomous defensive cyber solutions, including cyber vulnerability detection and mitigation; compromise detection and repair (self-healing); real-time response to threats; network and mission mapping; and anomaly resolution.

1.4 Command, Control and ISR

In the future, a far more integrated network of air, space, and cyber assets will operate in close coordination to provide desired effects. This network will be brought together by a federated system for command and control, fed by high levels of situation awareness drawn from across the various assets and components of air, space, and cyber, as well as dedicated ISR systems [5, 6]. Autonomy can perform a number of important functions to support this vision, including:

- Dynamic reconfiguration for maintaining an effective battlespace network, particularly in the face of anti-access/area-denial (A2AD) activities by potential adversaries,
- Integrating information across multiple sensors, platforms and sources,
- Fusing information in effective ways to provide not just data (level 1 situation awareness), but also meaningful understanding of the data in light of operational goals (level 2)

- situation awareness) and projections of future actions and events (level 3 situation awareness) matched to individual airman mission roles and decision needs,
- Intelligent flows of information across the network and information prioritization to ensure that needed information is provided to the right platforms and airmen in the system, and
 - Assistance in mission planning, re-planning, monitoring, and coordination activities.

1.5 Readiness and Sustainment

Significant benefits are available through the application of automation and autonomy to logistics, maintenance and other support activities, leveraging often currently available technology, to promote higher levels of operational readiness. Autonomous robotic vehicles can be used to deliver materials, and survey and repair runways. Automation can be used to integrate information across disparate logistics systems, create real-time health monitoring of platforms and assets, and to optimize the flow of parts, fuel, and expendables to needed locations. Automation can be used to reduce costs associated with aircraft maintenance, including end-to-end integration from the printing of parts to installation and system checking.

2.0 Definitions

A number of related concepts and technologies may be drawn upon to deliver these potential capabilities.

2.1 Automation

Automation has been applied in a wide variety of systems and generally includes the application of software to provide logical steps or operations to be performed. Traditional automation can be defined as that in which “*the system functions with no/little human operator involvement: however, the system performance is limited to the specific actions it has been designed to do*”[3]. Automation applied to aircraft systems has included the introduction of fly-by-wire technologies for flight control systems, data fusion for integrating information derived from multiple sensors, automation for guidance and navigation (e.g., flight management systems), and the more recent introduction of systems for automated recovery of aircraft in danger of an impending collision with the terrain. These examples show that many systems can be *semi-autonomous*, employing various levels of automation (from low level to more sophisticated) on one or more functions.

2.2 Autonomy

Recently, the term autonomy has gained favor in the computer science community. In general it involves the use of additional sensors and more complex software to provide higher levels of automated behaviors over a broader range of operating conditions and environmental factors, and over a wider range of functions or activities. Autonomy is often characterized in terms of the degree to which the system has the capability to achieve mission goals independently, performing well under significant uncertainties, for extended periods of time, with limited or non-existent communication, and with the ability to compensate for system failures, all without external intervention [7, 8].

To achieve this goal, autonomy incorporates “*systems which have a set of intelligence-based capabilities that allow it to respond to situations that were not programmed or anticipated in the design (i.e., decision-based responses). Autonomous systems have a degree of self-government and self-directed behavior (with the human’s proxy for decisions)*” [9]. Software approaches may extend beyond computational logic-based (or, more commonly, rule-based) approaches to include computational intelligence (e.g., fuzzy logic, neural networks, Bayesian networks) in which intelligent agents communicate and cooperate to achieved desired goals. In addition, learning algorithms can provide the capability to learn and adapt to changing circumstances [7]. Autonomy can be thought of as a significant extension of automation in which very high-level mission-oriented commands will be successfully executed under a variety of possibly not fully anticipated circumstances, much as we currently expect intelligent humans to operate when given adequate independence and task execution authority. Autonomy can be considered as well-designed and highly capable automation.

2.3 Remotely Controlled Vehicles

Unmanned air vehicles, unmanned ground vehicles and unmanned surface and underwater vehicles will form a significant part of future military operations. Currently, most of these systems involve human operators remotely controlling the vehicle with the assistance of fairly low levels of automation for some functions (e.g., the operator specifies waypoints to be followed by the platform). In the future, these remotely controlled vehicles may include more autonomous functionality, however the two concepts are actually orthogonal, meaning it is possible to have one without the other, or both together. Remotely controlled vehicles can be directly controlled by humans through teleoperation, may be semi-autonomous (employing some automated functions), or may be fully autonomous, and manned vehicles can contain software that allows various functions to be manual, semi-autonomous, or fully autonomous.

Autonomy can best be thought of as one potential end of the spectrum of control. However, over the next 30 years most applications will actually involve some level of semi-autonomous capabilities. That is, we will see a gradual evolution of system control, with intermediate levels of autonomy being applied to various functions. As the autonomy developed becomes more capable over time, can handle a greater range of functions, and can handle greater ranges of variability in the environment, systems will slowly evolve to more autonomous operations for longer periods of time. However, for most operations, a requirement will still exist for the autonomy to interact with airmen in order to receive commands, acquire operational needs, and coordinate actions.

3.0 Challenges for Use of Autonomous Systems

Over the past 30 years, extensive work has been undertaken to create intelligent software using various forms of artificial intelligence, including: (1) agent-based reasoning (e.g., rule-based expert systems, Bayesian belief networks, particle filtering, case-based reasoning, fuzzy logic), (2) biologically-inspired reasoning (e.g., neural networks, genetic algorithms, ant-colony optimization), (3) machine learning systems (e.g., data mining, supervised and unsupervised classifiers, “deep” neural networks), (4) naturalistic interfaces e.g., (natural language processing, semantic analysis, speech and gesture recognition), and (5) hybrid modeling approaches combining one or more of these technologies. In addition, more traditional automation software

has been widely applied across aviation and industrial applications for many years, from the earliest Sperry gyroscopic autopilot a century ago, to the latest flight management computers used in the commercial aviation on a daily basis. Research conducted on these systems provides a firm foundation for understanding the challenges involved in creating effective autonomy for future Air Force systems that must operate in complex, dynamic, and often unpredictable environments.

3.1 System Capabilities

Automation has traditionally provided the advantage of creating consistent, reliable and predictable performance of actions according to its programming, as shown in Figure 2. The challenge has been that the suitability of those actions is often limited to a constrained set of situations – ones that the designer has envisioned and the software developers have programmed for – and a constrained set of measurements available from a limited sensor suite that is limited in its ability to sense and understand the environment it is operating in. Creating systems that can accurately not only *sense* but also *understand* (recognize and categorize) objects detected, and their relationship to each other and broader system goals, has proven to be significantly challenging for automation, especially when unexpected (i.e., not designed for) objects, events, or situations are encountered. This capability is required for intelligent decision-making. Unfortunately, most automation to date has suffered from *brittleness*, that is, operating well for the range of situations it is designed and programmed to address, but needing human intervention to handle all the cases and situations it is not designed/programmed to handle.

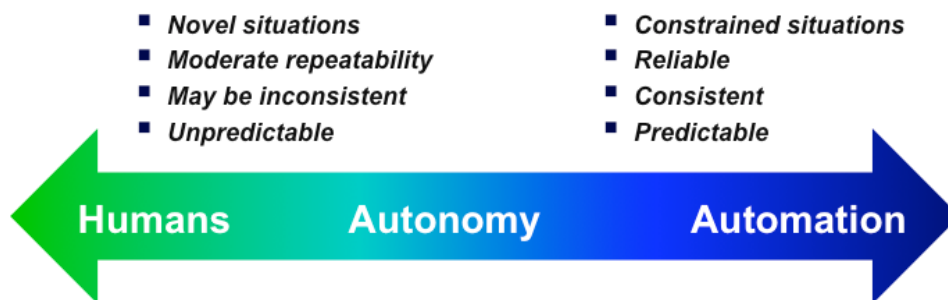


Figure 2. Relative advantages of humans and automation, with the characteristics of autonomy falling in between the two

Humans, on the other hand, while variable from person to person, are often credited with a superior ability to keep in mind the “big picture” (the overall mission objectives), assess the situation (context for action), to think on the fly, and to adapt to novel situations in a way that is not rule-based (like most software programs) but relies much more on pattern recognition, mental models, and analogical reasoning, sometimes at very abstract levels. In addition, in a warfare situation, the fact that humans can act in a fashion that is unpredictable by the adversary provides a significant advantage. On the other hand, humans are not as good at processing large volumes of data, quickly and consistently, nor of sustaining attention for long periods of time.

As the capabilities of autonomy increase, (including the ability to handle a broader range of situations and uncertainty) it is anticipated that the need for human intervention will decline, however, it is likely that some level of human-system interaction will continue to be required for the foreseeable future. This is to be expected for a number of reasons:

- As hardware complexity grows, there will be more opportunity for failures,
- As software complexity grows there will be more opportunities for bugs and vulnerabilities, and
- As these systems are injected into an adversarial environment, there will be opportunities for encountering situations that the original designers had never considered.

In short, considerable system complexity will be created as the software and hardware is expanded to try to cover more situations or modes of operations, and as the systems are used in more complex environments. This can result in: (1) reduced understandability of the system because of its complexity (why did it do that?), (2) reduced predictability in terms of how it will perform in any given situation, challenging the people who must interact with it, and (3) greater vulnerabilities through the communications links created for human intervention used to offset the first two items.

As a result, most or all Air Force operations conducted in the foreseeable future will require a combination of both humans and autonomy to get the job done in the face of a broad range of operational conditions and a determined adversary. Autonomy will be used to reduce manual data processing and integration, provide speed, and carry out actions within the capabilities of its software and hardware. Airmen will still be needed to provide the command guidance and control associated with directing the high level goals of the autonomy, for their ability to interject knowledge outside of its design boundaries and/or contextual awareness, to deal with novel situations, and for coordination with other forces and activities.

Based on extensive experience with automation over the past 50 years [10-16], a number of hurdles will need to be overcome to for successful implementation of autonomy in Air Force operations. The ability of people to effectively use automation has been strained by a number of factors, including reduced situation awareness, increased workload, increased decision making time, and difficulties in determining appropriate levels of trust. We anticipate similar issues to arise with the continuing introduction of increasingly autonomous systems.

3.2 Situation Awareness and Out-of-the Loop Performance Problems

Situation awareness when working with autonomous systems is critical for ensuring that they are operating in ways that are consistent with operational goals. A key challenge experienced by people overseeing automation is that they become *out-of-the-loop*, that is, slow to detect that a problem has occurred with the automation, or with the system being controlled by the automation, and then slow to come up to speed in diagnosing the problem to intervene appropriately [17, 18]. This occurs due to a fundamental lowering of situation awareness when using automation that occurs due to: (1) interfaces that do not provide needed information and often little feedback on system state, (2) systems that require extensive human monitoring, a skill that people do not excel at due to decrements in vigilance that can occur after as little as 30 minutes, and (3) a shift from active to passive processing of information [14]. Many aviation

accidents have occurred due to pilots becoming out-of-the-loop with respect to the automation and unable to intervene appropriately in a timely manner [10].

In addition, pilots have experienced considerable challenges in developing a good understanding of what the automation is doing, even when actively attending to it during normal operations. Misunderstandings of displayed information, sometimes due to misinterpretations of system mode, and an inability to accurately predict what the system will do in a given situation, creates inaccurate situation awareness and poor decision making [15]. Even highly trained pilots have been found not to fully understand all the modes of automated flight navigation and guidance control systems, creating real challenges for effective interaction with the automation [19]. Future systems will need to pay significant attention to the development of autonomy approaches that emphasize maintaining required levels of situation awareness for all airmen.

3.3 Optimal Workload Levels

While automation is often introduced with the goal of reducing manual workload (and thereby potentially reducing the manning required for an operation), it has often not accomplished that goal. It has been called the “irony of automation” that it often increases workload during high workload phases of flight (e.g., aircraft landing and take-off) and decreases workload during low workload phases of flight (e.g., en-route) [11]. Workload can often subtly shift from observable manual tasks to not so observable cognitive tasks as understanding and interacting with automation increases demands [10, 20]. Significant advances in creating autonomous systems that are easy to use, understand, and interact with will be required. In addition, significant attention needs to be paid to the selection of which tasks to automate or delegate to an autonomous system, so that the airman is provided with a coherent set of tasks that are suited to human capabilities, rather than leaving airmen with a disjointed set of left-over tasks that cannot be easily automated. This more human-centered approach will be important for creating optimal human-autonomy performance.

3.4 Integrating Human & Autonomy Decision Making

Autonomy is often directed at supporting human decision-making. Expert systems or decision support systems act to provide decision guidance, such as creating or rating courses of action, target cueing, or classifying detected targets. Effective decision support turns out to be difficult however [21-23]. While it is generally assumed that such systems will boost human decision making, particularly with difficult tasks, this is often not the case. Evidence shows that people actually take-in system assessments and recommendations which they then combine with their own knowledge and understanding of the situation [24]. A faulty decision aid can lead to people being more likely to make a mistake due to decision biasing by the aid. And the time required to make a decision can actually increase, as it is an additional source of information to take into account. Therefore, overall human/system decision accuracy and timeliness may not necessarily be increased by decision aiding systems if those systems are imperfect. While good advice can help, poor advice has a large effect on leading the decision maker astray, and overall mission performance can be significantly degraded.

Conversely, decision support systems that critique human decisions (e.g., point out potential problems with a planned course of action), have been found to work much better, eliminating the problems of biasing the human towards the computer’s solution because its inputs occur *after* the

human makes a decision rather than before [25]. It also plays to one of the computer's strengths: an ability to conduct fast-time simulations of a human-proposed solution (e.g., a course of action) to identify potential flaws or dead-ends across a multiplicity of environmental situations and adversarial actions. This provides better airman/automation synergy and improved overall combined performance.

As the use of intelligent agents and systems that make decisions increases in the future, significant care will need to be taken to develop cognitive interaction schemes that *enhance* rather than potentially degrade airman decision making [26]. In addition, the operational effects of such systems will need to be carefully tested based on combined human/system performance outcomes.

3.5 Informed, Situational Trust in Autonomy

In order for airmen to operate effectively with autonomous systems, they will need to be able to determine how much to trust the autonomy to perform its tasks. This trust is a function of not just the overall reliability of the system, but also a situationally determined assessment of how well it performs particular tasks in particular situations. For this, airmen need to develop informed trust – an accurate assessment of when and how much autonomy should be employed, and when to intervene. As shown in Figure 3, appropriately calibrating one's trust, along a scale from over-trust (complacency) to under-trust (resistance) can be difficult and is based on several factors [27-35]:

- **System factors** - including overall validity and reliability, subjective assessments of that reliability, the recency of a system failure (or inability to act appropriately in the situation) with trust lost quickly and slow to return after a problem is experienced, system understandability and predictability, timeliness and integrity,
- **Individual factors** – including the individual's perceived ability to perform the task, general willingness to trust, and other personal characteristics, and
- **Situational factors** – including time constraints, workload, effort required, and the need to attend to other competing tasks.

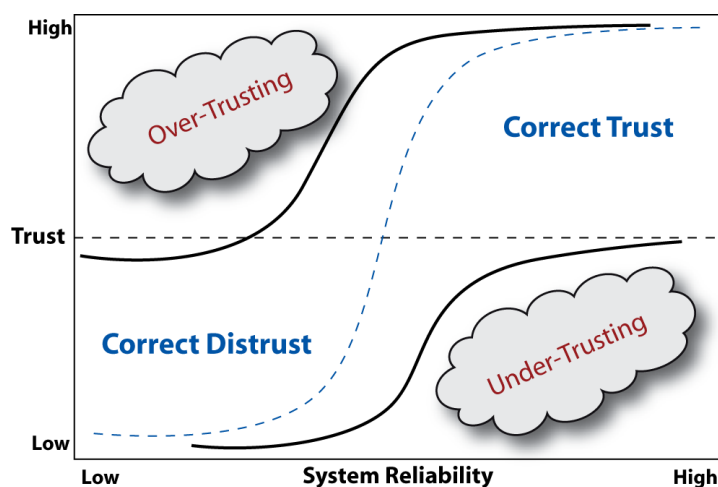


Figure 3. Appropriate calibration of trust in autonomy is critical

Distribution A. Approved for public release; distribution is unlimited. Public Release Case No 2015-0267

In the past airmen have made separate decisions about: (1) **Information** - how much to trust information inputs (e.g., confidence level based on the source of the information, reliability of a sensor, timeliness of the information, and confirmation from other sources), (2) **Others** - how much to trust other teammates, and (3) **Automation** - how much to trust a particular automated tool. However, with the expansion of autonomous systems in the future, the lines between these separate entities will become blurred. An autonomous system may be the source of information, the system that processes it and performs actions with it, *and* the “other” with whom the airman is interacting. Particular care will need to be taken to allow airmen to develop trust that is well informed so that they will know how much to trust the autonomous system for a particular task, at a particular time, for the particular situation.

4.0 Towards Symbiotic Human-Autonomy Systems

Effective teaming between airmen and autonomy will need to be designed into future autonomous systems for the Air Force. This is for two reasons. First, because of the limitations identified above, it is unlikely that autonomous systems, in the foreseeable future, will have the capabilities to act in a fully autonomous manner, and deal with the full range of mission, environmental and adversarial situations facing it. Second, command and control (C2) is essential for any effective military operations, and there will always be a need for controlling autonomous systems (if only at a task/mission level), assessing their task/mission success, and coordinating with other forces in the mission space. Supporting coordinated activities with airmen will require effective user interfaces to be successful. Approaches for autonomy therefore need to be human-centered – able to support the requirements for operator situation awareness, informed trust, manageable workload levels and ease of interaction that will be needed for airmen to work effectively within this paradigm.



Figure 4. Flexible autonomy should provide smooth, simple, seamless transition of functions between the airman and the system

4.1 Flexible Autonomy

Flexible autonomy forms a central tenet for future human-autonomy interaction. Because no autonomous system will work perfectly in all situations, any given function may need to be augmented or performed by the airman at different times. Control for functions, sub-systems, or even entire vehicles, will need to be able to pass back and forth between the airman and the autonomous system over time, Figure 4 [36-38]. Coordinated performance requires that this transition be smooth, simple and seamless; enabled by a high level of shared situation awareness, and efficient methods of interaction.

4.1.1 Levels of Autonomy

Different levels of autonomy may be appropriate at different times. Autonomy is not an all or nothing proposition. Rather there are intermediate levels of autonomy (often referred to as semi-autonomy) in-between fully manual and full autonomy, as shown in Figure 5. The level of

autonomy increases as the capability of the system increases for performing various components of any given function [39, 40]:

- **Execution of Actions**– the degree to which the autonomy carries out tasks and sequences of tasks,
- **Monitoring and Information Fusion** – the degree to which the autonomy assesses the state of the environment and the system, and integrates disparate data to characterize that situation,
- **Option Generation** – the ability of the system to generate potential options or courses of action for decisions, and
- **Decision Making** – the ability of the system to select between options to determine the actions to put into place.

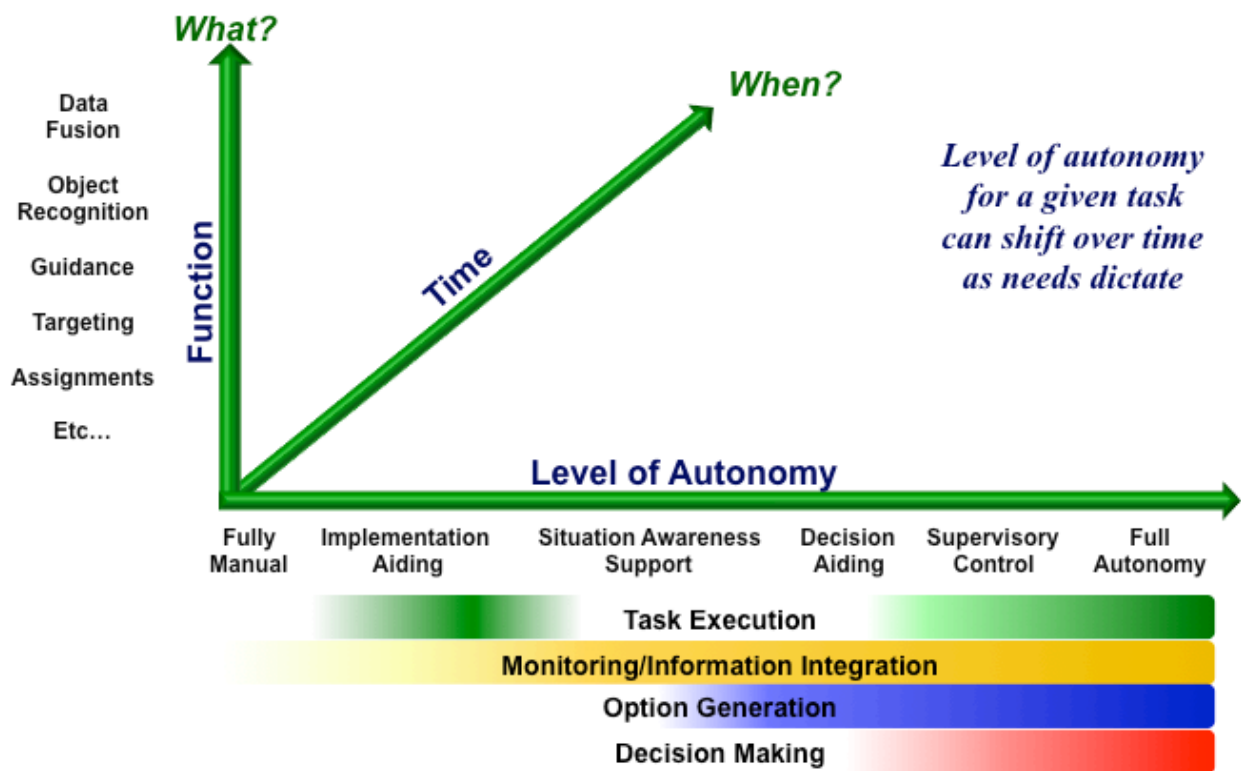


Figure 5. Flexible autonomy will involve shifts over time in the level autonomy that is applied to each function based on the situation.

Based on the capabilities of the system across these four components, different levels of autonomy (LOA) are possible. Some common LOAs include:

- **Fully manual** – where all aspects of task performance are completed by the human
- **Implementation aiding** – where the system carries out tasks for the human, such as flight management systems or smart weapons that follow human targeting, but the human makes all decisions,

- **Situation awareness support** – in which disparate data are fused to provide integrated information relevant to operator decisions and goal states,
- **Decision aiding**– where the system provides a list of potential options and rates or ranks those options as to their suitability, such as with a recommended target list or course of action assessment, (it may or may not select the best option),
- **Supervisory control** – where the system controls all aspects of a function automatically, including taking in information, deciding on correct actions and carrying out those actions, but the human can set goals and intervene as needed (also called on-the-loop control), and
- **Full autonomy** – provides full control over all aspects of a function, without human guidance or the ability to intervene. For example, the Automatic Ground Collision Avoidance System, currently fielded on F-16s, continuously monitors for impending ground impacts, projects potential escape trajectories, takes control and executes recoveries at the last possible instant, then returns wings-level control to the pilot.

Intermediate levels of automation have the advantage of inducing higher levels of situation awareness for the airman, as compared to full autonomy, by creating a more active role for the airman, as well as often being easier to implement technically [14]. In general, people have less difficulty in maintaining situation awareness with autonomy that aids with the gathering and presenting of relevant information and that aids in carrying out tasks, as compared to autonomy that involves higher level cognitive functions like generating options or selecting the best course of action [41]. This is particularly true if the autonomy is not perfectly reliable [24]. Autonomy for information acquisition and analysis supports better human/autonomy team coordination, leading to better team performance. Autonomy of decision making, however, leads to higher workload and is only generally effective under low workload conditions [42, 43].

4.1.2 Dynamic Autonomy Selection

The choice of when to use autonomy for a particular task or function, and what level of autonomy to use, will be a dynamic decision that can change over time. An airman may want the system to operate in a different autonomous fashion in certain circumstances as compared with others, or may want to take over a task manually in others. This shifting of control can depend on a number of factors, as shown in Figure 6. In general, the airman will be more willing to use autonomous systems (and at a higher level of autonomy) when:



Figure 6. Autonomy use shifts dynamically based situational factors

- **Trust** – the autonomy can be trusted to successfully perform the task in the current situation,
- **Verification** – it is possible to dynamically confirm or verify the autonomy’s performance in real-time,
- **Risk** – there is less risk of negative consequences (e.g., collateral damage), it is possible to mitigate risk (e.g., limit the scope or area of operation so as to avoid collateral damage), or in phases of the mission when more risk is acceptable (e.g., after the commencement of hostilities as compared to earlier phases when it is important to avoid potential provocation),
- **Need** – when the unique capabilities of autonomy are needed (e.g., to avoid putting airmen in situations where they are subject to hostile fire, to perform actions very quickly, or to extend capabilities),
- **Partnership** – when the autonomy is easy to interact with, understandable, and its future behaviors can be predicted.

In most cases, the decision to invoke autonomy should be in the hands of the airman. By being in control, the airman can stay in-the-loop and is less likely to fall prey to out-of-the-loop performance problems. However, in some cases it may be necessary for the autonomy to act on its own. This may occur in situations where the airman is unable to team with the system (e.g., loss of consciousness or communications), where there is not enough time available for the airman to be in-the-loop (e.g., defense against an incoming missile or cyber attack), or where serious consequences are imminent (e.g., collision with terrain or another vehicle). New developments in light weight, non-intrusive sensors that measure pilot physiology or neurological states may be integral to supporting such capability. In these cases, it is important that the autonomous system’s behaviors are clearly displayed to the airman and that the ability for manual override is available.

4.1.3 Limitations on Autonomy for Weapons Systems

Department of Defense Directive 3000.09 requires that autonomous and semi-autonomous systems that are developed for weapons systems be designed to allow commanders and operators to be able to exercise human judgment when the use of force is involved [44]. This includes both kinetic and non-kinetic systems and guided munitions that can independently select and discriminate targets. It requires that:

- *Semi-autonomous weapon systems that are onboard or integrated with unmanned platforms must be designed such that, in the event of degraded or lost communications, the system does not autonomously select and engage individual targets or specific target groups that have not been previously selected by an authorized human operator.*
- *The system design...addresses and minimizes the probability or consequences of failures that could lead to unintended engagements or to loss of control of the system.*
- *In order for operators to make informed and appropriate decisions in engaging targets, the interface between people and machines for autonomous and semi-autonomous weapon systems shall: (a) be readily understandable to trained operators, (b) provide traceable feedback on system status, and (c) provide clear procedures for trained operators to activate and deactivate system functions. [44]*

This directive is particularly important for minimizing the probability and consequences of failures in autonomous and semi-autonomous weapon systems that could lead to unintended engagements and for minimizing the potential for collateral damage. It also implies that the airman is ultimately in control of the system and should be both responsible for decisions and properly supported by the system to make those decisions. This will require a high degree of situation awareness over both the state of the system and the environment in which the system operates. With higher levels of autonomy, it will also require that future systems support shared situation awareness between the airman and the autonomy.

4.1.4 Human–Autonomy Interaction

In addition to the level of autonomy provided by a system (and the range of levels that can be accessed by the airman), several other features of the autonomy will have an impact on the creation of a smoothly functioning airman-autonomy team.

- **Robustness** – The degree to which the autonomy can sense, understand, and appropriately handle a wide range of conditions determines how robust it is. In the past, automation has often been brittle, only capable of handling preset assumptions about the world in which it operates and unable to easily handle boundary conditions. This creates a situation in which the airman must function as a trouble-shooter for difficult cases, often with limited situation awareness and understanding of what the autonomy has been doing. To the degree that future autonomous systems are more robust, able to handle a wider variety of conditions, able to adapt to changing conditions, and more inclusive of the airman in providing a joint solution space, the better will be the overall mission performance.
- **Span of Control** - The span of control allotted to an autonomous system can vary from only very specific tasks for specific functions, up to autonomy that controls a wide range of functions on a system. For example, autonomy with a wide span of control may act to diagnose system problems and take actions to correct them, order maintenance actions, and modify flight performance to compensate for changes. The wider the span of control, the more autonomous the system, but also the greater the need for effective communication with the airman on actions taken by the autonomy.
- **Control Granularity** - Autonomy can also vary in terms of the level of detail in the breakdown of tasks for control it requires (i.e., the necessary level of micro-management vs macro-management over the autonomy), shown in Figure 7. The system may require the airman to program each detailed task that is required and specify parameters for task performance (e.g., set up the automation each time), to select from a playbook of pre-set sequences or behaviors [45], or it may only require a high level command or goal to be provided to the system, with the autonomy undertaking a complex set of tasks to meet that goal. The airman’s workload should generally



Figure 7. Level of Control Granularity

decrease with less control granularity (e.g., goal-based command guidance). However, it should be noted that increased queuing of tasks (setting up a sequence of tasks to be performed over time) has been shown to decrease situation awareness, leaving people more out-of-the-loop [41]. Therefore as less control granularity is required, extra measures will be needed to keep the airman informed on autonomy performance in carrying out its tasks and projections of future actions.

4.2 Building Shared Situation Awareness to Support Airman-Autonomy Teams

Airmen working with autonomy will need to be able to answer a number of questions to properly oversee the system and to determine when interventions or shifts in level of autonomy are needed. They need information to support a number of assessments:

- How much confidence to place in the autonomous system?
- Is the autonomous system working properly?
- Is it getting good data?
- Is it operating within the envelope of situations it is programmed to handle?
- Will the system's actions meet the operational goals?

Conversely, the autonomy may need to make the same assessments regarding the airmen it is working with. These questions require that the user-interface for the system be carefully designed to create a high level of situation awareness.

4.2.1 Human-Interaction Guidelines

Several decades of research on automation-induced human error provide a good basis for the development of interfaces that will avoid critical problems and mission failures. A number of key design guidelines should be followed to support the need for situation awareness and effective control of autonomous systems so that the airman can function effectively [24].

- ***Automate only where necessary and at the lowest level necessary*** – In many cases, adding autonomy to reduce manual workload is not needed. Rather, significant benefits can be derived from applying good human factors design principles to the user interface in order to reduce workload and improve human performance with the system. Only after the interface for system has been optimized should autonomy be applied to reduce workload. This is because higher levels of autonomy are expensive to develop and can create new challenges by adding complexity and cognitive workload and run the risk of lowering situation awareness. If it is possible to meet mission performance objectives with lower levels of autonomy, this is preferable. Autonomy that carries out routine tasks or that integrates information to support airman situation awareness provides significant benefits and should be encouraged. Autonomy that generates options and makes decisions is far more complex and difficult for the airman to understand and interact with, and should be regarded cautiously.
- ***Provide Automation Transparency*** – The state of the autonomy and its intended actions must be made highly transparent to the airman. The current goals and assumptions of the autonomy, its current and projected actions, and how much confidence should be placed in its data and algorithms should be clearly represented.

- ***Keep operator in control and in the loop*** – Airmen will be more effective at interacting with the system if they are in-the-loop and active in making decisions about the autonomy and controlling its operation. Situations where the autonomy is activated without specific input from the airman should be minimized to situations of imminent danger (e.g., aircraft collision, defensive actions of a cyber system) where the airman is either unable to respond, due to being incapacitated for example, or unable to make a decision in the extremely short timeframes available.
- ***Avoid the proliferation of automation modes*** – One method of adding capabilities to automated systems has been to create modes of operation that can handle different situations the system may encounter, or different operator preferences. It will operate slightly differently, with different rules, in each mode. Modes, however, add significant complexity to the system, making it harder for airmen to understand what the system is doing, and should be avoided as much as possible.
- ***Make modes and system states salient*** – When automation modes are included, it is imperative that the mode state be made highly salient to avoid situation awareness errors that can occur when the airman misinterprets the current or future actions of the system based on a misunderstanding of its mode.
- ***Enforce automation consistency*** – Consistency in terminology, information placement and functioning of the autonomous system is critical for creating autonomy that is both understandable and predictable. The more consistent the logic throughout its functions, the better. In cases where autonomy functioning varies in different situations or modes, it is very important to make such behaviors transparent to the airman.
- ***Avoid advanced queuing of tasks*** – Autonomy that queues up multiple tasks (e.g., a flight management system that presequences multiple waypoints) will put airmen more out-of-the-loop than systems that involve the airmen at each step. If the system does provide advanced queuing, it will be even more critical to emphasize interface transparency and steps should be taken to keep the airman engaged.
- ***Avoid the use of information cueing unless highly reliable*** – Automation that attempts to reduce clutter by cueing the airman's attention to certain parts of the scene or subsets of information must be highly reliable to be effective. If such cueing is unreliable, airmen will be significantly hampered in overcoming its deficiencies.
- ***Use methods of decision support that create human/system symbiosis*** – Since airmen may be biased by decision support systems and unable to overcome poor advice, decision support systems that work through critiquing, which support the consideration of different options, interpretations, or contingency planning, and those that integrate information to provide situation understanding and projections are preferable.

4.2.2 Shared Situation Awareness

Adherence to these guidelines for creating effective airman interfaces for autonomous systems will significantly improve the degree to which airmen can be successful in working with the system to achieve mission objectives. As more capable autonomy is developed, which has more advanced intelligence designed to cover a far wider range of situations and functions, the ability of the airman to understand what it is doing so as to interact with it properly will be highly taxed. With future autonomous systems, it will be critical to create advanced interfaces that support the need for shared situation awareness between the airman and the autonomy.

Shared situation awareness is fundamental to supporting coordinated actions across multiple parties who are involved in achieving the same goal and who have inter-related functions such as those that occur with flexible autonomy.

Situation awareness is defined as *“the perception of the elements in the environment, within a volume of space and time, the comprehension of their meaning, and the projection of their status in the near future”* [46]. As such it involves more than just low level data (Level 1 SA), but also the integration of that data to provide an understanding of the significance of that information for mission goals (Level 2 SA) and the ability to project what is likely or possible to happen in the near future (Level 3 SA) that is important for proactive decision making. It involves information from the external environment, the systems a person is operating, others (blue, red or civilian), and task or mission status, which often must be integrated into a meaningful and useful picture of the situation upon which decisions are made.

Work on developing shared situation awareness between human teammates can be leveraged as a model for supporting shared human-autonomy situation awareness, Figure 8. Shared situation awareness is *“the degree to which team members possess the same situation awareness on shared situation awareness requirements”* [47], which are those common aspects of the situation that are needed for decision making across both roles. With human teammates, even people who get the same input from the same displays, and who are co-located in the same environment, can have challenges in achieving shared situation awareness because they may interpret information differently or form different projections of the future, based on different goals and different mental models of the system and the environment.

Autonomous systems will have computer models (which are likely to be different from operator mental models) for interpreting the information that they take in from their sensors and input sources (which may be different than the information made available to the airman). Thus there is a significant potential for the autonomy and the airman to have very different assessments of the world driving their decisions. To overcome this challenge, it is critical that there be effective two-way communications of the situation models between the airman and the autonomy. This means not only sharing the low-level data upon which each is operating, but also how that information has been interpreted (the comprehension of the meaning of that information in light of operational goals) and what future projections each have made.

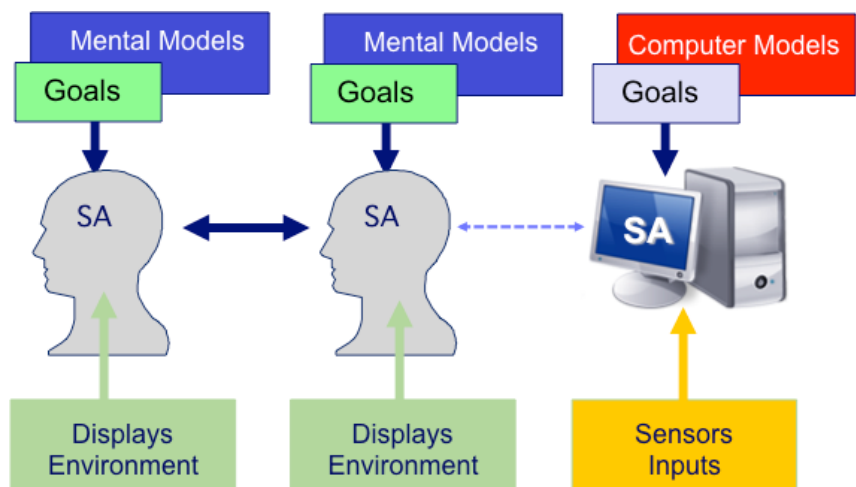


Figure 8. Supporting shared situation awareness between autonomy and human teammates for effective coordination.

Flexible autonomy will require a high level of shared situation awareness to support a number of fundamental operations involved in this concept, Figure 9:

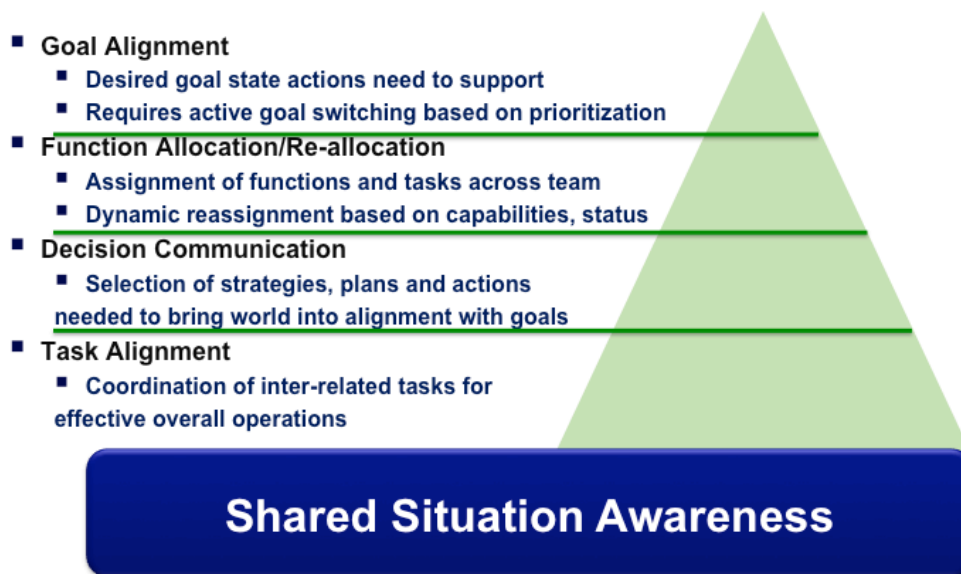


Figure 9. Shared situation awareness underpins flexible autonomy

- **Goals** – The airman and the autonomy need to be focused on supporting the same goals which can dynamically change. For example, it would be a problem if the pilot’s goal is to perform a go-around while the autonomy is trying to land at an airport. As priorities change, goals change, and shared situation awareness is needed to make sure that the autonomy and the airman’s goals are aligned.
- **Function allocation and re-allocation** – Flexible autonomy will involve an ongoing assignment and reassignment of functions to the airman and the autonomous system. Keeping up with who is doing what will be critical, as well as the relative capabilities and status of both the airman and the autonomy for performing various functions.
- **Decision communication** – As the airman and the autonomy make decisions about how to perform their various functions, it will be important that these decisions (including strategies, plans and actions) be shared between them so that actions on related functions can be coordinated.
- **Task alignment** – Tasks being performed by the autonomy and the airman are likely to be highly inter-related and often inter-dependent. Each will need to maintain an ongoing understanding of what actions have been taken by the other and how successful those actions are at achieving shared goals.

To support the need for the operator and the autonomy to successfully coordinate and collaborate across these fundamental aspects of shared performance, it is important that they are each able to align their world views. That is, that they both have an accurate and shared understanding the state of the environment and the system, per the elements in Figure 10. Effective team performance depends on both teammates having a shared understanding of how

well each other are performing, how the actions of the other party are affecting their own tasks, and what the other is planning to do in the future. In addition, to support the needed inter-predictability between the collaborating partners, autonomous systems will require models of their human collaborators' model of the battle space thus allowing them to estimate the meaning of perceived information to the airman (as well as vice-versa).

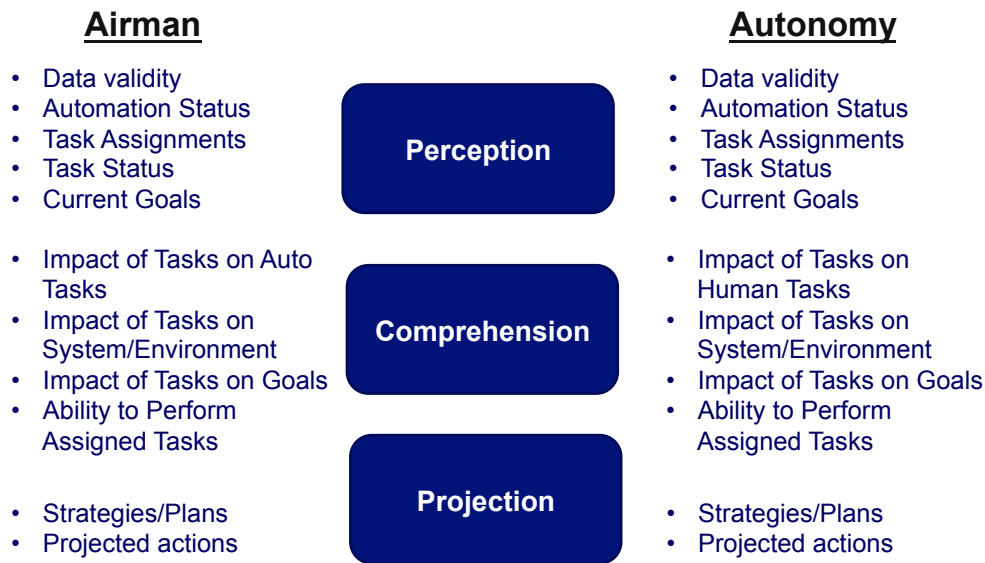


Figure 10. Elements for supporting shared situation awareness across airman and autonomy teams

4.2.3 Supporting Unmanned Operations

Many unmanned system control stations have neglected the interfaces for supporting the airman who is responsible for their operation. The challenges for operating remote aircraft or other vehicles are significant due to [24]:

- Time-lags of as much as two seconds or more in the control loop which can significantly degrade direct manual control of the vehicle,
- Loss of direct sensory information (visual, auditory and haptic), with all control information coming through limited visual displays, overloading limited visual attention,
- Intermittent and noisy data associated with data links that may not always be present and that can be subject to interference from adversaries,
- Perceptually demanding tasks that involve interpreting often degraded sensor imagery and limited fields-of-view, leading to poor situation awareness,
- Multi-tasking across many displays that frequently can overload the airman,
- Poor support for team interaction across the distributed team involved in unmanned vehicle operations (e.g., pilots, analysts, commanders on the ground, etc...)

As autonomy is added to the control station for unmanned aircraft, some of these challenges may be alleviated (e.g., direct control with time-lags, loss of communications, intermittent data links), however, many of these challenges will remain and the additional challenge of

understanding and interacting with the autonomy will be incurred. Significant attention to the development of control stations for unmanned vehicles will be needed to address these challenges. Future systems should include:

- Pilot control interfaces that adhere to military standards for human factors of vehicle control systems,
- Data integration to reduce workload,
- Multi-sensory cues to compensate for loss of haptic and auditory information,
- Improved spatial awareness of the environment and relevant objects in the environment to include self-orientation, wayfinding, contextual awareness and an understanding of situation awareness limitations,
- Predictive displays to assist with compensating for control time-lags,
- Displays to support understanding and projection of autonomous systems operations (including monitoring, diagnosis and mission and payload management), rapid shifts in level of control between the pilot and autonomy for various functions, real-time assessments of trust in the autonomy, and
- Displays to support coordinated action with manned aircraft, multiple unmanned aircraft, and other teammates (e.g., analysts and commanders).

As with manned systems, careful development and testing of the pilot control stations is a critical part of system development and needs focused attention [26, 48, 49].

4.2.4 Trustworthy Autonomy

Helping the airman to develop appropriate levels of trust in autonomous systems is critical for ensuring that neither over-trust (complacency), nor under-trust (disuse), hamper the successful use of autonomous systems. Conversely, future autonomous systems may need to infer a level of trust for the people it interacts with (e.g., whether the airman is able to act effectively given current workload or loss of attention or consciousness).

A number of guidelines should be applied to the design of the autonomy so that its level of trustworthiness can be readily determined. In addition to enhancing the overall system competence, reliability, and robustness, a number of key system attributes should be supported, to encourage *appropriate* levels of trust to be formed.

- **Support cognitive congruence with the airman** – Designs that allow for cognitive congruence or analogical reasoning on the part of the autonomous system will help to support the understandability of the system. This can be achieved by architecting the system, at a high level, to be congruent with the way humans parse problems and form solutions, for example. Automation knowledge management processes should also be designed similarly to the way humans solve problems. Simple tools that help the autonomy and the airman to align goals, courses of action, and functions can help in this effort.
- **Avoid anthropomorphism** - Autonomy designs that look too “human-like” on the surface (e.g., via life-like avatars, facial expressions, hand gestures, body language, etc...) should be avoided since they have the potential to induce people to over-estimate the system’s

capabilities. A glib conversational interface, for example, can create a misattribution of capability, leading to a later loss of trust which can be hard to regain.

- ***Design for transparency and traceability*** - The use of systems that are based on simple logical rules, while likely slow and brittle, provide the advantage of explanatory power regarding system behaviors given state of the incoming data. In contrast, a deep-learning neural network might be fast and adaptive, however, there may be little in the way of an explanation of how it arrives at the conclusions. Trade-offs between transparency for the human and optimality in the decision space should be an explicit consideration in the design of the autonomy. The system must be able to explicitly explain its reasoning in a concise and usable format (either visual or textual), in order to support the airman's need for trust determination.
- ***Support robust visualizations*** – Autonomy designs should feature contextual overviews and visualizations at different levels of resolution. Context is critical for ensuring the robustness of decisions, both for humans and systems. An ability to provide context to the system via high-level inputs (preferably in a natural language context for ease of operation by the airman), and to the airman (via visualizations that support variable levels of resolution and abstraction), can provide both members of the airman-autonomy team with the needed shared situation awareness.
- ***Create system self-health assessment*** – Autonomy will need to achieve self-awareness of its health integrity. This calls for the maintenance of meta-information on the system's data, information, and knowledge (e.g., staleness, reliability, etc...). Health management subsystems should also monitor the communications channels, knowledge bases, and software applications used for potential contextual violations of the underlying assumptions used in the system's design, and inconsistencies in communications commands, as well as for potential cyber breaches. The autonomy must go far beyond simple database integrity checking, to provide consistency checkers at more abstract levels, similar to health monitoring systems used for flight management systems.
- ***Support airman-autonomy joint training*** - Mixed-initiative team training will be needed as part of any system development and deployment effort. Extensive airman-autonomy team training should emphasize both the nominal design envelope, and situations outside that envelope (e.g. unusual events, and areas where the system is less capable). This will help in developing an understanding of common team objectives, the separate roles of the airman and the autonomous system, and the ways in which they are co-dependent. This will provide the airman with an understanding of the system's limits of operation, as well as telltale behaviors might be associated with the system's behavior as it approaches its limits. If the autonomous system is also capable of assessing its human partner, this mixed-initiative training will help them to develop mutual mental models of each other, creating expectations for competence, dependability, predictability, timeliness, and uncertainty reduction.

5.0 Challenges for the Development of Autonomous Systems

There are a number of technical challenges associated with developing successful autonomous systems. In particular future systems will need to be more robust, avoiding the problem of brittleness that has significantly limited previous systems. The Air Force Research Laboratory Autonomy Science and Technology Strategy describes several key goals for

addressing these challenges, including: (1) deliver flexible autonomy systems with highly effective human-machine teaming, (2) create actively coordinated teams of multiple machines to achieve mission goals, (3) ensure operations in complex contested environments, and (4) ensure safe and effective systems in unanticipated and dynamic environments [9]. To support these goals, the development of intelligent autonomy that operates on the basis of an accurate situation model (analogous to human situation awareness) will be needed.

5.1 Situation Models for Autonomy

For autonomy to be successful, it will have to evolve beyond simple computational logic to systems that can reason based on a more complete understanding of its evolving mission and environment. As shown in Figure 11, such systems will need to incorporate [50, 51]:

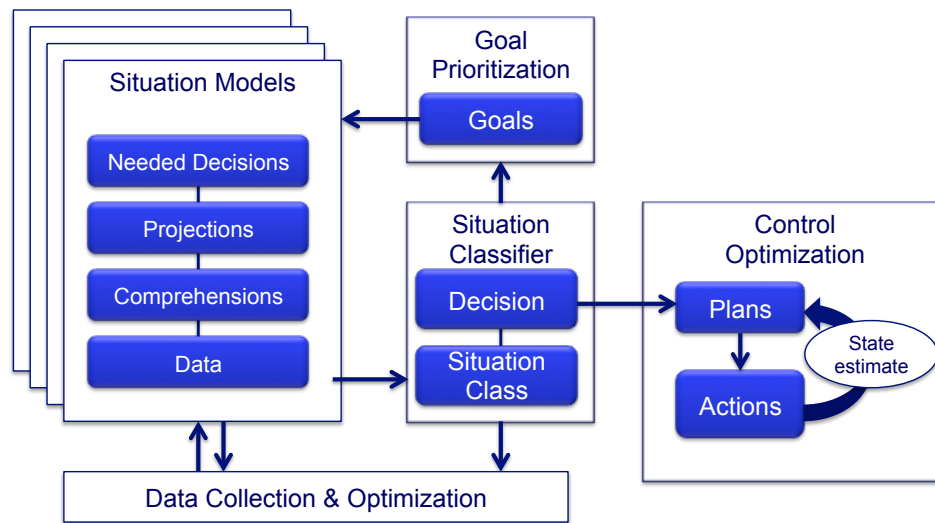


Figure 11. Cognitively inspired architectures for autonomy situation models

- A computer model that allows for an ongoing representation of the current situation, fusing multiple sensory inputs into situation comprehension and projection to support decision making,
- Pattern matching that maps the current situation against learned situation categories which in turn correspond to stored plans and actions,
- The use of situation models to provide expectations for directing attention to relevant information and interpreting information, for providing defaults for missing information in the face of partial data,
- A context model for representing uncertainty in situation representations,
- Goal-based behavior that directs the search for information and interpretation of that information (goal-driven behavior),
- The ability to recognize critical environmental cues that map to different situation categories,
- The ability to handle multiple goals and dynamic goal re-prioritization based on the comparative status of situation categories and projections of future states, providing data-driven behavior,

- Models of systems, the environment, and other actors (e.g., teammates, adversaries, civilians) to allow reasoning when pattern matching does not provide a good fit with existing situation categories,
- Active learning and refinement of situation categories and models,
- The ability to create plans to achieve goals and dynamically replan as needed.
- The ability to optimize the collection of data from the environment, systems and others as needed to support the ongoing requirements of the situation models, and
- Although not illustrated in Figure 11, the ability to interact with human operators or other autonomous systems on any of these functions (e.g., alignment of goals, situation models, decisions, function allocations and prioritizations, and plans) to achieve coordinated and approved actions.

5.2 Learning Systems

The use of learning systems in the development of system autonomy introduces certain advantages and disadvantages that will need to be addressed if they are to be successful. On the positive side, systems that apply learning based algorithms may be capable of addressing a wide range of contextual factors (e.g., environmental, mission-specific, and adversary induced) to create more robust solutions over a broader scope of situations. In such cases, the designers and developers do not need to anticipate all possible conditions in advance (which has proven very difficult to do, both theoretically and in practice) and pre-specify what should be done in each possible combination of conditions. Rather, they create a structure that allows the autonomy to organize itself and learn and adapt to changing situations. This requires that the autonomy be able to comprehend, learn, and reason [7]. Several approaches, inspired by human cognition and biology, may be leveraged:

- **Fuzzy Logic** – Provides a framework for mapping one or more continuous state variables into useful categories for reasoning and decision making,
- **Neural Networks** – Solutions are learned by the network through a mathematical framework that represents knowledge through variable interconnection weights which are learned by the program during training with a large dataset of exemplar cases,
- **Genetic and Evolutionary Algorithms** – Inspired by evolutionary genetics, repeated iterations of simulations are used to narrow down across a large set of potential options to find optimal solutions.

Each of these techniques may be applied to the challenges of autonomy, in conjunction with appropriate system architectures, to acquire, encode, represent, store, process, and recall knowledge. Given the inherent complexity of the world, such approaches may be far more tenable than approaches that require extensive detailing of contingencies by experts.

Learning approaches are not without their challenges, however. They often require significant work to determine the relevant parameters and information to provide to the learning systems, and significant work to create appropriate systems architectures for learning and organizing outputs. In addition, learning systems will create new challenges:

- **Understandability** - The logic and behavior of such systems can be quite opaque to the airman, and often the system developers do not fully understand how the autonomy will behave, although there are techniques for deriving rules for characterizing the

predominant characteristics within the “black box” of the algorithms [52]. These insights may be incomplete, however, and not represent the full complexity of behavior the system may exhibit.

- **Validation** - Methods for successful verification and validation of autonomous systems developed through learning techniques will be critical if they are to be accepted in safety critical Air Force operations. Current techniques are generally insufficient for this challenge.
- **Standardization** - If the learning algorithms are able to continue to evolve in practice, significant new challenges will also be introduced in terms of consistency. Will the lessons learned by one system be transmitted to others so that they operate with some level of consistency? What will be the criteria for ensuring that appropriate lessons are learned in each case, and the generalizability of those lessons to other environments? Or will different systems behave differently, creating challenges for airmen to be able to understand and predict the operation of the autonomy so as to be able to interact with it appropriately and correctly calibrate trust? If learning algorithms are frozen after the training period, will they fall prey to the same issues of brittleness as other approaches, unable to learn and adapt in a changing world?

These challenges are not insignificant and must be addressed if learning systems are to be used to develop future autonomous systems.

5.3 Verification & Validation

Verification and validation of software for advanced systems for the Air Force is critical for ensuring that they are able to operate safely and consistently for their intended operations. New methods for verification and validation of autonomy software will be needed [53]. Traditional methods are based on requirements tracing and fail to address the complexities associated with autonomy software. There are simply too many possible states and combination of states to be able to exhaustively test each one, and understanding where the boundary conditions are will be difficult. The ability of the system to degrade gracefully and to support human-autonomy interaction will form an important aspect of successful autonomy implementation and will need to be expressly incorporated into validation testing.

5.4 Cyber

Autonomy has the potential to solve many challenges in Air Force operations, but it brings with it new vulnerabilities for cyber-attack, just like any other software system [54]. Because of the complexity of the autonomy, it may be fundamentally more difficult to detect inadvertent bugs or deliberately embedded malware. Greater contextual awareness can help with this challenge, as can self-health monitoring systems. In addition, methods for creating cyber resilience, including the ability to detect and repel, or work around, cyber-attacks on the autonomy, and on the rest of the system it is imbedded in, are critical. Cyber resilience is not a feature that can be added on to a system after it is developed. Rather, a consideration of cyber resilience must be foundational to the development approach for any autonomous system. A more complete discussion on the issues associated with mission assurance in a cyber-contested environment is provided in Cyber Vision 2025 [54].

6.0 Conclusions

Many Air Force systems will experience an evolution towards increasing levels of autonomy over the next several decades. These advances will only be successful in achieving their goals of increased range and speed of operations, increased mission capabilities, increased reliability, persistence and resilience, or reduced manning loads if they take careful consideration of the need for effective human-autonomy teaming. Past paradigms that created brittle automation, with limited capabilities and limited consideration of human operators, will be replaced by an explicit focus on synergistic human-autonomy teams. This new paradigm will directly support high levels of shared situation awareness between the airman and the autonomy, creating situationally relevant informed trust, ease of interaction and control, and the manageable workload levels needed for mission success. By focusing on airman-autonomy teaming, the Air Force will create successful systems that get the best benefits of autonomous software along with the innovation of empowered airmen.

7.0 References

1. U.S. Air Force, *America's Air Force: A call to the future*. 2014, Author: Washington, DC.
2. Lockheed Martin. *F-35 Lightning II*. 2015; Available from: <https://http://www.f35.com/about/life-cycle/software>.
3. Defense Science Board, *The role of autonomy in DoD systems*. 2012, Department of Defense, Washington, DC.
4. Air Force Scientific Advisory Board, *Operating next-generation remotely piloted aircraft for irregular warfare*. 2011, U.S. Air Force: Washington, DC.
5. U.S. Air Force, *Battlespace networking: ISR horizons future vision*. 2015, Deputy Chief of Staff for Intelligence Surveillance and Reconnaissance: Washington, DC.
6. U.S. Air Force, *Sensing as a service: ISR horizons future vision*. 2014, Deputy Chief of Staff for Intelligence Surveillance and Reconnaissance: Washington, DC.
7. Krogmann, U., *From automation to autonomy: Trends towards autonomous combat systems, in Advances in Vehicle Systems Concepts and Integration (RTO MP-44)*. 1999, NATO Research and Technology Organization: Neuilly Sur-Seine Cedex, France.
8. Schooley, L.C., et al., *High autonomy control of space resource processing plants*. IEEE Control Systems, 1993(June): p. 29-39.
9. U.S. Air Force, *Autonomy science and technology strategy*. 2013, Air Force Research Laboratory, Dayton, OH.
10. Wiener, E.L. and R.E. Curry, *Flight deck automation: Promises and problems*. Ergonomics, 1980. **23**(10): p. 995-1011.
11. Bainbridge, L., *Ironies of automation*. Automatica, 1983. **19**: p. 775-779.
12. Billings, C.E., *Toward human centered automation, in Flight deck automation: Promises and realities*, S.D. Norman and H.W. Orlady, Editors. 1988, NASA-Ames Research Center: Moffet Field, CA. p. 167-190.
13. Sheridan, T., *Telerobotics, automation and human supervisory control*. 1992, Cambridge, MA: MIT Press.
14. Endsley, M.R. and E.O. Kiris, *The out-of-the-loop performance problem and level of control in automation*. Human Factors, 1995. **37**(2): p. 381-394.

15. Sarter, N.B. and D.D. Woods, *"How in the world did I ever get into that mode": Mode error and awareness in supervisory control*. Human Factors, 1995. **37**(1): p. 5-19.
16. Parasuraman, R. and V. Riley, *Humans and automation: Use, misuse, disuse and abuse*. Human Factors, 1997. **39**(2): p. 230-253.
17. Wickens, C.D. and C. Kessel, *The effect of participatory mode and task workload on the detection of dynamic system failures*. IEEE Transactions on Systems, Man and Cybernetics, 1979. **SMC-9**(1): p. 24-34.
18. Young, L.R.A., *On adaptive manual control*. Ergonomics, 1969. **12**(4): p. 635-657.
19. McClumpha, A. and M. James, *Understanding automated aircraft*, in *Human performance in automated systems: Current research and trends*, M. Mouloua and R. Parasuraman, Editors. 1994, LEA: Hillsdale, NJ. p. 183-190.
20. Wickens, C.D., *Engineering psychology and human performance*. 1st ed. 1984, Columbus, Ohio: Charles E. Merrill Publishing Co.
21. Kidd, A.L. and M.B. Cooper, *Man machine interface issues in the construction and use of an expert system*. International Journal of Man Machine Studies, 1985. **22**: p. 91-102.
22. Pritchett, A.R. and R.J. Hansman. *Pilot non-conformance to alerting system commands*. in *Ninth International Symposium on Aviation Psychology*. 1997. Columbus, OH: Ohio State University.
23. Selcon, S.J. *Decision support in the cockpit: Probably a good thing?* in *Human Factors Society 34th Annual Meeting*. 1990. Santa Monica, CA: Human Factors Society.
24. Endsley, M.R. and D.G. Jones, *Designing for situation awareness: An approach to human-centered design*. 2nd ed. 2012, London: Taylor & Francis.
25. Layton, C., P.J. Smith, and C.E. McCoy, *Design of a cooperative problem-solving system for en-route flight planning: An empirical evaluation*. Human Factors, 1994. **36**(1): p. 94-119.
26. Pew, R.W. and A.S. Mavor, eds. *Human system integration in the system development process: A new look*. 2007, National Academic Press: Washington, DC.
27. Dzindolet, M.T., et al. *Misuse and disuse of automated aids*. in *Human Factors and Ergonomics Society 43rd Annual Meeting*. 1999. Santa Monica, CA: Human Factors and Ergonomics Society.
28. Lee, J.D. and N. Moray, *Trust, self-confidence, and operations' adaptation of automation*. International Journal of Human-Computer Studies, 1994. **40**: p. 153-184.
29. Lee, J.D. and K.A. See, *Trust in automation: Designing for appropriate reliance*. Human Factors, 2004. **46**(1): p. 50-80.
30. Masalonis, A.J. and R. Parasuraman. *Effects of situation-specific reliability on trust and usage of automated air traffic control decision aids*. in *Human Factors and Ergonomics Society 47th Annual Meeting*. 2003. Santa Monica, CA Human Factors and Ergonomics Society.
31. Merritt, S.M. and D.R. Ilgen, *Not all trust is created equal: Dispositional and history-based trust in human-automation interactions*. Human Factors, 2008. **50**(2): p. 194-210.
32. Muir, B.M., *Trust in automation: Part I. Theoretical issues in the study of trust and human intervention in automated systems*. Ergonomics, 1994. **37**(11): p. 1905-1922.
33. Riley, V., *A theory of operator reliance on automation*, in *Human performance in automated systems: Current research and trends*, M. Mouloua and R. Parasuraman, Editors. 1994, LEA: Hillsdale, NJ. p. 8-14.
34. Seong, Y. and A.M. Bisantz. *Judgment and trust in conjunction with automated decision aids: A theoretical model and empirical investigation*. in *Human Factors and Ergonomics Distribution A. Approved for public release; distribution is unlimited. Public Release Case No 2015-0267*

- Society 46th Annual Meeting*. 2002. Santa Monica, CA Human Factors and Ergonomics Society.
35. Wiegmann, D.A., A. Rich, and H. Zhang, *Automated diagnostic aids: The effects of aid reliability on users' trust and reliance*. *Theoretical Issues in Ergonomics Science*, 2001. **2**(4): p. 352-367.
 36. Miller, C.A., M. Pelican, and R. Goldman. *Tasking interfaces for flexible interaction with automation: keeping the operator in control*. in *International Conference on Intelligent user interfaces*. 1999. Redondo Beach, CA.
 37. Johnson, M., et al., *Seven cardinal virtues of human-machine teamwork: Examples from the DARPA robotic challenge*. *IEEE Intelligent Systems*, 2014. **November/December**: p. 74-80.
 38. Parasuraman, R., M. Mouloua, and R. Molloy, *Effects of adaptive task allocation on monitoring of automated systems*. *Human Factors*, 1996. **38**(4): p. 665-679.
 39. Kaber, D.B. and M.R. Endsley, *Out-of-the-loop performance problems and the use of intermediate levels of automation for improved control system functioning and safety*. *Process Safety Progress*, 1997. **16**(3): p. 126-131.
 40. Parasuraman, R., T.B. Sheridan, and C.D. Wickens, *A model of types and levels of human interaction with automation*. *IEEE Transactions on Systems, Man and Cybernetics*, 2000. **30**(3): p. 286-297.
 41. Kaber, D.B. and M.R. Endsley, *The Effects of Level of Automation and Adaptive Automation on Human Performance, Situation Awareness and Workload in a Dynamic Control Task*. *Theoretical Issues in Ergonomic Science*, 2004. **5**(2): p. 113-153.
 42. Rovira, E., K. McGarry, and R. Parasuraman, *Effects of imperfect automation on decision making in a simulated command and control task*. *Human Factors*, 2007. **49**(1): p. 76-87.
 43. Wright, M.C. and D.B. Kaber, *Effects of automation of information-processing functions on teamwork*. *Human Factors*, 2005. **47**(1): p. 50-66.
 44. Department of Defense, *DoD Directive 3000.09: Autonomy in Weapon Systems*. 2012, Author: Washington, DC.
 45. Miller, C. *From the Microsoft paperclip to the rotorcraft pilot's associate: Lessons learned from fielding adaptive automation systems*. in *Human Performance, Situation Awareness and Automation: User-Centered Design for the New Millennium Conference*. 2000. Savannah, GA: SA Technologies, Inc.
 46. Endsley, M.R. *Design and evaluation for situation awareness enhancement*. in *Human Factors Society 32nd Annual Meeting*. 1988. Santa Monica, CA: Human Factors Society.
 47. Endsley, M.R. and W.M. Jones, *A model of inter- and intrateam situation awareness: Implications for design, training and measurement*, in *New trends in cooperative activities: Understanding system dynamics in complex environments*, M. McNeese, E. Salas, and M. Endsley, Editors. 2001, Human Factors and Ergonomics Society: Santa Monica, CA. p. 46-67.
 48. U.S. Air Force Medical Service, *Human systems integration guidebook*. 2014.
 49. Boehm-Davis, D.A., F.T. Durso, and J.D. Lee, eds. *APA handbook of human system integration*. 2015, American Psychological Association: Washington, DC.
 50. Endsley, M.R., *Bringing cognitive engineering to the information fusion problem: Creating systems that understand situations*, in *Fusion 2011*. 2011, International Society of Information Fusion: Chicago, IL.
 51. Endsley, M.R., *Toward a theory of situation awareness in dynamic systems*. *Human Factors*, 1995. **37**(1): p. 32-64.

52. Huang, S.H. and M.R. Endsley, *Providing understanding of the behavior of feedforward neural networks*. IEEE Transactions: Systems, Man and Cybernetics, 1997. **27**(3): p. 465-474.
53. U.S. Air Force, *Global horizons (AF/ST TR 13-01)*. 2013, Office of the Chief Scientist: Washington, DC.
54. U.S. Air Force, *Cyber Vision 2025 (AF/ST TR 12-01)*. 2012, Office of the Chief Scientist: Washington, DC.