

A Cognitive Systems Engineering Approach to Developing HMI Requirements
for New Technologies

DISSERTATION

Presented in Partial Fulfillment of the Requirements for the Degree Doctor of Philosophy
in the Graduate School of The Ohio State University

By

LisaCarolynn Fern, M.S.

Graduate Program in Industrial and Systems Engineering

The Ohio State University

2016

Dissertation Committee:

David D. Woods, Advisor

Philip Smith

Jozef C. Raadschelders

Copyrighted by
Lisa Carolynn Fern
2016

Abstract

This document examines the challenges inherent in designing and regulating to support human-automation interaction for new technologies that will be deployed into complex systems. A key question for new technologies, is how work will be accomplished by the human and machine agents. This question has traditionally been framed as how functions should be allocated between humans and machines. Such framing misses the coordination and synchronization that is needed for the different human and machine roles in the system to accomplish their goals. Coordination and synchronization demands are driven by the underlying human-automation architecture of the new technology, which are typically not specified explicitly by the designers. The human machine interface (HMI) which is intended to facilitate human-machine interaction and cooperation, however, typically is defined explicitly and therefore serves as a proxy for human-automation cooperation requirements with respect to technical standards for technologies. Unfortunately, mismatches between the HMI and the coordination and synchronization demands of the underlying human-automation architecture, can lead to system breakdowns. A methodology is needed that both designers and regulators can utilize to evaluate the expected performance of a new technology given potential human-automation architectures. Three experiments were conducted to inform the minimum HMI requirements a detect and avoid system for

unmanned aircraft systems (UAS). The results of the experiments provided empirical input to specific minimum operational performance standards that UAS manufacturers will have to meet in order to operate UAS in the National Airspace System (NAS). These studies represent a success story for how to objectively and systematically evaluate prototype technologies as part of the process for developing regulatory requirements. They also provide an opportunity to reflect on the lessons learned from a recent research effort in order to improve the methodology for defining technology requirements for regulators in the future. The biggest shortcoming of the presented research program was the absence of the explicit definition, generation and analysis of potential human-automation architectures. Failure to execute this step in the research process resulted in less efficient evaluation of the candidate prototypes technologies in addition to the complete absence of different approaches to human-automation cooperation. For example, all of the prototype technologies that were evaluated in the research program assumed a human-automation architecture that relied on serial processing from the automation to the human. While this type of human-automation architecture is typical across many different technologies and in many different domains, it ignores different architectures where humans and automation work in parallel. Defining potential human-automation architectures a priori also allows regulators to develop scenarios that will stress the performance boundaries of the technology during the evaluation phase. The importance of adding this step of generating and evaluating candidate human-automation architectures prior to formal empirical evaluation is discussed.

Acknowledgments

First and foremost, I wish to thank my advisor, Prof. David Woods for the continuous support of my graduate studies and related research, for his patience, motivation, and immense knowledge, and also for the endless energy and enthusiasm for the field of Cognitive Systems Engineering that inspires not only his students but researchers from all backgrounds around the world.

Besides my advisor, I would like to thank the rest of my PhD committee: Prof. Philip Smith, Prof. Jozef Raadschelders, and Prof. James Todd, for their insightful comments and encouragement, but also for the hard questions which incited me to widen my research perspectives.

This dissertation would not have been possible without the support of my NASA family. I am forever grateful to Mr. Jay Shively for his mentorship and guidance in my career, and for allowing me the opportunity to execute the research detailed in this document. I thank Mrs. Debra Randall for her enthusiastic encouragement and motivation. I am indebted to Mr. Confesor Santiago and Mr. Conrad Rorie for the technical expertise and attention to detail that they brought to bear on this research effort. Thank you to the entire UAS Integration into the NAS Project for their continued support of my graduate and professional efforts.

Vita

1998.....Riverglen School
2005.....B.S. Psychology, University of Calgary
2008.....M.S. Industrial and Systems Engineering,
The Ohio State University
2008 to present.....Senior Research Engineer, San Jose State
University at NASA Ames Research Center

Publications

Fern, L., & Rorie, R. C. (2016). An examination of unmanned aircraft systems pilots' interaction with air traffic control while responding to detect and avoid conflicts. *Air Traffic Control Quarterly*.

Fields of Study

Major Field: Industrial and Systems Engineering

Table of Contents

Abstract.....	ii
Acknowledgments.....	iv
Vita.....	v
Publications.....	v
Fields of Study	v
Table of Contents.....	vi
List of Figures.....	xiii
List of Equations.....	xxi
List of Abbreviations	xxii
Chapter 1: Introduction.....	1
Human-Automation Interaction Design and Function Allocation	3
The Myth of Function Allocation by Substitution	8
Capturing Evolving Systems as Changes in Human-Automation Architectures.....	9
How Regulators Can Envision Future Operations: Learning from the Past	11
Chapter 2: Background.....	17
Detect and Avoid Systems for Unmanned Aircraft Systems (UAS)	17

DAA Event Timeline	20
DAA Roles and Responsibilities.....	23
Development of Minimum Operational Performance Standards for DAA Systems	37
Minimum Operational Performance Standards for a DAA HMI.....	39
Chapter 3: Research Overview	43
Metrics of Human Performance	44
Relevant Environment.....	52
Chapter 4: Experiment One	56
Overview.....	56
Experimental Design and HMI Configurations	58
Method	74
Measures.....	83
Results.....	83
Results Summary	92
Chapter 5: Experiment Two.....	93
Overview.....	93
Purpose.....	93
Experimental Design and HMI Configurations	95
Method	102

Measures.....	107
Results.....	107
Results Summary	114
Chapter 6: Experiment 3	116
Overview	116
Purpose	122
Experimental Design and HMI Configurations	130
Method	139
Measures.....	145
Results.....	146
Results Summary	155
Chapter 7: Research Results and Implications for HMI MOPS	157
Measured Response.....	157
Loss of Well Clear.....	178
Research Summary.....	190
Implications for the MOPS	195
Chapter 8: Discussion	205
Reflecting on the Methodology for Determining HMI Requirements for DAA Systems from a CSE Perspective.....	211

Improving the Methodology for Developing HMI Requirements for New Technologies	217
.....
Conclusions.....	227
References.....	230

List of Tables

Table 1. The Fitts (1951) MABA-MABA list.	4
Table 2. Scale of degrees of automation (Sheridan, 1987).....	6
Table 3. Stages of the pilot-DAA timeline.	45
Table 4. Alerting, information elements and display location for the four display configurations in Experiment 1.....	61
Table 5. The multi-level alerting structure used in Experiment 1. Only the advanced display configurations, Advanced Standalone and Advanced Integrated, included the “Predicted Collision Avoidance” alert level; it was not included in the basic display condition.	63
Table 6. The minimum information elements for the basic display condition.	66
Table 7. Measured response means by display configuration for Experiment 1.....	84
Table 8. Proportion of encounters that resulted in losses of well clear by information level and display location.	90
Table 9. Mean separation index across all display configurations.	91
Table 10. Alerting, information elements and display location for the four display configurations in Experiment 1.....	95

Table 11. The multi-level alerting structure used in Experiment 2 for all display configurations.	97
Table 12. The minimum information set for the basic display condition.	98
Table 13. Measured response means by display configuration for Experiment 2.	108
Table 14. Number, proportion, and mean separation index of losses of well clear by display condition.	113
Table 15. Proposed SC-228 MOPS alert structure.	120
Table 16. Multi-level alerting structure and separation criteria used in Experiment 3 for all display configurations.	133
Table 17. Comparison of the information elements used in Experiments 1, 2 and 3.	134
Table 18. Measured response means by display configuration for Experiment 3.	146
Table 19. Proportion of uploads with prior clearance approval by alert type and display.	148
Table 20. Total response time means and medians (in brackets) by display configuration for corrective DAA and DAA warning alerts.	152
Table 21. Number of each category of loss of well clear by display condition.	154
Table 22. Mean and median separation index across all display configurations.	155
Table 23. Measured response means by display configuration for Experiments 1, 2 and 3.	159
Table 24. Loss of well clear metrics for Experiment 1 and Experiment 2.	179
Table 25. Quick look overview of the significant differences and trends from the three reported experiments. The green checkmarks indicate displays that were found to have	

significantly better performance compared to another display(s), which are marked with red exes. The black checkmarks and exes denote the better and worse performing displays given strong trends in the results. 191

Table 26. *DAA Draft MOPS alerting thresholds. Reprinted with permission from: RTCA (2015a), DRAFT Minimum Operational Performance Standards (MOPS) for Unmanned Aircraft Systems (UAS) Detect and Avoid (DAA) Systems.*..... 199

Table 27. How the determine resolution maneuver function is accomplished for each of the three evaluated HMI configurations. 221

List of Figures

Figure 1. How well clear functions are currently accomplished for manned aircraft.....	19
Figure 2. DAA event timeline.....	21
Figure 3. How will well clear functions be accomplished when coordinated between UAS pilot on the ground and onboard and/or ground-based automation.	28
Figure 4. Envisioned DAA system design where well clear functions are accomplished solely by the onboard automation which is delegated authority to maneuver without pilot input.	29
Figure 5. DAA human-automation architecture in which the human accomplishes all three well clear functions.	30
Figure 6. True DAA human-automation architecture in which the machine is allocated all three well clear functions, but pilot remains ultimate responsibility for safety of aircraft.	33
Figure 7. Levels and interactions of DAA roles and responsibilities.	34
Figure 8. Relationship between stages of the pilot-DAA timeline and measured response metrics.....	46
Figure 9. DAA Timeline broken into time segments.....	47

Figure 10. The four DAA displays that were evaluated in Experiment 1: Standalone Basic (top left), Integrated Basic (top right), Standalone Advanced (bottom left), and Integrated Advanced (bottom right)..... 59

Figure 11. DAA timeline with the well clear and DAA alerting thresholds from Experiment 1 62

Figure 12. Visual depiction of DAA alert levels and thresholds 65

Figure 13. Screenshot of the Advanced Standalone DAA display condition during a self separation alert. The maneuver recommended by Autoresolver-AD is shown in the upper right-hand box (“Descend to: 14500 Feet”), the lateral trial planner is indicated by the magenta colored flight plan, the vertical trial planner (i.e., altitude tape) is located on the lower right side of the display, and the time to the predicted closest point of approach with the self separation alert is located in the bottom left. 69

Figure 14. Screenshot of the Advanced Integrated DAA display condition with an active predicted collision avoidance threat. The maneuver recommended by Autoresolver-AD is shown in the upper box (“Fly Heading 122”), the lateral trial planner is indicated by the arrow pointing to heading 122 off the nose of the ownship icon (center), the vertical trial planner is located on the far right side of the TSD, and the vertical situation display is shown in the lower quarter of the display. 72

Figure 15. UAS GCS. CSD (not displayed), bottom left; TSD, bottom center; out-the-window view, top center; health and status panel, bottom right..... 76

Figure 16. Vigilant Spirit Control Station tactical situation display (AFRL/RH).
Distribution A: Approved for public release; distribution unlimited, 3/18/2013; 88ABW-2013-1303. 77

Figure 17. Vigilant Spirit Control Station’s compass rose interface. 77

Figure 18. The Cockpit Situation Display (CSD) developed by the Flight Deck Display Research Laboratory. 79

Figure 19. Notification time by information level and display location..... 85

Figure 20. Initial response time by information level and display location..... 86

Figure 21. Initial edit time by display location and information level..... 87

Figure 22. Total edit time by information level..... 88

Figure 23. Total response time by information level..... 89

Figure 24. Proportion of encounters that became losses of well clear by information level (left) and display location (right). Reprinted with permission from Santiago & Muller (2015), *Pilot Evaluation of a UAS Detect-And-Avoid System’s Effectiveness in Remaining Well Clear*. 90

Figure 25. Separation index distribution by information level (left) and display location (right). The lower and upper boxes represent the bottom 25th and 75th percentiles, respectively, the median is indicated by the line intersection the boxes, and the minimum and maximum values are shown by the lower and upper whiskers..... 91

Figure 26. The four DAA displays that were evaluated in Experiment 2: Info Only (top left), Info + Vector (top right), Info + AR (bottom left), and Info + Vector + AR (bottom right)..... 96

Figure 27. The lateral (left) and vertical (right) vector planner tools. The white arrow on the lateral vector planner tool indicates that a small right turn would change current conflict from a self separation threat to preventive threat. The green-bordered altitude box indicates that a descent to 8000 ft MSL would result in a conflict free trajectory; the yellow-bordered altitudes indicate that altitudes from 8500 to 10,000 ft MSL would result in a self separation alert level..... 101

Figure 28. The auto-resolution text box providing directive guidance, here, “Fly heading 343.” 102

Figure 29. UAS GCS set up for Experiment 2; TSD, bottom center; out-the-window view, top center; health and status panel, bottom right and bottom left. 104

Figure 30. Initial edit times by display configuration..... 110

Figure 31. Mean total edit times by display configuration. 111

Figure 32. Mean total response times by display configuration. 112

Figure 33. Proportion of encounters that became losses of well clear by display condition. *Reprinted with permission from Santiago & Muller (2015), Pilot Evaluation of a UAS Detect-And-Avoid System’s Effectiveness in Remaining Well Clear.* 113

Figure 34. Separation index distribution by display configuration. The lower and upper boxes represent the bottom 25th and 75th percentiles, respectively, the median is indicated by the line intersection the boxes, and the minimum and maximum values are shown by the lower and upper whiskers. 114

Figure 35. *SARP well clear threshold recommendation. Reprinted with permission from: RTCA (2015c), DRAFT Minimum Operational Performance Standards (MOPS) for*

Unmanned Aircraft Systems (UAS) Detect and Avoid (DAA) Systems, Appendix C, Development of DAA Well Clear.....	117
Figure 36. <i>FAA whitepaper well clear threshold recommendation. Reprinted with permission from: RTCA (2015c), DRAFT Minimum Operational Performance Standards (MOPS) for Unmanned Aircraft Systems (UAS) Detect and Avoid (DAA) Systems, Appendix C, Development of DAA Well Clear</i>	118
Figure 37. TCAS II resolution advisory display implemented on a round-dial instantaneous vertical speed indicator, with the required vertical rate indicated by a green “band” and the vertical rate to be avoided indicated by a red “band”. Reprinted from <i>DOT (2011), Introduction to TCAS II Version 7.1.</i>	126
Figure 38. TCAS II resolution advisory display implemented on the primary flight display. The right shows the implementation of the red and green bands on the vertical speed tape. The left shows the pitch cue implementation (red pitch lines in the center of the attitude direction indicator). Reprinted from <i>DOT (2011), Introduction to TCAS II Version 7.1.</i>	127
Figure 39. Illustration of NASA Langley Research Center’s conflict prevention heading (right) and vertical speed (left) bands. Conflict prevention bands were generated by the prototype autonomous operations planner. Images reprinted from Mondoloni, Palmer & Wing (2002), Development of a Prototype Airborne Conflict Detection and Resolution Simulation Capability.	128

Figure 40. Illustration of NLR’s conflict heading bands generated by the predictive airborne separation assurance system. Reprinted from *Hoekstra, van Gent & Ruigrok (2002), Designing for safety: the ‘free flight’ air traffic management concept*..... 129

Figure 41. The four DAA displays that were evaluated in Experiment 3: Info Only (top left), No-Fly Bands (top right), Omni Bands (bottom left), and Vector Planning Tool (bottom right)..... 131

Figure 42. The No-Fly Bands suggestive maneuver guidance shows the pilot which headings and vertical speeds are to be avoided with yellow bands. The heading bands are shown on the inner range ring of the primary GCS display (left) and the vertical speed bands are down on the vertical speed indicator located on the right side of the primary GCS display (right). 136

Figure 43. The Omni Bands maneuver guidance shows pilots which heading and altitudes are predicted be safe versus those that are predicted to lead to a conflict. The heading bands are shown on the inner range ring of the primary GCS display (left) and the altitude bands are shown on altitude blocks on the right of the primary display (right). The green bands indicate a range of heading and altitude options (e.g., 8000 ft) that are not predicted to lead to a preventive DAA alert level or higher..... 137

Figure 44. The Vector Planner Tools maneuver guidance shows the predicted safety level of individual heading and altitude options when probed by the pilot. The heading options are probed by dragging an arrow along the inner range ring of the primary GCS display (right). The altitude options are probed by clicking on individual 500 ft blocks in the altitude table located on the far right side of the primary display (left). 139

Figure 45. Experimental ground control station. TSD, bottom-center; health and status panels, bottom-left and right; out-the-window view, top center.....	141
Figure 46. Proportion of uploads with prior clearance approval by alert type and display.	148
Figure 47. Initial response time by display configuration.	149
Figure 48. Mean initial and total edit times by display configuration.....	150
Figure 49. Mean total response times by display configuration.	151
Figure 50. Total response time by alert type at first appearance, by display.....	152
Figure 51. Proportion of encounters that were predicted to lose well clear that became actual losses of well clear by display condition.....	153
Figure 52. Separation index distribution by display configuration. The lower and upper boxes represent the bottom 25 th and 75 th percentiles, respectively, the median is indicated by the line intersection the boxes, and the minimum and maximum values are shown by the lower and upper whiskers.	155
Figure 53. Relationship between total response time for all four displays evaluated in Experiment 1, and the entire DAA timeline.	170
Figure 54. Relationship between total response time for all four displays evaluated in Experiment 2, and the entire DAA timeline.	170
Figure 55. Relationship between total response time for all four displays evaluated in Experiment 3, and the entire DAA timeline.	173
Figure 56. Total response time means by display configuration for Experiments 1, 2 and 3.....	174

Figure 57. Proportion of actual to predicted losses of well clear by display configuration for Experiments 1, 2 and 3.....	180
Figure 58. Distribution of separation index for all display configurations in Experiments 1, 2, and 3. The lower and upper boxes represent the bottom 25 th and 75 th percentiles, respectively, the median is indicated by the line intersection the boxes, and the minimum and maximum values are shown by the lower and upper whiskers.....	185
Figure 59. DAA timeline with expected pilot response time, ATC interaction time, and aircraft maneuver time.	197
Figure 60. How the well clear functions are accomplished according to the draft DAA MOPS (RTCA, 2015a).	210
Figure 61. Methodology implemented for the DAA HMI requirements development. ..	212
Figure 62. How the well clear functions are accomplished when coordinated between pilot and DAA system for UAS (left) compared to how they are accomplished by manned aircraft (right).....	213
Figure 63. Modified methodology for developing HMI requirements for new technologies.	218
Figure 64. Baseline human-automation architecture assumed by SC-228 DAA MOPS.	220

List of Equations

Equation 1. <i>Separation index</i>	52
---	----

List of Abbreviations

ATC	Air traffic control
CDTI	Cockpit display of traffic information
CPA	Closest point of approach
CSE	Cognitive systems engineering
DAA	Detect and avoid
DoD	Department of Defense
FAA	Federal Aviation Administration
GCS	Ground control station
HMI	Human machine interface
JADEM	Java Architecture for DAA Modeling and Extensibility
MABA-MABA	Machines are better at, men are better at
MOPS	Minimum operational performance standards
NAS	National Airspace System
NMAC	Near mid air collision
SC-228	(RTCA) Special Committee 228, MOPS for UAS DAA Systems
TCAS	Traffic alert and Collision Avoidance System
UAS	Unmanned aircraft system

Chapter 1: Introduction

As current systems evolve into new ones, advancements in technological capabilities provide new opportunities to introduce ever increasingly sophisticated forms of automation. During this evolution from old to new, the age-old question about what to automate will be addressed by engineers either implicitly, through feature creep, or explicitly, through formal design methods that include the allocation of functions between men and machines. The way that the design of automation capabilities, function allocation, and coordination between humans and machines is handled during the development phase will greatly impact the success of the future deployed system. More importantly, the degree to which the coordination of functions between the human and machine agents are supported by the human machine interface (HMI) will have direct consequences on overall system performance, and in safety critical domains, on system safety. While issues of function allocation and human-automation interaction have been a concern for human factors and cognitive systems engineers supporting practitioners for decades, rapid technology development in safety critical domains has forced regulators to also explore these issues.

This document will examine, from a cognitive systems engineering (CSE) perspective, a recent research effort aimed at defining the minimum HMI design standards for an emerging technology: detect and avoid (DAA) systems for unmanned

aircraft systems (UAS). DAA systems are a critical technological requirement to enable UAS operations in civil airspace by providing UAS pilots on the ground with capability of maintaining “well clear” from other aircraft. To fulfill this requirement, a DAA system will have to facilitate the performance of three well clear functions by the ground-based pilot: detect potential conflicts, determine a resolution maneuver, and execute the maneuver. A critical question for the regulators of this new technology is how these well clear functions will be accomplished by, and coordinated between, the pilot and DAA system. The answer to this question will help determine the final set of technical standards for DAA systems, especially those for the DAA HMI, which serves as the main conduit for human-machine interaction and cooperation.

In order to understand the challenges underlying the design and regulation of this new technology a brief look at how human-automation interaction design has traditionally been approached is provided. This is followed by an overview of how the three well clear functions could potentially be coordinated between the pilot and the DAA system, and what the allocation of well clear functions between humans and machines implies about the underlying human-automation architecture. Finally, a discussion about the current effort to establish minimum operational performance standards (MOPS) for DAA systems provides the necessary context for understanding a series of experiments that were designed specifically to inform the minimum information requirements for the DAA HMI. These three experiments are then evaluated, from the perspective of CSE, as a methodology that could be utilized by designers and regulators

to ensure effective human-automation coordination and safety assurance in the deployment of new technologies.

Human-Automation Interaction Design and Function Allocation

Human-automation interaction design for complex systems has traditionally been treated as a question about how to allocate functions or tasks between human and machine agents (no distinction is made here between “tasks” and “functions” or between “machines” and “automation”). In practice, function allocation is inherently difficult because most tasks are not independent, and there are often an infinite variety of ways that tasks can be allocated or shared between the human and machine agents. The field of human-automation interaction has a rich history of research on, and methods for, determining how best to allocate functions in complex human-machine systems.

A report authored by Fitts and his co-contributors (1951) was the first to propose a systematic method for the allocation of functions between humans and machines based on a list identifying the relative strengths and weaknesses of human and machines, commonly referred to as “Fitts list” (Table 1). This approach is now frequently called the “machines are better at, men are better at” or MABA-MABA approach, and variations of this list have shown up in the literature over the years. Under this seemingly elegant solution for designing complex human-machine systems, engineers need only to identify the functions for which machines are superior to humans and allocate those to the machine, and vice versa for humans. Birmingham and Taylor (1954) expanded this line of thinking by claiming that since machines are more reliable in performing intricate

computations, “man is best when doing least,” and that humans should be replaced by machines in complex systems whenever it is practically feasible.

Table 1. The Fitts (1951) MABA-MABA list.

Men are better at:	Machines are better at:
<ul style="list-style-type: none"> • Detecting small amounts of visual, auditory, or chemical energy • Perceiving patterns of light or sound • Improvising and using flexible procedures • Storing information for long periods of time, and recalling appropriate parts • Reasoning inductively • Exercising judgment 	<ul style="list-style-type: none"> • Responding quickly to control signals • Applying great force smoothly and precisely • Storing information briefly, erasing it completely • Reasoning deductively

This statement by Birmingham and Taylor highlights a major criticism to these early approaches of function allocation: that they rely on a comparison of machine capabilities to human capabilities. This type of human-machine comparison ignores core human competencies that contribute to overall system performance and results in the majority of system functions being reduced to a mathematical equation (Jordan, 1963). Not surprisingly, machines often come out in this comparison as being superior to humans. Ironically, even Taylor and Birmingham recognized that humans were infinitely adaptable and contribute to maximizing system output, yet still managed to reach the conclusion that humans should be designed out of the system whenever feasible! Thus, these MABA-MABA approaches to allocating functions between humans and machines typically ended with the conclusion that anything that could be automated should be

allocated to the machine. Anything left over (usually those functions too hard to automatize) should be left to the human, who is typically be expected to adapt to the machine functioning. Since automating everything was clearly unrealistic in practice, the MABA-MABA list saw limited utility for engineers that attempted to use them.

Recognizing the limitations of the MABA-MABA lists on their own in guiding function allocation, Sheridan (1997) laid out a number of considerations to aid in the task allocation process. While among these considerations included looking for the most obvious allocation of tasks (i.e., according to the MABA-MABA lists) as well as the “extreme” options (i.e., give all tasks to the machine or the human), Sheridan recommended looking at the levels or degrees of computerization or automation based on a hierarchy from his own supervisory control framework (Sheridan, 1987; Table 2). The supervisory control paradigm, he argued, is useful for “classifying human functions with respect to computer functions” (p. 92) and for defining a categorization for all of the different ways that a human can supervise a machine. (Note the emphasis on defining the human role *relative* to the machine.) Supervisory control continues to be a popular topic in human-automation interaction literature in a wide range of domains including surface transportation, aviation, spacecraft, ships and power plants. Research in supervisory control of robots (surface, underwater, and air) has received much attention and funding in the past several years as the Department of Defense (DoD) looks for ways to reduce the manpower burden and increase the operations of these relatively new technologies [see Chen, Barnes & Harper-Sciarini (2011) for a recent review of research on supervisory control of robots].

Table 2. Scale of degrees of automation (Sheridan, 1987)

1.	The computer offers no assistance: the human must do it all
2.	The computer offers a complete set of action alternatives, and
3.	narrows the selection down to a few, or
4.	suggests one alternative, and
5.	executes that suggestion of the human approves, or
6.	allows the human a restricted time to veto before automatic execution, or
7.	executes automatically, then necessarily informs the human, or
8.	informs the human only if asked, or
9.	informs the human only if it, the computer, decides to
10.	The computer decides everything and acts autonomously ignoring the human

More recently, research on human-automation interaction has begun to shift toward adaptable and adaptive automation. These newer human-automation interaction paradigms attempt to overcome the limited view of fixed human and machine strengths and weaknesses that plague both the MABA-MABA and the levels or degrees of automation approaches by implementing multiple levels or degrees of automation within a system that can change dynamically. Adaptable automation typically refers to the ability of the human to adapt the level of machine automation according to the human's needs, current level of workload, etc. For an example of recent research on adaptable automation, see the flexible levels of interaction approach to multiple unmanned aerial vehicle control by Draper, Miller, Calhoun, Ruff, Hammel and Barry (2013). Adaptive automation, on the other hand, is automation that is self-adapting based on real-time information of the human's physical and cognitive demands, such as workload or mission

performance (see Hancock, Jagacinski, Parasuraman, Wickens, Wilson, & Kaber, 2013 for a recent discussion of adaptive automation research).

While adaptable and adaptive automation reduce the rigidity of applying single levels of automation to various system functions such as those prescribed by Parasuraman, Sheridan & Wickens (2000), they still largely subscribe to machine-centric principles of function allocation – that is, how can we maximize what the automation does? In many ways both adaptable and adaptive automation rely on going back to the MABA-MABA and levels of automation lists with the added challenge of finding elegant ways to vary what the machine is doing without disrupting the performance of the human. Any function allocation based on machine competencies assumes that the task environment is predictable and controllable and that the nature of the task demands consistent performance (Jordan, 1963). However, to the extent that the task environment is dynamic and unpredictable, humans are necessary for coping with contingencies, a function that can only be carried out effectively when human-machine cooperation is properly supported; another reason that these traditional function allocation techniques fail is that they ignore the coordination and synchronization of activities that has to be accomplished between the human and machine agents. If complex human-machine systems are designed based solely on the allocation of functions without taking into consideration the human-machine cooperation necessary for optimal system functioning, at best the system will fail to meet performance expectations, at worst clumsy designs will lead to system failure with potentially catastrophic results.

The Myth of Function Allocation by Substitution

Another limitation of function allocation according to MABA-MABA and levels of automation (whether implemented in a rigid or flexible manner) is that the allocation of functions to machines is frequently framed as simply substituting a function that was previously carried out by a human operator – a practice termed “function allocation by substitution” (Hollnagel, 1999). This practice reflects the wide held and persistent belief in the ‘substitution myth’: that a machine can substitute a human function within the system while preserving the current functioning of that system. This myth is perpetuated in some definitions of automation, such as Parasuraman et al.’s (2000) definition that “automation refers to the full or partial replacement of a function previously carried out by the human operator” (p. 287).

This understanding of new automation as changing only a single role in the system ignores the reality that the introduction of new technology changes multiple roles as well as the interactions between roles. This is why the substitution myth is an oversimplification fallacy; observations of real work in complex systems reveal that the addition of new technology changes the roles of humans and introduces new ways of doing things, new demands on their cognitive activities, new ways of interacting, and new opportunities for failures (e.g., Roth et al., 1987, Sarter & Woods, 1994, Sarter, Woods & Billing, 1997). When a function, previously carried out by a human, is allocated to a machine, the function is not seamlessly absorbed into the system and eliminated from the human’s list of responsibilities. Rather, it creates new functions for the human, often in the form of meeting new coordination and synchronization demands

with the machine. These new functions include, but are not limited to, data entry and acquisition, directing attention, detecting potential problems, and deciding how and when to intervene in the automation's activity. All of these functions are required to keep the system functioning properly and all of which become challenged by poor human-automation cooperation, a detail that is often overlooked by designers and managers alike (Dekker & Woods, 2002, Woods & Hollnagel, 2006). The result is that introduction of new automation frequently falls short of fulfilling the promises (e.g., reduced workload, lower workforce costs, increased reliability, fewer errors, etc.) that the substitution of humans by machines was supposed to bring.

Introducing new technology into highly regulated domains, where safety is a chief concern of regulators, provides a unique challenge. Here, a key question is how to ensure that safety is preserved or improved in the existing system as roles adapt and the way that work is done changes. While the consequences of these changes can be predicted, many are unanticipated. Despite promises for better performance, the past shows us that the implementation of new forms of automation has also resulted in new paths to failure that were previously thought impossible. The accident at Three Mile Island (Kemeny, 1979; Perrow, 1984) provides a stark example of system outcomes that were incomprehensible to designers, operators, managers, and other stakeholders.

Capturing Evolving Systems as Changes in Human-Automation Architectures

In order for automation capabilities to be exploited while also preserving, or even exceeding, the current level of safety within the target domain, an understanding of the human and automation roles is needed; the human machine roles and responsibilities, and

coordination between roles, define the underlying human-automation architecture. Specifically, we need to understand how the roles and responsibilities of humans and automation shift with changing architectures in order to design the new system such that these roles are properly supported, a critical requirement if the system is to perform at a desired level of efficiency and safety. As more and more forms of automation are introduced, the change in roles and responsibilities needs to be traced so that both designers and regulators can anticipate new challenges in the coordination and synchronization of the human and machine roles.

For regulators, this process of tracing the various changes in human-automation architectures is especially critical as they anticipate the reverberations caused by introduction of the new technology to the overall system: impacts to operational goals, the need to change regulatory policy, and most critically, new opportunities for system failure. The regulatory body is responsible for ensuring that the integration of any new technology is as seamless as possible and does not adversely impact the safety of, or efficiency of operations within, the larger domain system. In order to do this, the regulators must have sufficient confidence in the reliability and performance of the new technology so that they can make considerations about how its introduction to the domain might create new challenges and risks. As part of this, they will also need to understand how the new technology will interact with other system components. Essentially, regulators must envision the effect of the future operations of the new technology in the future domain environment, a task that, like allocating functions to machines, is also susceptible to oversimplification. Engineers, customers, and managers frequently struggle

to accurately predict the real consequences of introducing new technology into the larger system, typically falling prey to envisioning *only* those consequences that they have predicted a priori (e.g., lower workload, high accuracy; Roesler, Feil & Woods, 2001, Dekker & Woods, 2002). Further, advocates for the new technology can become overly confident in their predictions, a miscalibration of how well they know what they know. This same problem of only anticipating and preparing for consequences that are predicted prior to implementation of a new technology is likely to befall regulators of the system, setting them up to be unprepared for unpredictable outcomes and potentially making less than desirable regulatory decisions.

How Regulators Can Envision Future Operations: Learning from the Past

In order to overcome the oversimplification shortfalls inherent in envisioning future operations, a methodology is needed that will support better anticipation of the reverberations and side effects of change. In order to be successful, this methodology must occur earlier in the design or standards development process than is typically done today. Human factors and CSE have traditionally been utilized to study the impact of new technologies near the tail end of the design and development processes, often in the form of verification and validation activities or as part of a late-step tuning process, leaving little or no opportunity to make significant changes to the system. As Woods and Dekker observed “it is found, over and over again, that providing such empirical testing roles provides too little information, too late in the design process, at too great a cost” (p. 275). This problem is exacerbated in regulatory environments as design decisions become accepted standards for all manufacturers of that technology, and design problems, such as

those that result in poor human-automation cooperation, are propagated through the system. Instead, CSE needs to be applied early in the design and standards development process, collecting empirical human performance data in a relevant context to ensure successful human-automation cooperation and better predictability of the effect of introducing the new technology into the system. It is much easier to address potential side effects or consequences of new technology earlier on in the design and standards development processes, when the field is already undergoing a period of change and commitments to specific designs and/or requirements have not yet been made.

The research detailed in this document utilizes a specific domain, aviation, to extract lessons learned and to improve upon a methodology that can be applied to the process of developing regulatory requirements for new technologies. Aviation is a domain that has a history of numerous changes to its human-automation architectures for which the regulatory guidelines have not kept pace, especially with respect to cockpit design. In aviation, like in many other domains, “brand new” technologies typically do not appear. Instead, new technologies tend to be more machine-capable generations of previous technologies and they are introduced as a simple replacement of the previous generation. In the past several decades, new technological capabilities have allowed for the increasing use of automation for both tactical and strategic control and navigation functions, as well as the display and management of information.

With each new advance in cockpit automation, the ensuing reverberations have included an increase in safety in some areas but also a decrease in safety in other areas as new risks have appeared. For instance, the implementation of these automation

capabilities has resulted in reduced workload associated with manual piloting tasks and increased airworthiness, reliability and operational efficiency of aircraft. However, new demands on pilots and new risks to safety have also resulted. Pilot tasks have shifted from largely inner-control loop tasks that demand a high level of psychomotor skill and continuous attention to more supervisory and management tasks, which while not entailing continuous attention, place new requirements on attentional, knowledge and information-processing resources (Billings, 1997).

A major challenge for pilots of modern aircraft has been adjusting to the new role of monitoring the automation functioning – understanding what it is currently doing and what it will do next, and intervening when problems arise (Sarter & Woods, 1994, Sarter, Woods & Billing, 1997). A recent investigation into the use of highly automated flight path management systems found a number of challenges for pilots managing these complex systems which in turn result in new risks to aviation safety (FAA, 2013b). Some examples of identified vulnerabilities for pilots include: mode confusion errors; overreliance on, and reluctance to intervene in, automation functioning; degraded manual flying skills including the ability to prevent, recognize and recover from upset conditions, stalls and unusual attitudes; and, insufficient knowledge of system functioning or flightcrew procedures. Lest these findings lead one to fall into the trap of believing that humans are the weak link in the system, the working group also found that “pilots mitigate safety and operational risks on a frequent basis, and the aviation system is designed to rely on that mitigation” (p.29). Undoubtedly, these issues point to an absence of sufficient support for human-automation coordination, an issue that could potentially

be addressed in the design and standards development processes. Perhaps not surprisingly, the same report cited issues with flight deck design and standardization, as well as the knowledge and skills of the regulators and the regulatory process. Many of these concerns surrounded the human factors skills, knowledge, expertise, training, and/or application of both the designers and the regulators.

These clear shortcomings of the existing cockpit automation technologies and the inability of regulators to keep pace with evolving automation capabilities, provide many areas of opportunity to improve the human-automation cooperation in various aviation systems. One such opportunity is to address the challenges faced by both designers and regulators in supporting effective human-automation coordination by utilizing a CSE approach to developing performance standards. Although this approach is much harder to implement with manned aircraft given the complexity and maturity of both the current systems and standards, the opportunity exists for a new aviation technology, unmanned aircraft systems (UAS).

The purpose of this document is to present a recent research program that was implemented in order to develop the minimum HMI requirements for UAS DAA system, a newly transitioning technology for which there are no existing standards. The research program is then evaluated from the perspective of CSE, with the goal of improving upon that methodology so that both regulators and designers can implement it to systematically assess and ensure the safety performance of human-automation architectures for future systems. While this research program was devised, not with CSE principles specifically in mind, but rather to address near term goals to provide input on the design of this new

system, it provides a broader opportunity to look at how human-automation interaction in new technologies can be better supported by more explicitly defining and analyzing candidate human-automation architectures. The research program presented here exhibits the typical situation faced by both designers and CSE professionals whereby the human-automation architecture is implicitly accepted and can only be narrowly tuned within defined boundaries. An improvement to the methodology, in turn, focuses on identifying potential variations of the human-automation architecture a priori and using those to guide the systematic testing of candidate prototypes.

Two key conditions need to be met for any methodology to aid in the transition of new technologies to future, highly regulated worlds. First, it will need uncover changes in human and automation roles and responsibilities so that impacts of the new technology on the existing system can be identified. Of critical importance is understanding new coordination and synchronization requirements, which come with new opportunities for breakdowns and failures. Analyzing these changes allows stakeholders to more accurately predict or anticipate the reverberations of technological change in a mature system *before* it is deployed in order to steer design toward better human-automation cooperation (Woods & Dekker, 2000). Second, the methodology will need to provide the ability to make predictions based on empirical results that have been abstracted from observations embedded within a relevant context (Roesler, Feil & Woods, 2001).

The research program detailed in later chapters provides a first step toward such a methodology. It utilized the systematic evaluation of various system configurations or prototypes embedded in a future envisioned environment in order to provide empirical

data from which regulators can make future predictions about operational performance. What it lacked, however, was the formal definition and analysis of the underlying human-automation architecture of each prototype system. By missing this step in the overall methodology, the research program missed a significant opportunity to identify the critical coordination and synchronization demands within the system a priori and to allow those demands to drive both the candidate HMI designs as well as the scenarios used in system evaluation. Despite these limitations, the research presented here still provides a radical departure from the status quo process for developing HMI requirements. Instead of relying heavily on subject matter expertise with limited empirical study late in the development phase, this research program was employed early and relied heavily on the analysis of objective data collected through systematic investigation. Thus it provides an excellent starting point from which to improve on the development a new methodology for designing and regulating future systems.

The next chapter provides the necessary context for understanding and interpreting the research detailed in subsequent chapters. While the research described in this document focuses on a specific system (UAS) embedded within a specific context (developing HMI requirements today for near term civil airspace operations), the lessons learned from this process can be extracted to a larger understanding of how to better design or regulate any future human-automation architecture in any domain.

Chapter 2: Background

Detect and Avoid Systems for Unmanned Aircraft Systems (UAS)

The introduction of UAS to the National Airspace System (NAS) provides a rich opportunity to examine the nature of evolving systems in a highly regulated environment. Today's UAS represents a transition from the traditional, manned aviation human-automation architecture where the pilot sits in the cockpit located on board the aircraft to a new human-automation architecture where the aircraft is remotely piloted from a ground control station (GCS); a command and control link transmits control and navigation commands from the GCS to the aircraft which in turn downlinks onboard sensor and status information. The immaturity of today's UAS technologies coupled with the change in human-automation architecture provides a rich opportunity for the aviation and technology industries to exploit automation capabilities as various users envision future worlds where autonomous aircraft will operate. As with the previous promises in aviation and other domains, the anticipated benefits of both current and future UAS technology is lower costs (due to reduced manpower) and better performance (more reliable and consistent than fallible human pilots). The Department of Defense (DoD) is the biggest proponent of employing highly automated or autonomous UAS as they face constant pressures to reduce operating and manpower budgets, as well as reduce risk to life (Dahmn, 2010).

Despite the push by DoD and other users who want to see more broad operational and autonomous application of UAS, the Federal Aviation Administration (FAA) has the sole authority and responsibility to ensure the safety, and minimize the system-wide impact, of operations of UAS in the NAS. Arguably one of the biggest changes associated with the human-automation architecture of UAS compared to traditionally piloted aircraft, and one of the two biggest technological challenges currently being addressed by the FAA (FAA, 2013a), is a shift in the roles and responsibilities for complying with the “see and avoid” requirements specified by the General Operating and Flight Rules (2004) in section 91.113 of the Title 14 Code of Federal Regulations (14CFR), for a pilot of a manned aircraft to remain “well clear” of other aircraft. (The second biggest technological concern of the FAA with respect to UAS is defining and ensuring acceptable performance of the command and control technology that links the GCS and the aircraft). This rule is one of many regulations in the FAA’s layered approach to aviation safety.

In today’s manned aircraft, pilots have sole responsibility for seeing and avoiding other aircraft under this rule. Even when an aircraft is under ATC separation services, the pilot in command maintains full responsibility for the safety of the aircraft and for maintaining a well clear separation distance from other aircraft in order to minimize collision risk. In order to maintain well clear, the pilot must carry out three major functions: 1) detect potential conflicts, 2) determine a safe conflict resolution maneuver, and 3) execute the safe conflict resolution maneuver. In addition to the functions required specifically for maintaining well clear, a pilot must also coordinate maneuvers with ATC

whenever practicable to maintain interoperability with the rest of the users of the airspace. This function is more of a general pilot function, but potentially has critical impacts to the larger system especially during well clear events where a pilot is required to make a maneuver. Figure 1 provides a simplified depiction of how the three well clear functions are accomplished in manned aircraft. To carry out these functions, the pilot relies upon visual, out-the-window acquisition of other aircraft to which s/he applies a subjective judgment of well clear to in determining whether another aircraft is a potential conflict (i.e., will penetrate the subjectively assessed well clear separation distance), and if it is, how to maneuver safely to avoid it.




Figure 1. How well clear functions are currently accomplished for manned aircraft.

With UAS, the pilot seated at a GCS is no longer capable of visually detecting other aircraft from the cockpit. Therefore, in order to be able to fulfill the regulatory

requirement to see and avoid other aircraft, some type of technology is needed to assist, or even replace, the pilot in maintaining well clear. Within the UAS community, the function of maintaining well clear via electronic detection of aircraft is now commonly referred to as “self-separation”¹ or “detect and avoid” (versus the manned aviation’s “see and avoid”) and the technology that supports this function are referred to as a DAA system. (From a joint cognitive systems point of view, the “DAA system” necessarily includes the pilot, however, from the perspective of designers and regulators it typically only refers to the machine components. To avoid confusion to the reader, in this document “DAA system” will be used only to reference the machine components.)

DAA Event Timeline

Figure 2 provides a general overview of a DAA, or loss of well clear, event. The onset of the event occurs with the detection of a potential loss of well clear conflict. A potential conflict may occur simply due to the normal passing of aircraft in the airspace. In these cases, aircraft are usually not on a collision course and a maneuver may not be required to avoid a loss of well clear. Whether automation or the human does the detection, there will be some error in the estimation of a loss of well clear or collision risk. In other cases, a potential conflict may be the result of both anticipated and unanticipated events in the system, such as: violation of ATC separation criteria in

¹ The term “self-separation,” while still commonly in use amongst the UAS community, both in the United States and internationally, has been rejected as an official term by the FAA to describe the function of a UAS remaining well clear of other aircraft because of legacy connotations of “self-separation” in the air traffic management domain, which implies full delegation of separation authority and responsibility to the pilot. UAS are envisioned to operate under instrument flight rules, which mean they will operate under air traffic control (ATC) separation rules.

controlled airspace (by ATC or pilots), aircraft not maintaining right of way rules, late detection by the DAA system, and pilots of manned aircraft not visually acquiring nearby aircraft and failing to maintain their own well clear.

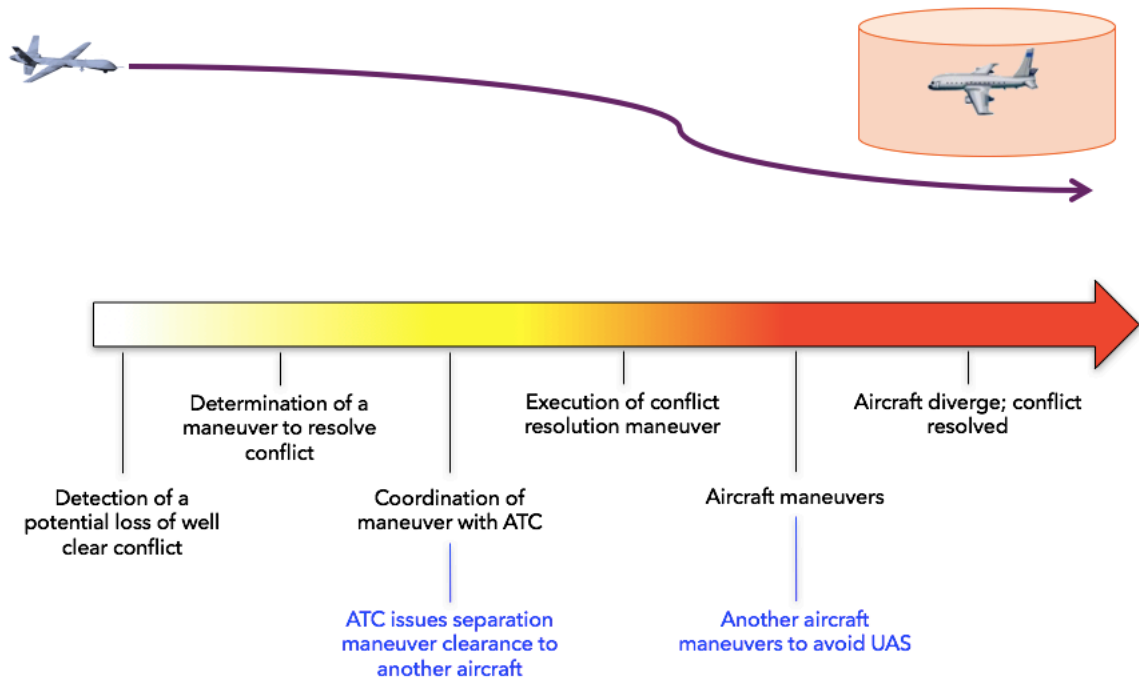


Figure 2. DAA event timeline.

Following the detection of a potential conflict, the DAA system and/or pilot must then determine a maneuver that will resolve the conflict. That maneuver then has to be coordinated with ATC before it is executed, except in exceptional cases. Finally, once the maneuver is executed, the aircraft begins to maneuver away from the conflicting aircraft until the aircraft are on diverging trajectories and the conflict is resolved. In this simple overview of the DAA even timeline, there are also two cascading events resulting from

the coordination and execution of a resolution maneuver (marked in blue text in Figure 2). These are the maneuvers by other aircraft, either initiated by the pilots of those aircraft or ATC, resulting from the well clear maneuver by the UAS.

The arrow in Figure 2 represents the temporal aspect of a DAA event as the aircraft approach the closest point of approach (CPA). The timeline for a DAA event is necessarily constrained, since they only occur when aircraft have gotten too close to each other, relative to a separation standard. In addition, system constraints influence how far out a conflict can be detected, while operational constraints influence when it is acceptable for pilots to maneuver to maintain well clear – which in turn affects when a potential conflict may be “flagged” as a threat by the system. A tradeoff has to be made between providing a safe amount of time for a DAA system to execute a well clear maneuver, and both technological capabilities and operational suitability. The left side of the timeline represents that maximum time allowed by the operational context and system capabilities to detect a potential threat. Here the temporal urgency for executing the stages of the DAA event is low. As you move further right, the temporal urgency increases, and there is less time for the DAA system to avoid a loss of well clear.

Overall, the DAA event timeline is largely driven by the three well clear functions (detect, determine, execute); a failure in execution of any of these could result in a loss of well clear and potential collision. Coordination with ATC is also an important function because of the potential effects on the rest of the airspace system. However, elimination of the function is highly unlikely to be catastrophic, while on the other hand, due to the temporal nature of a DAA event, taking time to execute this function could delay

maneuvering and therefore increase collision risk – a critical tradeoff decision that a pilot has to make in real time.

DAA Roles and Responsibilities

The concept of DAA (both the function and technology) can be understood at three levels or scales within the aviation domain: tactical, operational, and regulatory. At each level are individuals with specific roles relevant to DAA.

Regulatory

At the regulatory level, the FAA has been granted the authority to regulate the safe use of the U.S. airspace (Federal Aviation Act, 1958). Another key role of the FAA is to ensure the efficiency of the airspace so that it is useable by the public – safety decisions that effectively shut down the airspace or even severely limit the operations available, will not be acceptable to the U.S. public. Thus, their primary role with respect to DAA is to ensure that UAS operations do not negatively impact the safety or operational efficiency of U.S. airspace through the development of DAA technology that enables UAS to remain well clear of other aircraft so as to reduce the risk of a near mid air collision (NMAC). As part of their role, the FAA will issue technical standards for manufacturers of DAA systems that are consistent with those responsibilities. At a broader level, the FAA will also develop operational rules and pilot certification requirements for UAS to operate in U.S. civil airspace, however, the DAA technical standards are a critical prerequisite before those rules can be written. In addition, the FAA will monitor the safety and operational performance of future deployed UAS systems to ensure that there is no degradation in either of these areas to the NAS. Where

UAS operations are found to degrade airspace efficiency or safety, they will modify the regulations to address those degradations.

Safety and operational efficiency are not the FAA's only areas of obligation; the powers and responsibilities granted by the Federal Aviation Act of 1958 also include the promotion, encouragement, and development of civil aeronautics, and meeting the requirements of national defense. These additional responsibilities are contextually important with respect to DAA: both the UAS industry and the DoD have a vested interest in the content of future UAS regulations. In fact, both groups were strong advocates pushing for the FAA to establish policy that would allow for expanded operations of UAS in the NAS, which led to the current effort to develop technical requirements for DAA systems. The FAA typically employs the use of advisory committees to draft technical standards that will later be invoked through rulemaking. Advisory committees include, and require the consensus of, other industry and government organizations that choose to participate. In the case of DAA standards, representatives from UAS and avionics manufacturers, pilot lobby groups, DoD (primarily the Air Force), and NASA have all shown up to participate in the advisory committee process. So while the FAA is primarily concerned with the operational and safety impacts of UAS flying in the NAS when developing the DAA standards, the participation of other groups, especially those from industry, results in competing opinions about what the minimum technical requirements for DAA equipment should be. These opinions may be about strictly technical issues, for example, the best method for building or assessing hardware and software or even determining what is technological

feasible (there's no point in specifying requirements for a system that can't be built). However, beliefs about which requirements should be employed may also be based on non-technical issues such as economic concerns by manufacturers of the cost of building certain technologies, which cannot immediately be dismissed by the regulators since the final set of standards will have to be agreed to by those same manufacturers. Thus, decisions will have to be made by group consensus and will represent tradeoffs between competing wants within the advisory committee.

Operational

At the operational level, air traffic controllers are responsible for strategic operational efficiency and safety. Here, the role of controllers is to try to maximize throughput of the NAS given operational constraints (e.g., weather, demand, etc.) while also contributing to safety by ensuring that aircraft maintain minimum separation criteria according to the operational rules established by the FAA. With respect to DAA, the role of air traffic control (ATC) could be impacted by the deployment of a DAA system that does not interoperate smoothly with the current system. Since operational rules apply not only to how ATC operates, but also to how pilots operate, controllers have expectations about how pilots behave in the NAS. The ability of ATC to predict pilot behavior allows them to more efficiently and safely manage the airspace. These expectations center on what routes pilots fly and how they coordinate initial routes and changes to their route, as well as how pilots will interpret and follow clearances issued by ATC. UAS are expected to fly instrument flight rules flight plans under ATC separation services. This means that they will be expected to coordinate and receive ATC clearance approval for any

deviations off of their filed flight plan except in the case of an emergency. If the DAA system is built such that UAS pilots make frequent deviations from their flight plan, the frequent resulting maneuvers could be disruptive to the airspace environment. This is especially true if well clear maneuvering by a UAS causes other aircraft in the system to maneuver (on their own assessment of well clear) or be maneuvered by ATC (based on ATC separation criteria), highlighting the interaction of DAA roles beyond just the pilot and GCS. This problem could be further compounded if UAS pilots are not coordinating their maneuvers with ATC, which would cause a reduction in predictability of how UAS behave. If ATC does not feel they can reasonably predict the behavior of UAS pilots operating in their airspace, they may start employing unofficial procedures such as increased separation minima for UAS in order to ensure safety. This type of task tailoring would potentially have a negative impact on the missions of both manned and unmanned aircraft as well as affecting overall operational efficiency.

Tactical

At the tactical level pilots are responsible for fulfilling their operational missions while complying with airspace regulations and maintaining the safety of their aircraft. The change in human-automation architecture for UAS means that pilots now control the aircraft via the GCS that sends command and control information to the aircraft and receives back onboard sensor information. Maintaining well clear with the assistance of both onboard and ground-based DAA technology will be a critical responsibility that UAS pilots will need to perform. While pilots of manned aircraft execute this function primarily through visual acquisition of aircraft and input directly to the onboard control

and navigation interfaces, which in some cases may be aided by machine capabilities, UAS pilots will be responsible for executing this function in conjunction with DAA machine capabilities via the ground-based control station. Pilots will still be responsible for also following operational rules when executing this function, specifically, as outlined in the previous section, pilots will be expected to coordinate maneuvers off of their filed flight plan with ATC prior to executing well clear maneuvers, except in the case where an immediate safety of flight concern exists. A DAA system should support the pilot in meeting both of these responsibilities related to maintaining well clear in the operational environment.

At a minimum, the DAA system will require a suite of onboard surveillance equipment that is capable of electronically detecting nearby aircraft. The removal of the pilot from a cockpit on board the aircraft means that it is physically impossible for the pilot to do this without the aid of machines. In addition, the DAA system will notionally consist of other hardware and software components that together provide the necessary information to support the ability of the UAS pilot to maintain well clear. These potential components include: data fusion and/or correlation logic, threat detection and resolution logic, the display of traffic and resolution maneuver guidance information, and control and navigation interfaces to execute a resolution maneuver. However, there are a number of possible configurations of the specific machine components, and the distribution of those components on board the aircraft or at the GCS is dependent on how the well clear functions are allocated or shared between the pilot and machines. A key question for both designers and regulators of this DAA system is the see and avoid functions to remain

well clear will be coordinated between, and accomplished by, the pilot and DAA capabilities. In other words, what human-automation architecture will be implemented for DAA systems for UAS (Figure 3). Note that coordination with ATC is still considered a pilot function to be carried out by the pilot in command in conjunction with the well clear functions.

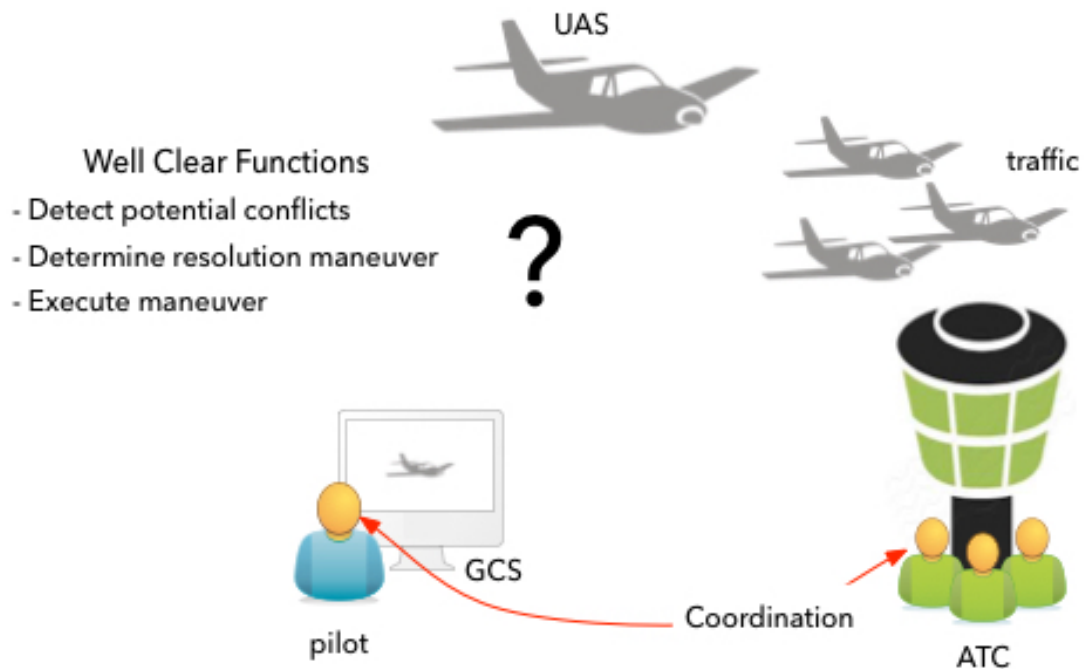


Figure 3. How will well clear functions be accomplished when coordinated between UAS pilot on the ground and onboard and/or ground-based automation.

On one end of the spectrum, the three well clear functions could be accomplished entirely by the machine, leaving the pilot entirely “out-of-the-loop” (Figure 4). In this human-automation architecture all of the DAA hardware and software components that carry out the determine and execution functions would likely be located on board the

aircraft. The pilot on the ground would be responsible for monitoring the aircraft and coordinating with ATC based on the aircraft's behavior (however, in many cases, coordination with ATC is disregarded altogether). This design assumes that not only are the well clear functions accomplished by the onboard DAA system, but that *authority* to maneuver is also delegated to the DAA system. The authority here is delegated by the designers who intentionally remove the links between the well clear functions and pilots under the assumption that the technology will perform as expected and only residual monitoring by the pilot is necessary. Unfortunately, the HMI is typically not designed to properly support the necessary monitoring and possible intervention or redirection of the automation's activities, or the coordination of activities with ATC.

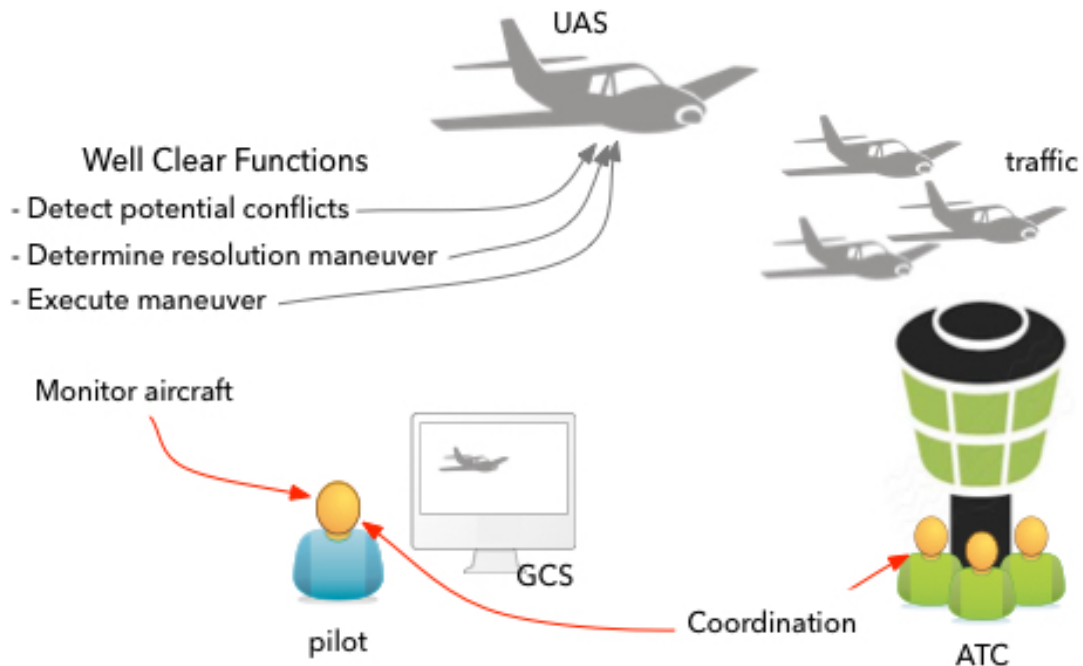


Figure 4. Envisioned DAA system design where well clear functions are accomplished solely by the onboard automation which is delegated authority to maneuver without pilot input.

On the other end of the spectrum, the pilot could accomplish determining and executing resolution maneuvers without machine assistance (Figure 5). This is the most fully “manual” DAA architecture possible. Note that this human-automation architecture, and any for which the human shares or is fully responsible for either of the other two well clear functions, requires pilots to accomplish the function of detecting potential conflicts with the assistance onboard surveillance equipment. In this DAA design, the onboard equipment simply passes surveillance information to the pilot via the DAA HMI. Note that when the pilot accomplishes the DAA function of executing a resolution maneuver, it is essentially by instructing the flight control automation to execute the selected maneuver.

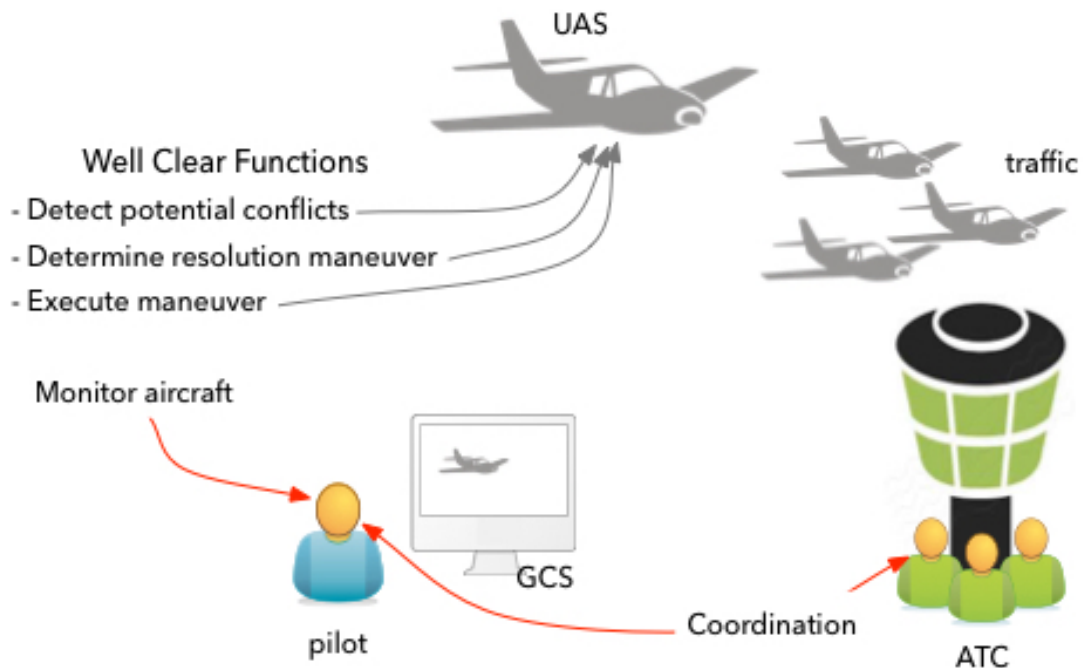


Figure 5. DAA human-automation architecture in which the human accomplishes all three well clear functions.

Between these two ends of the spectrum there are multiple ways for allocating partial or full responsibility for the well clear functions between onboard and ground-based machines and the ground-based pilot. While each option utilizes the same high-level architecture, whereby the pilot remotely operates the aircraft from the GCS, each permutation of function allocation fundamentally changes the details of the human-automation architecture and with it the coordination and synchronization requirements to support effective human-machine cooperation. This fact underscores the danger in taking only a simple, high-level view of the human-automation architecture and not fully analyzing the roles and responsibilities within that architecture and the resulting human-machine cooperation needs.

For all of the possible human-automation architectures of the DAA system, even for the fully autonomous option shown in Figure 4, an HMI is required to support the coordination of the shared responsibility of maintaining well clear between the pilot and machines, which becomes a critical component of the architecture. This is because the current regulatory environment requires that the pilot in command of the aircraft maintain ultimate responsibility for the safety of flight of the aircraft, regardless of whether that pilot is located on board the aircraft or in a GCS. In addition, the UAS pilot must fulfill the operational responsibility of coordinating any maneuvers off of their approved flight plan with ATC under most conditions. Thus, even in the fully autonomous option, where the onboard automation has authority to execute a well clear maneuver without pilot input, a pilot must still be capable of monitoring the activities of the aircraft and have

sufficient information about its operations that s/he can coordinate with ATC about the automation's functioning and intervene in that functioning if necessary.

In the fully autonomous DAA case, these pilot responsibilities add at least one additional (though likely more) well clear function(s) to monitor the automation's activity. In order to effectively assess whether the automation's behavior is appropriate in a given context, the pilot must also be able to monitor the automation's accomplishment of the detecting and determining functions (Figure 6). In this "fully autonomous" case, the human-machine interface must provide sufficient information so that the pilot can coordinate his/her activity to intervene if necessary. In addition, that information must be timely enough relative to an evolving conflict in order to meet the synchronization demands of intervening soon enough to maintain well clear, which as will become apparent in later sections, always takes place within a constrained timeline. The differences between Figure 4 and Figure 6 highlights the oversimplification fallacy associated with function allocation by substitution: while designers may envision the human-automation architecture depicted in Figure 4, the one depicted Figure 6 is much more realistic – and much more complicated, even in these simplified diagrams. The predictable result of this oversimplification when envisioning the human-automation architecture is that the functions that the pilot is responsible is for, which require coordination and synchronization with the machine components, are ill supported, resulting in coordination breakdowns between humans and machines and possible system failure. The takeaway from this discussion and the figures provided, is that despite all of the human-automation architectures for proposed UAS DAA systems having the same

high-level structure inherent in a remotely piloted aircraft, they are all different at a functional level and will therefore all require very different human-machine interfaces to support the human-machine cooperation.

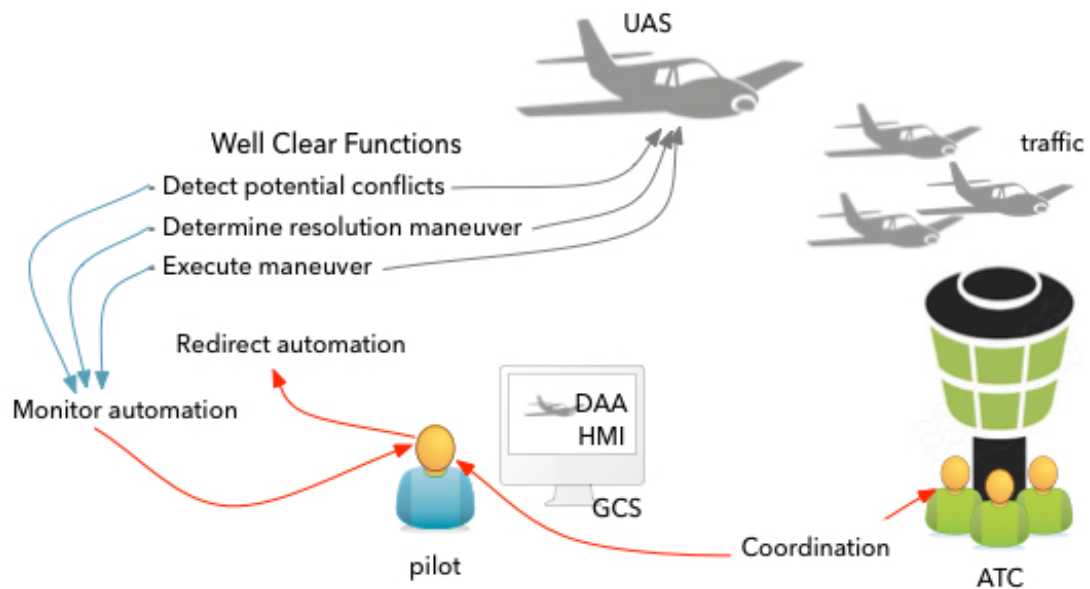


Figure 6. True DAA human-automation architecture in which the machine is allocated all three well clear functions, but pilot remains ultimate responsibility for safety of aircraft.

DAA Roles and Responsibilities Summary

Figure 7 provides a high level overview of how the three levels and their associated roles are interrelated with respect to DAA. Each level essentially represents different scales that the DAA system can be viewed at. It is important to note that the human roles at each level are critical parts of the system. In fact, each level should be considered a joint cognitive system with multiple, interacting, intelligent agents (both human and machine). While this fact may seem self-evident on the regulatory and

operational levels where humans carry out key functions, it is sometimes lost on engineers and regulators that pilots should also be considered part of the DAA system. When the machine components of the any system are developed without consideration of the human role(s) associated with that system, predictable human-cooperation problems occur.

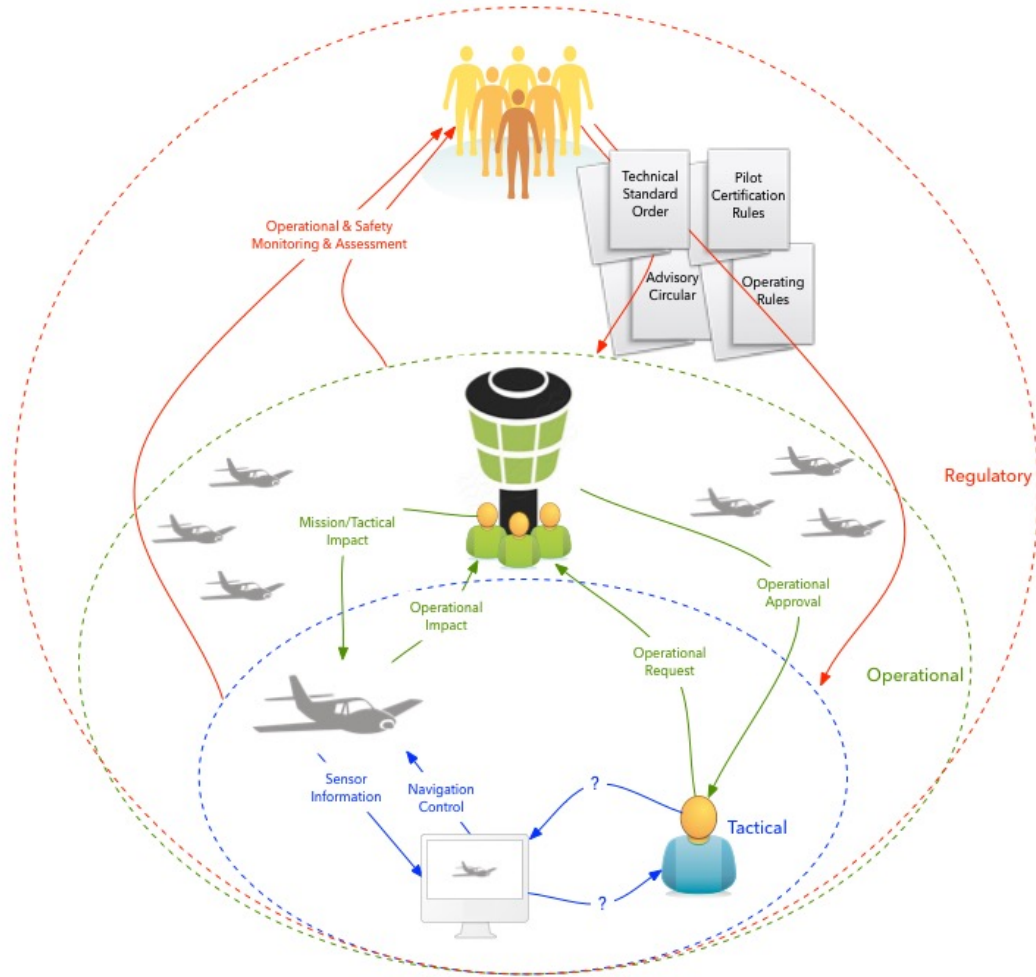


Figure 7. Levels and interactions of DAA roles and responsibilities.

Another feature of the DAA levels shown in Figure 7 is that they are both nested and hierarchical. The nested nature of the levels represents how the activities of each inner level are contained within the activities of the outer level. The hierarchical nature of the nested levels shows the authority that the outer levels have on the inner levels. At the top outmost level, the FAA as the regulators develop regulatory documents that specify the technical standards for DAA systems as well as operational rules and pilot certification requirements that affect the execution of roles at the operational and tactical levels. In addition, the regulators will monitor and assess the activities in each of the two lower levels to determine whether changes to various DAA-related regulations are required as a result of observed or anticipated safety or efficiency degradations.

In the middle level, ATC is responsible for executing the rules defined by the FAA within the contextual factors of the operational environment in order to ensure safety through maintenance of standardized separation minima while at the same time maximizing efficiency (i.e., throughput). For UAS equipped with DAA systems, a major factor influencing ATC's ability to meet the demands of safety and efficiency is the ability to predict the behavior of UAS pilots in maintaining well clear. ATC's ability to carry out their responsibilities is facilitated to the extent that UAS pilot's maneuvers are predictable and coordinated beforehand. A well clear maneuver by any aircraft in the ATC sector can result in maneuvers by other aircraft in the sector, which can result in maneuvers by other aircraft, and so on. One can imagine a situation where aircraft maneuvers propagate through the system like dominos. These reverberations are exacerbated when maneuvers are frequent and/or unpredictable.

In return for predictable behavior, ATC can more easily facilitate the execution of the UAS's mission and those of other aircraft in their airspace. To the extent that UAS pilots' execution of well clear maneuvers with the assistance of the DAA system interfere with ATC's ability to maintain separation minima and maximize throughput of the airspace, ATC may place operational restrictions on UAS pilots which degrades their ability to carry out their own missions.

Finally, at the inner most level, UAS pilots are responsible for the safety of the aircraft by utilizing the DAA system to assist them in maintaining well clear of other aircraft. Their ability to carry out this function is directly impacted by the technical equipment standards defined by the FAA, and their actions in carrying out this function directly impact the ability of ATC to carry out their role at the operational level. While ATC does not directly impact the pilot's ability to execute this particular function, formal or informal ATC rules can impact the pilot's ability to execute their missions effectively. The key questions remaining in order to completely define the roles and responsibilities within the system, is to specify the human-automation architecture between the pilot and the DAA system at the tactical level. Once that is determined, the human-machine interface requirements to support the coordination and synchronization within the joint cognitive system can be developed. Unfortunately, human-automation architectures for systems are rarely specified in the design or requirements development phases, and instead must be derived implicitly from the stated system assumptions or requirements.

Development of Minimum Operational Performance Standards for DAA Systems

Today the FAA faces a number of challenges in determining the minimum equipment standards or requirements for a UAS DAA system to ensure that UAS will maintain a safe separation distance (i.e., well clear) from other aircraft at an acceptable level of performance so that they are consistent with, and do not negatively impact, the current civil air traffic level of safety. The set of minimum standards that the FAA develops will indirectly specify the human-automation architecture of the DAA system through explicitly stated assumptions about the system as well as the final equipment requirements. The HMI requirements in particular, which are intended to support the pilot in command's ability to maintain well clear at the specified safety or performance level, will provide some insight into the intended human-automation architecture. If the human-automation architecture is not well understood during the requirements development process, which is likely given that human-automation architectures are not typically well-specified by engineers or regulators in the design phase, then the HMI requirements are at a high risk of not being well matched to support the human-machine coordination and synchronization needed for effective performance.

A second challenge for the FAA is that the regulations for DAA technologies need to be determined in advance of UAS being deployed operationally. No UAS currently has a DAA capability – those that are currently operating in civil or military airspace fly via segregated airspace procedures. In fact, many of the hardware and software technologies required by a DAA system are currently only in the research or prototype stage of development. At a high level, even the function of maintaining well

clear based on the electronic detection of other aircraft is entirely new to the aviation domain including both the manned and unmanned pilot populations. Because UAS present such a fundamental change in the human-automation architecture that has not been seen before in the history of aviation, this puts increased pressure on the ability of the FAA to accurately envision the future operational world and the impact that UAS employing DAA technologies will have on it *before* they allow these systems to be deployed. The degree to which methodologies utilized by the FAA in developing and testing equipment requirements accurately represent the future operational environment will be directly correlated to the true safety outcomes of introducing UAS into the NAS.

RTCA Special Committee 228 (SC-228) is the advisory committee that was established in order to help the FAA develop the MOPS for DAA systems for UAS (RTCA, 2013b). This committee is a consortium of government, industry, and academic engineers and subject matter experts whose technical expertise includes: surveillance and avionics equipment, aerospace engineering, human factors, air traffic management and operations, and UAS manufacturing and operations. The final DAA MOPS will cover requirements for all equipment and processes required for a DAA system. The research detailed in following chapters was conducted in support of developing the MOPS for the HMI of the DAA system. The methodologies employed in this research were developed to address the following needs of the committee: 1) evaluation of potential DAA HMIs in order to determine the minimum requirements for a DAA system; 2) development of metrics that can be used to assess and predict the performance of the prototype interfaces

in the envisioned world; and, 3) testing in a relevant environment that captures the envisioned future operational context.

Minimum Operational Performance Standards for a DAA HMI

The primary HMI in the DAA system, and the focus of the reported research, is the DAA traffic display, which may provide the pilot with three key informational features: traffic information elements, alerting, and maneuver guidance. These features support the pilot's ability to detect potential conflicts and determine maneuvers to resolve them. Information elements provide key data from nearby aircraft, such as altitude, heading, and speed. This data may be derived directly from surveillance sources, or they may be the result of computations by onboard or ground-based algorithms. An algorithm that uses aircraft state and/or intent information to extrapolate the time and location of CPA between the ownship and an intruder provides alerting information. The algorithm then applies a threat or alert level based on specified alerting thresholds; this alert level is presented to the pilot on the DAA display. Finally, maneuver guidance is an algorithm-based decision-aiding tool that provides the pilot with one or more trajectories that would be conflict free. Maneuver guidance may be directive or suggestive. Directive guidance supports allocation of this function entirely to the machine, and is based on an algorithm that provides a single, specific recommended resolution. In terms of maneuver guidance for maintaining well clear, the algorithm may choose a maneuver based on pre-determined heuristics, such as maximizing separation at CPA or minimizing time to gain (or regain) well clear. Suggestive guidance supports shared allocation of the determine function between the human and machine, it is based on an algorithm that provides a

range of potential resolutions from which the pilot can select the one that s/he deems is most appropriate.

Information elements and alerting are common features of manned cockpit display of traffic information (CDTIs), and directive maneuver guidance is utilized in TCAS II displays, however, these features are intended to assist pilot's normal visual acquisition of other aircraft. In fact, pilots are specifically prohibited from using information from a TCAS II traffic display as a basis for maneuvering without visual acquisition of the traffic they are maneuvering for (Department of Transportation, 2013), underscoring the novelty of the UAS DAA function of remaining well clear based solely on electronic systems. Thus, there is very little information about minimum information requirements for maintaining well clear on the basis of electronic displays of traffic information.

Previous Research

While there have been several recent human-in-the-loop evaluations of traffic displays for UAS (Calhoun, Miller, Hughes & Draper, 2014; Fern, Flaherty, Shively & Turpin, 2011; Fern & Shively, 2011), most have focused on military airspace operations rather than civil airspace operations. In addition, only a limited number of studies have examined traffic displays in the context of specific information requirements and their effect on pilot performance; that is, most studies have examined *how* to present a pre-determined set of information rather than *what* specific information is required.

A recent study examined the minimum visual information requirements for a UAS DAA system by comparing four display categories with different, and progressively

more, information: Position, Direction, Prediction, and Rate (Friedman-Berg, Rein & Racine, 2014). The Position display provided the pilot participants with instantaneous intruder positions and included aircraft identification, range, bearing, relative altitude, absolute altitude and numeric range. The Direction display added directionality as well as a heading chevron, numeric heading, and a vertical trend arrow. The Prediction display added yellow and red alert color-coding, and 30 second dead-reckoning vector lines to intruders. Finally, the Rate display added ground speed, history trails, and climb/descent rates. This study found that the Prediction display performed as good as, or better, than the Position and Direction displays, and no different than the Rate displays across a number of performance and workload metrics. Perhaps most importantly, the Prediction display resulted in significantly less NMACs than the two lower level displays (Position and Direction), and was not significantly different from the Rate display. The authors concluded that the information included in the Prediction display was the minimum visual information required for a DAA pilot display.

A survey conducted by Draper, Pack, Darrah, Moulton & Calhoun (2014), found that the majority of pilots surveyed indicated that the following information should be present at all times on a DAA display: intruder identification, intruder location, intruder relative position, intruder threat/alert level, DAA task priorities and status, DAA maneuver recommendations, flight restrictions, weather, navigation data, and visual alerts. While Friedman-Berg et al. (2014) and Draper et al. (2014) sought to establish a minimum information set, Bell, Drury, Estes & Reynolds (2012) compared a basic display (similar to the Direction display described in the Friedman-Berg et al. study) to

two “advanced” concepts: a display that depicted the relative CPA between ownship and intruder, and a display that depicted ownship avoidance areas with polygon shapes. This study found no significant differences in the frequency of violations of the defined well clear threshold between the three display concepts, however, there was a significant difference in the duration of violation events – the basic display had significantly longer violation durations than both of the other two advanced displays. Together, these three studies provided a starting point for the systematic evaluation of candidate DAA displays, specifically with respect to determining the minimum information requirements for such displays. The next chapter will provide an overview of the experiments that were conducted in order to inform the HMI requirements for future DAA systems, each of which will be presented in subsequent chapters.

Chapter 3: Research Overview

Three separate experiments were conducted to systematically evaluate potential DAA HMI configurations in order to inform the MOPS for UAS DAA systems. Each study evaluated the effect of various DAA HMI configurations on pilot performance of the function of maintaining well clear. The experiments reported in the following chapters were conducted iteratively, with each consecutive study building upon the lessons learned from the study that came before. Therefore, each experiment shares certain similar elements that are summarized below, such as: HMI components, human performance metrics, and the simulation of a relevant environment. They are all slightly different however, since new developments and decisions from the SC-228 advisory committee had to be incorporated prior to each experiment so that the resulting DAA human machine requirements would be consistent with other DAA requirements. In addition to documenting the experimental design for each experiment, the previous lessons learned and SC-228 decisions that informed the research question will be outlined. The results of each study will be discussed independently in each experiment chapter relevant to the research question being asked. The discussion will address all three experiments taken together. While some aspects and results of each of these three studies have appeared in other articles and reports, this document provides the first comprehensive description of how these studies link to the development of regulatory requirements.

Metrics of Human Performance

Two key sets of metrics were developed in order to quantify and assess human performance: measured response and loss of well clear. Measured response metrics support two goals of the research. First, these metrics allow for the quantification of the pilot contribution to the DAA timeline shown in Figure 2. The portion of the timeline that captures only the pilot's role in a DAA event is referred to as the "pilot-DAA timeline". Second, measured response metrics provide a means for comparing and evaluating potential DAA HMIs. The loss of well clear metrics also support the ability to evaluate potential architectures and interfaces, but one that is more operationally relevant to safety considerations by the regulators.

Measured Response

Measured response has been used to quantify the end-to-end response time for a UAS pilot to complete an Air Traffic Control (ATC) clearance (Shively, Vu & Buker, 2013; Vu, Morales, Chiappe, Strybel, Battiste, Shively & Buker, 2013; Rorie & Fern, 2014). By breaking down the end-to-end response into discrete stages, from the issuance of the clearance until the UAS completes the maneuver, it is possible to extract discrete response time metrics, such as the time required for a pilot to initiate a control input into the GCS, or the time required to complete a control input in the GCS. A study by Rorie and Fern (2014), found these response time metrics to be sensitive to differences in GCS command and control input interfaces. For the series of experiments detailed in this document, measured response was adapted for the DAA pilot task by quantifying the end-to-end response time for a UAS pilot to complete a well clear maneuver in response

to an alert presented on the DAA display. The primary stages of the pilot-DAA timeline extend from the time that a DAA alert appears on the pilot’s display to when the aircraft completes the subsequent avoidance maneuver.

Table 3 shows the stages of the pilot-DAA timeline. From this timeline, a number of different metrics can be extracted, such as the time it takes pilots to initiate a maneuver response in the GCS following the appearance of the alert ($T_3 - T_0$; i.e., initial response time), or the time it takes the pilot to upload the final maneuver response using GCS control interfaces ($T_{4b} - T_3$, i.e., total edit time). Figure 8 illustrates the relationship between several of the metrics that can be generated using the pilot-DAA timeline. Note that this timeline assumes that the pilot is responsible for executing the resolution maneuver (stages T_{4a} and T_{4b}). For human-automation architectures where automation is responsible for this, the stage would simply have to be renamed to reflect that, however, those stages would still have to be executed by the system.

Table 3. Stages of the pilot-DAA timeline.

Stage	Description
T_0	DAA (self separation or collision avoidance) alert appears on the display
T_1	Pilot notifies ATC and requests a maneuver clearance
T_2	ATC provides maneuver clearance
T_3	Pilot initiates an edit in GCS to maneuver
T_{4a}	Pilot uploads 1 st maneuver to aircraft
T_{4b}	Pilot uploads final maneuver to aircraft
T_5	Traffic alert is removed from display
T_6	UAS completes maneuver

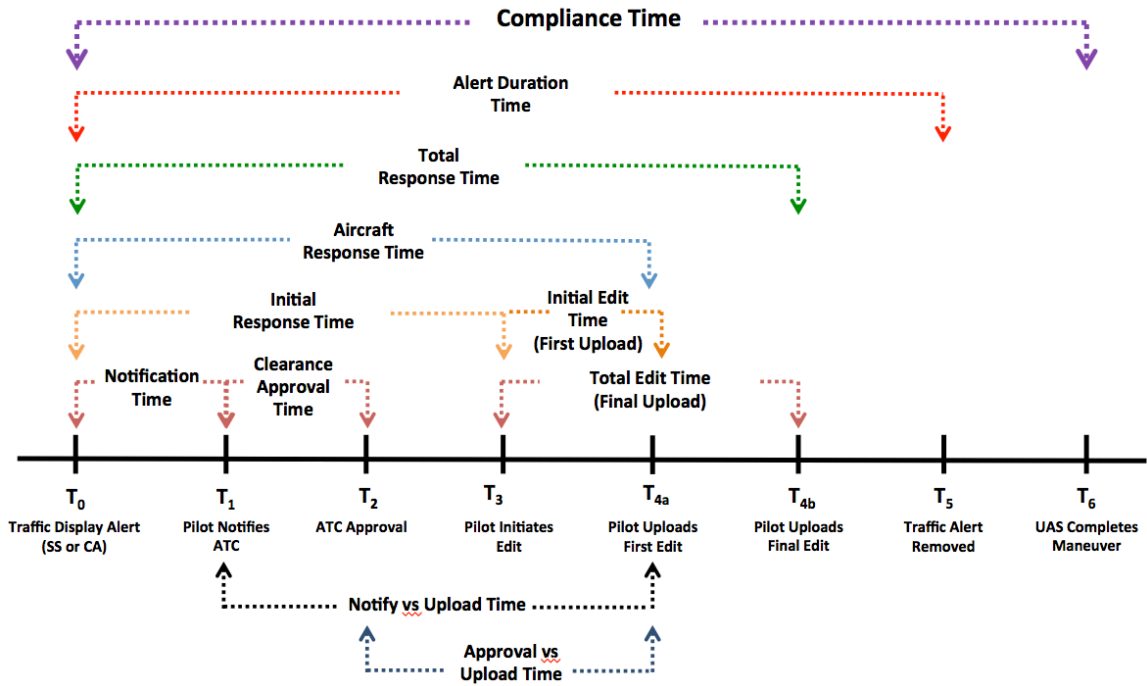


Figure 8. Relationship between stages of the pilot-DAA timeline and measured response metrics.

Measured response allows for the pilot’s contribution to the overall DAA timeline to be quantified. Figure 9 shows the DAA event timeline broken down into time segments. Pilot response time is critical to performing the DAA function – delayed pilot responses could result in delayed aircraft maneuvering, which in turn could cause losses of well clear. A delayed response may be due to a pilot’s inability to carry out one of the three major DAA sub-functions: detecting potential conflicts, determining a resolution maneuver, or executing a resolution maneuver.

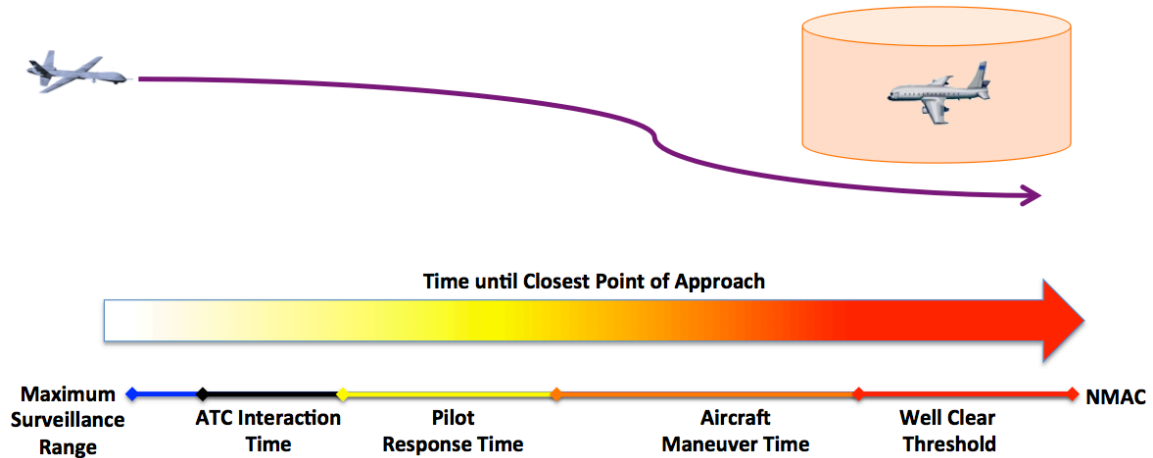


Figure 9. DAA Timeline broken into time segments.

Alternatively, system technology may be a limiting factor for pilot response times. Surveillance ranges dictate the furthest distance, and therefore time, that intruders can be detected. Active on-board radars, which are required to detect non-cooperating aircraft, are a newly developed technology for which the state-of-the-art detection ranges are still limited compared to cooperative sensor technology. Aircraft performance will also limit pilot response times, as enough time must be allotted in the DAA timeline for any conflict resolution maneuver to be executed. With respect to the technological limitations on pilot response times, there are obviously tradeoffs – better surveillance and aircraft performance can accommodate longer pilot response times.

Unfortunately, the operational environment may constrain the amount of time that pilots have to detect and/or respond to potential conflicts in ways that are not as easily changed as the technological constraints. Pilots may be constrained in when they can detect threats; late detections may result from various encounter geometries, especially

late maneuvering aircraft (e.g., last minute changes in altitude or heading), which cause intruders to cross alerting thresholds at closer distances than the maximum surveillance ranges allow. Pilots may also be constrained in when they can maneuver against threats by operational rules. As previously mentioned, pilots are required to coordinate any maneuvers off of their approved flight plan under normal conditions, which adds time to the timeline. In addition, ATC may find it unacceptable and disruptive for pilots to maneuver too early against a potential threat. Consequently, considerations for interoperability with ATC must be taken into account.

Figure 9 provides a nominal overview of the DAA timeline, taking these factors into account. The timeline is bound on the left by the maximum surveillance range that can be achieved and on the right by an NMAC, and has 5 major components (from left to right): 1) maximum surveillance range, 2) ATC interaction time (i.e., the time it takes pilots to coordinate a maneuver off of their approved flight plan with ATC), 3) pilot response time, 4) aircraft maneuver time, and 5) the time component of the well clear threshold. Overall, pilots are constrained in the maximum time that they have to respond to a potential threat by intruder trajectories, ATC interoperability requirements and the capabilities of the surveillance technology on one side of the DAA timeline, and aircraft performance on the other. Under normal conditions, a pilot may have the maximum allotted time in which to detect, determine and execute a well clear maneuver. However, pilots may have only a very short timeline in which to respond due to late maneuvering aircraft and other off nominal situations that must be taken into account when developing the minimum system requirements.

Comparing Figure 8 and Figure 9, the ATC interaction and pilot response times in the DAA timeline together correlate to T_0 to T_{4b} in the pilot-DAA timeline. Thus, measured response is not only able to capture the response times for pilots interacting with a DAA system; it also captures the effect of the end-to-end system performance. In addition, the two timelines above capture the synchronization requirements of the DAA system. That is, it highlights were activities need to be synchronized in order for a UAS to effectively maneuver to remain well clear of potential threats. The DAA equipment must synchronize the detection and display of threat and resolution information timely enough for pilots to determine and/or execute a maneuver to avoid the conflict. Further, pilots need to coordinate their maneuver with ATC in order to support interoperability with them and other users of the airspace.

Measured response metrics can also be used to evaluate the effect of different HMIs on human performance. Since the pilot-DAA timeline shown in Figure 8 is divided into several stages in order to capture response times for different pilot interactions with the DAA system, different DAA systems can be evaluated for how they affect these various measures. For example, certain DAA architectures may degrade response times in initially responding to the appearance of an alert (i.e., initial response time) while other architectures may degrade how long it takes for the pilot (or automation) to determine a resolution maneuver (i.e., total edit time). Having multiple metrics that capture different stages of the timeline allow for these types of comparisons to be made across various potential systems.

Given potentially limited timelines, pilots will be expected to respond in a timely manner, and that response needs to be taken into account when designing various components of the DAA system. Thus, the ability to evaluate the anticipated pilot response times for a given human-automation architecture and assess how that affects the rest of the DAA system and its ability to meet the safety requirement of remaining well clear of other aircraft is critical for developing the DAA performance standards. However, measured response metrics do not provide insight into how well the system is performing in remaining well clear, in order to do that, loss of well clear metrics need to be analyzed also.

The following seven measured response metrics are reported for all three experiments.

Notification Time ($T_1 - T_0$). A measure of the time it takes a pilot to notify ATC of the appearance of a self separation or collision avoidance alert (T_0) and the need to execute a maneuver (T_1). Calculated as the difference between the first appearance of the alert and the beginning of the pilot's transmission on the radio to notify the controller.

Clearance Approval Time ($T_2 - T_1$). A measure of the time it takes for ATC to approve a pilot's request to maneuver in response to a DAA alert. Calculated as the time between the beginning of the pilot's transmission to ATC (T_1) to the beginning of the ATC's response (T_2).

Proportion of Uploads with Clearance Approval (T_2 versus T_{4a}). The percentage of encounters where pilots received ATC approval (T_2) *prior* to executing a resolution maneuver in response to a DAA alert (T_{4a}).

Initial Response Time ($T_3 - T_0$). A measure of the time it takes a pilot to initiate a maneuver response, or edit, in the GCS command and control interface (T_3) in response to a traffic display alert (T_0). Calculated as the difference between the first appearance of the alert and the start of an edit.

Initial Edit Time ($T_{4a} - T_3$). A measure of the time it takes a pilot to input an initial edit into the GCS in order to maneuver in response to a traffic alert. Calculated as the time between initiating an edit (T_3) and the first upload to the aircraft (T_{4a}). This metric is only relevant if a pilot uploaded multiple edits to the aircraft; when pilots made only one edit, initial edit time is equivalent to total edit time.

Total Edit Time ($T_{4b} - T_3$). A measure of the time it takes a pilot to complete an edit into the GCS in order to maneuver in response to a traffic alert. Calculated as the time between initiating an edit (T_3) and the final upload to the aircraft (T_{4b}). The final upload is assumed to be the pilot's final resolution decision.

Total Response Time ($T_{4b} - T_0$). A measure of the time it takes a pilot to upload a final maneuver resolution to the aircraft in response to a traffic alert. Calculated as the time between the initial appearance of the traffic alert (T_0) and the final maneuver upload to the aircraft (T_{4b}).

Loss of Well Clear

The loss of well clear metrics are an operational measure of the DAA system's performance in maintaining well clear from other aircraft. The primary measure is the occurrence or rate of losses of well clear (i.e., penetration of the defined well clear threshold) given a particular instantiation of a DAA system. A measure of the severity of losses of well clear when they occur is also used to help to differentiate the performance of different DAA configurations. The severity metric reported in the studies in this document is an index of the separation between the aircraft at the CPA. The separation severity index, S_{index} , (see Equation 1) is defined as the larger of the horizontal and vertical separations normalized by the required separation in each dimension where the h_sep_{CA} and v_sep_{CA} are the geometric portions of the well clear definition. Thus, the severity metrics captures the proportion of the spatial well clear volume that was penetrated when the loss of well clear happened.

Equation 1. *Separation index.*

$$S_{index} = \min_t \left\{ \max \left[\frac{horz.range(t_i)}{h_sep_{CA}}, \frac{vert.range(t_i)}{v_sep_{CA}} \right] \right\}$$

Relevant Environment

The measured response and loss of well clear metrics of pilot performance were developed in order to provide a means for evaluating and comparing the performance of various DAA HMI configurations. As these prototypes are evaluated, the metrics become predictions for how they will perform in the future operational environment and are

critical to the safety assurance of the system. However, the metrics are valid for predicting future performance only to the extent that the environment that the systems are tested in encompasses the key operational features and demands that exist in the future environment. Therefore, in order to be useful in providing safety assurance to regulators who assess the results of the system evaluations, the critical features of the operational environment must be captured in a relevant environment.

Fortunately, the DAA Operational Services and Environment Description (RTCA, 2015b) document developed by the SC-228 committee provided substantial information on the target operational environment from which the simulation environment utilized in the DAA HMI experiments was modeled on. Key aspects of the operational environment that were modeled in the simulation environment included: airspace class, airspace density, operational and ATC interoperability rules, and separation requirements.

All of the experiments reported in this document utilized a simulated UAS operating in a full airspace sector with other manned aircraft and a live air traffic controller. One benefit of conducting research with UAS compared to manned aircraft, is that it is much easier to create a realistic simulation environment since sophisticated motion and visual techniques are not needed to replicate a true operational setting. Utilizing a live controller to manage the simulated ATC sector was considered critical since interactions with ATC are such a critical part of the environment that UAS will operate in, and because of concerns that UAS will maneuver without prior coordination based on the DAA system. In addition, confederate “pseudo” pilots managing the simulated manned aircraft in the airspace sector provided realistic communications on the

ATC communication channel and created representative situations where UAS pilots had to wait for openings on a busy channel to request clearances to maneuver from the controller. Finally, “full mission” scenarios were employed whereby pilots operated a simulated UAS for longer periods of time than just a single encounter (which is typical for “part task” experimental designs) and were tasked with not only maintaining well clear of other aircraft, but also with carrying out other secondary mission-related tasks which took their attention away from the DAA display. As pilots executed their UAS missions, they would encounter a number of conflicts with different aircraft.

In order to test the performance of pilots’ ability to maintain well clear given various DAA HMI designs, a number of different encounter types between the UAS ownship and other simulated manned aircraft were generated. Encounters varied in their predicted distance at CPA with the ownship in order to trigger various alert levels of the DAA system. Varying distances at CPA causes different alert levels and test whether pilots will respond as expected given training. For example, some alert levels indicate that pilots needs to make an imminent maneuver while others indicate no maneuver is required. If pilots maneuver for the latter, or don’t maneuver for the former, then this behavior points to issues in the training and/or design of the alert levels or other HMI features. Thus, scenarios were specially designed within the constraints of the simulated environment to elicit the activation of particular display elements in order to gather empirical data on pilot performance, in terms of measured response and loss of well clear metrics, while maintaining well clear of other aircraft.

Now that an overview of the general approach to method and measures for the three experiments has been provided, the following chapters will present each separate experiment in detail.

Chapter 4: Experiment One

Overview

The first meeting of SC-228 was held on 30-31 July 2013. The final DAA white paper that defines the assumptions, approach, and core requirements for a DAA system for UAS integration into the NAS was completed on 1 November of the same year (RTCA, 2013a), though the problem space of defining requirements for a future system was still largely indeterminate. One such case in point, although the DAA white paper provided the following assumption: “UAS pilots will be authorized to use their DAA system to remain “well clear” using new or revised rules analogous to 14 CFR §91.113 and §91.181,” (p. 10), there existed no agreed upon mathematical definition of well clear. Shortly after the white paper was published, the first experimental attempt to identify the minimum human machine interface requirements for a DAA system was in its planning stages.

Purpose

The goal of the first experiment was to evaluate candidate DAA display and algorithms with respect to maintaining well clear and avoiding collisions. Three specific research questions were developed:

1. What are the minimum information requirements for DAA displays?

2. What display features improve accuracy and expediency in determining, negotiating and executing traffic avoidance maneuvers?
3. Is there a performance difference between integrated and standalone displays?

The first two research questions addressed the main goal of the SC-228 DAA working group with respect to the DAA traffic display – identification of the minimum information requirements that support acceptable pilot performance in maintaining well clear from other aircraft. Since there is no truly comparable task in manned aviation where pilots rely solely on an electronic display of traffic information to make decisions about maneuvering, it was not clear if an informative display, which lacks maneuver guidance or decision aiding for the pilot, would be sufficient for pilots to effectively carry out the DAA tasks.

The third research question above, driven by near term technological considerations, was aimed at determining whether DAA the human machine interface, or DAA displays, should be integrated into the primary displays of a UAS GCS, or whether they could be ‘bootstrapped’ on a separate, standalone display within the GCS. Display location has the potential to significantly impact pilot performance in maintaining well clear and collision avoidance from other aircraft. A standalone display is considered a near term technology solution for existing UAS, because it is easier to develop, certify and field an independent, separate display compared to modify existing GCS software. However, the standalone display concept has some disadvantages. First, a standalone traffic display is unlikely to be integrated with the GCS command and control interface, thus a pilot would have to identify a potential threat and resolution maneuver, and then

translate that to the command and control interface on another display in order to execute the maneuver, and confusion could arise if the displays are at different zoom levels or orientations (e.g., track up versus north up). In addition, a standalone display is likely to be produced by a different manufacturer than the GCS, which could result in inconsistencies in the presentation of similar information between the DAA and GCS displays. Finally, pilot response times could be slower as pilots have to switch attention (and possibly interaction) between two different displays; previous UAS research has shown that switching between different information sources in the GCS can disrupt pilot performance (Draper, Calhoun, Ruff, Mullins, Lefebvre, Ayala, & Wright, 2008). While a DAA display that is integrated into the GCS primary display could overcome several of these disadvantages, in addition to the large increase in overhead in terms of resources to develop, a major risk of an integrated display is increased clutter on the primary display. Thus, display location in was considered a critical issue that needed to be understood in addition to minimum information elements.

Experimental Design and HMI Configurations

This study utilized a within-subjects, repeated measures factorial design to compare the effect of information level and display location on UAS pilots' performance on maintaining well clear and collision avoidance from other aircraft while operating in civil airspace. Two levels of information (basic, advanced) were compared across two levels of display location (standalone, integrated) for a total of four displays (Figure 10): 1) Basic Standalone, 2) Basic Integrated, 3) Advanced Standalone, and 4) Advanced Integrated.



Figure 10. The four DAA displays that were evaluated in Experiment 1: Standalone Basic (top left), Integrated Basic (top right), Standalone Advanced (bottom left), and Integrated Advanced (bottom right).

Display Location

For this experiment, the standalone DAA display condition could only receive ownership state and trajectory information from the navigation system, and could not send any information to it (e.g., to the command and control interface). In order to replicate a

completely independent system, the standalone display was configured to receive information from the UAS navigation system but was not capable of sending any information directly to the aircraft. All command and control changes to the aircraft were required to be made through the primary command and control interface in the GCS. The integrated display condition saw the DAA display features integrated directly into the primary display of the GCS. The DAA display was also integrated with the navigation system. Where possible, display features were integrated with the command and control interface. The alerting, information elements and maneuver guidance contained in each of the four different displays is described next.

Information Level

Two information levels were compared, basic and advanced, which differed in the information elements, maneuver guidance and alerting that was presented to pilots on the DAA display. For both basic display configurations, Standalone Basic and Integrated Basic, the same information elements were presented. However, the advanced display configurations, Standalone Advanced and Integrated Advanced, presented different informational elements based on different display capabilities of the underlying technology. Table 4 shows the alerting, information elements, and maneuver guidance for each display configuration.

Table 4. Alerting, information elements and display location for the four display configurations in Experiment 1.

	Basic Standalone	Basic Integrated	Advanced Standalone	Advanced Integrated
Alerting				
<i>Self Separation Alert</i>	✓	✓	✓	✓
<i>Predicted Collision Avoidance Alert</i>			✓	✓
<i>Collision Avoidance Alert</i>	✓	✓	✓	✓
Information Elements				
<i>Minimum Information Elements (Table 6)</i>	✓	✓	✓	✓
<i>CPA Location</i>			✓	✓
<i>Time to CPA</i>			✓	✓
<i>Well Clear Ring</i>				✓
<i>Vertical Situation Display</i>				✓
Maneuver Guidance				
<i>Auto-Resolutions</i>			✓	✓
<i>Trial Planner Tools</i>			✓	✓

Alerting

Thresholds. The alerting structure used in this experiment used a well clear definition of 0.8 nm lateral, 400 ft vertical, and 40 s to CPA. The DAA threshold for this experiment was 110 sec – the time to CPA of the self separation² and the predicted collision avoidance alerts, which alert to predicted losses of well clear. The 110 sec

² The term “self separation” was used frequently early in the development of the SC-228 MOPS as an alternate description of the pilot task to remain well clear (i.e., maintain self separation). Due to concerns from ATC organizations regarding the use of “separation” as a function that falls outside the scope of ATC responsibilities, the term was officially removed from the draft MOPS after the review official review in August 2015. The term has been replaced with “detect and avoid (DAA)” or “maintain well clear”. That convention has been followed in this document, except where the experiments reported have already been published with labels or names that include “self-separation” that are important to maintain for cross-reference purposes, such as alert names.

threshold for these two alerts is based on a human-in-the-loop (HITL) study with air traffic controllers that indicated that 120 sec was the maximum time from CPA that air traffic controllers felt comfortable allowing UAS pilots to maneuver against other aircraft; maneuvers that occurred at 120 sec or greater to CPA were considered disruptive to the (Mueller, Isaacson, & Stevens, unpublished). Taken together the well clear and DAA time thresholds, 40 and 110 sec, respectively, help to bound the DAA timeline, leaving a maximum 70 sec (depending on the encounter) for aircraft maneuvering, ATC interaction, and pilot response (Figure 11).

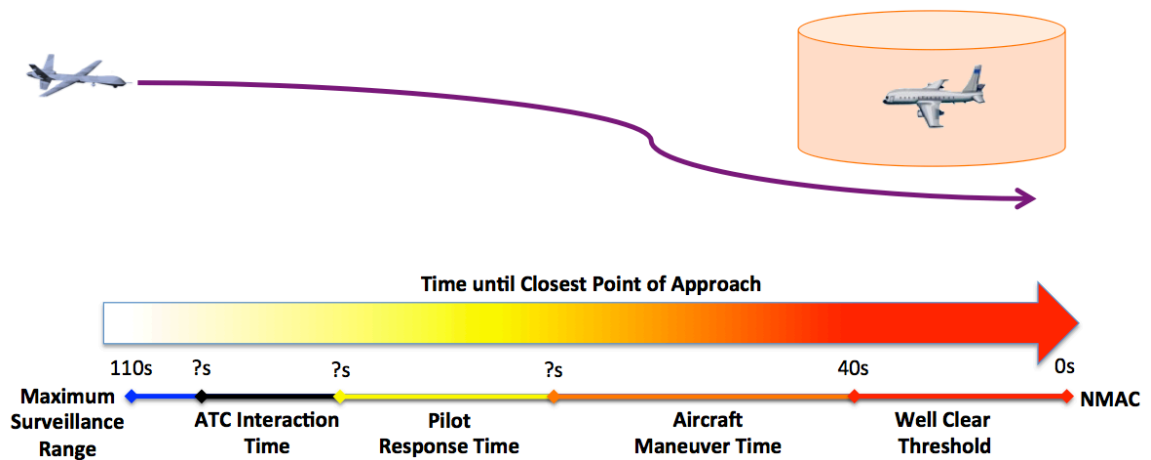


Figure 11. DAA timeline with the well clear and DAA alerting thresholds from Experiment 1.

Alert Levels.

Table 5 presents the multi-level alerting structure that was used in this experiment. Threat level was based on the location of and time to the CPA between the ownship and an intruder aircraft. Both the location and time thresholds had to be met for

an intruder to be assigned a threat level. Location of the CPA was measured by both the lateral distance, in nautical miles (nm), and vertical distance, in feet (ft), to the predicted CPA location. The time to CPA was measured in seconds (sec). For this experiment, the collision avoidance and well clear thresholds were treated as equivalent.

Table 5. The multi-level alerting structure used in Experiment 1. Only the advanced display configurations, Advanced Standalone and Advanced Integrated, included the “Predicted Collision Avoidance” alert level; it was not included in the basic display condition.

Alert/Threat Level	CPA Distance from Ownship		Time to	Color
	Lateral	Vertical	CPA	
Proximal	> 2 nm	> 900 ft	N/A	Grey
Preventative	< 2 nm	< 900 ft	< 120 sec	White
Self Separation	< 1.2 nm	< 900 ft	< 110 sec	Yellow
Predicted Collision Avoidance*	< 0.8 nm	< 400 ft	< 110 sec	Yellow, Red Border
Collision Avoidance	< 0.8 nm	< 400 ft	< 40 sec	Red

* Advanced Display configurations only

In this experiment, the proximal threat level indicated aircraft that were within the simulated surveillance range of ownship, but not within one of the higher-level alert thresholds. The preventive alert level indicated aircraft within a close range of ownship that should be monitored for changes in state that could cause a higher-level alert. The self separation alert indicated aircraft who’s CPA was predicted to be within a buffered well clear volume within the alerting threshold time. The buffered well clear added 0.4

nm laterally to the well clear threshold. Finally, the collision avoidance alert indicated an aircraft that had penetrated the well clear threshold.

The predicted collision avoidance alert level was introduced in the advanced display configurations to indicate the severity of a self separation alert. This alert was intended to aid pilots in discriminating between critical and non-critical self separation threats. The existence of a buffered lateral threshold for self separation alerts compared to collision avoidance alerts meant that some self separation alerts never progressed to collision avoidance alerts. The predicted collision avoidance alert level indicated critical self separation alerts that were predicted to progress to a collision avoidance alert, compared to the self separation alerts that were predicted to remain outside of the 0.8nm well clear threshold, but within the 1.2nm well clear buffer. In the Advanced Standalone display, the predicted collision avoidance alert was depicted with a red outline and yellow fill on the CPA location icon, while in the Advanced Integrated display condition, it was depicted with a red outline and yellow fill on both the intruder aircraft icon and the CPA location icon. Figure 12 depicts the difference between the self separation and predicted collision avoidance alert thresholds.

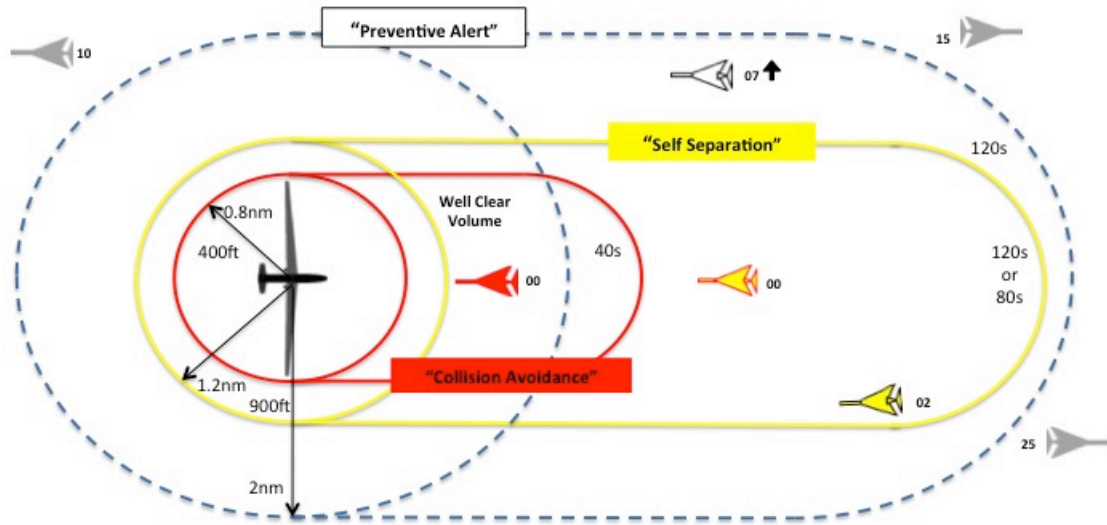


Figure 12. Visual depiction of DAA alert levels and thresholds

Information Elements

The results from Friedman-Berg et al. (2014) and Draper et al. (2014) discussed in the previous chapter, were crosschecked against existing relevant references and documents [e.g., Minimum Operational Performance Standards for Aircraft Surveillance Applications System (RTCA, 2011)], in order to come up with a set of minimum information elements that could be evaluated in this experiment as the basic condition. This baseline information set that made up the basic information condition is listed in Table 6.

Table 6. The minimum information elements for the basic display condition.

Intruder Information	Visibility
Location	Always visible
Range	Always visible
Bearing	Always visible
Heading	Always visible
Relative Altitude	Always visible
Vertical Trend	Always visible
Heading Predictor	Always visible
Vertical Velocity	Within data tag
Absolute Altitude	Within data tag
Ground Speed	Within data tag
Aircraft ID	Within data tag

The visibility of each information element identified as a minimum requirement was also specified, that is, whether it should be always visible on the DAA display, or whether it could be shown in a data tag. Information shown within a data tag was only displayed if: 1) the data tag was selected for an intruder with a Proximal or Preventative threat level

Table 5), or 2) an aircraft had a self separation or collision avoidance threat level (at which point the data tag information was automatically visible).

The advanced information configurations provided additional information elements to the minimum information set listed in Table 6. Both the advanced display configurations depicted the predicted CPA location and time to CPA, while only the Advanced Integrated DAA display configuration also included a well clear ring and a vertical situation display.

CPA Location. At the onset of a self separation alert, the predicted physical location of the intruder's CPA to ownship was depicted by a color-coded circle. In the Advanced Integrated condition, the ownship's predicted location at CPA could also be displayed by hovering over the aircraft icon of the potential intruder. The CPA location remained on the display as long as a threat was active and automatically disappeared once an alert was cleared.

Time to CPA. A countdown timer was triggered by the onset of a self separation alert, indicating the time remaining until CPA was reached. In the Advanced Standalone configuration, the time was displayed in data block the bottom left-hand corner of the display. In the Advanced Integrated configuration, the timer was displayed in the data tag of the relevant aircraft. The time disappeared from both the data tag and data block as soon as an alert was cleared.

Well Clear Ring: The presence of a self separation alert enabled the appearance of a "well clear ring" around ownship. With a radius of 0.8 nm, the ring gave pilots a visual reference of the lateral collision avoidance threshold.

Vertical Situation Display. A panel at the bottom of the Advanced Integrated configuration displayed a vertical profile of traffic +/- 1,000 ft vertically from ownship. The vertical situation showed traffic icons, heading predictors, CPA location, and appropriate color-coding for alert level. Range rings within the vertical situation display were designed to align with the range rings on the primary display of the GCS.

Maneuver Guidance

Maneuver guidance was provided only in the advanced display configurations. The maneuver guidance provided both directive guidance, which provides a single recommended maneuver solution, in a text format, and suggestive guidance, which provides a range of potential solutions, in the form of trial planning tools. Pilots were not required to use or follow the guidance provided by the resolution tools included in this condition. Since they were implemented substantially differently within the Advanced Standalone (Figure 13) and Advanced Integrated (Figure 14) displays, they are described separately. The implementation of the auto-resolutions and trial planning tools in the Advanced Standalone DAA display condition is described below (Figure 13).

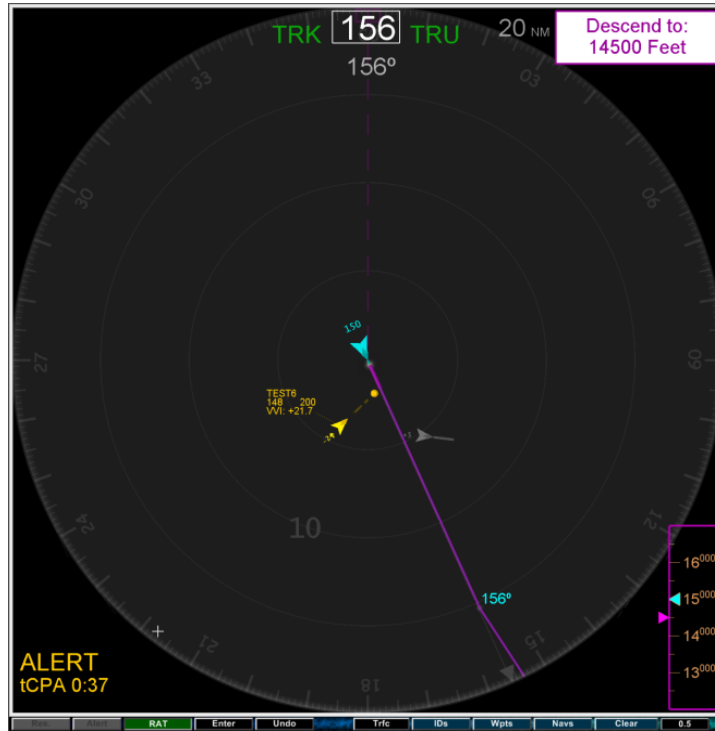


Figure 13. Screenshot of the Advanced Standalone DAA display condition during a self separation alert. The maneuver recommended by Autoresolver-AD is shown in the upper right-hand box (“Descend to: 14500 Feet”), the lateral trial planner is indicated by the magenta colored flight plan, the vertical trial planner (i.e., altitude tape) is located on the lower right side of the display, and the time to the predicted closest point of approach with the self separation alert is located in the bottom left.

Auto-Resolutions. At the onset of a self separation alert, Autoresolver-AD (described in the next section) provided pilots with a recommended resolution maneuver. The maneuver appeared in a text box in the upper right hand of the display. If Autoresolver-AD computed a more effective maneuver, a “Refresh” button flashed continuously at the bottom of the display. When pressed, the new resolution maneuver replaced the previous one.

Trial Planning Tools. Two separate trial planning tools, each intended to give pilots the ability to ‘test’ various heading and altitude vectors by providing immediate feedback as to the predicted threat level of a proposed maneuver, were included. The tools were engaged automatically during self separation alerts and populated with the maneuver recommended by Autoresolver-AD. However, pilots were also able to engage the trial planning tools manually by selecting a dedicated button (“RAT”) at the bottom of the display.

The lateral trial planning tool allowed pilots to superimpose a ‘proposed’ route line on top of their active route. The proposed route line could be manipulated without any impact to their active route. Using the superimposed route line, pilots could “trial plan” different heading vectors by clicking and dragging on a waypoint off the nose of the aircraft. As pilots moved a proposed waypoint away from their active route, a heading readout appeared adjacent to the waypoint, informing pilots of the exact heading vector being trial planned. The proposed route line was color-coded based on the predicted alert level for the trial-planned heading. Proposed headings that were predicted to lead to only proximal alerts turned the line magenta, while headings predicted to lead to at least one preventive alert turned the line white. Magenta was used instead of the corresponding grey color for proximal alerts, in order to be more visually salient to the pilot and distinctive from the actual current trajectory. Proposed headings that were predicted to lead to at least one self separation alert or at least one collision avoidance alert turned the route line yellow or red, respectively. While trial planning, halos appeared around

intruders that indicated their predicted threat level given the proposed trajectory. The halos were color-coded according to the predicted alert level of the associated intruder.

Similar to the lateral tool, the vertical trial planning tool allowed pilots to probe different altitudes for their relative safety. In order to trial plan different altitudes, pilots manipulated an altitude tape positioned in the bottom right-hand corner of the display. The altitude tape included three different altitude 'bugs': current, commanded, and trial planning. The altitude tape was centered on the trial planning altitude bug, allowing pilots to test various altitudes by clicking and dragging the altitude tape's surface. The trial plan altitude bug and the border of the altitude tape were then color coded using the same alerting logic described for the lateral trial planning tool.

The implementation of the auto-resolutions and trial planning tools in the Advanced Integrated DAA display condition is described below (Figure 14).

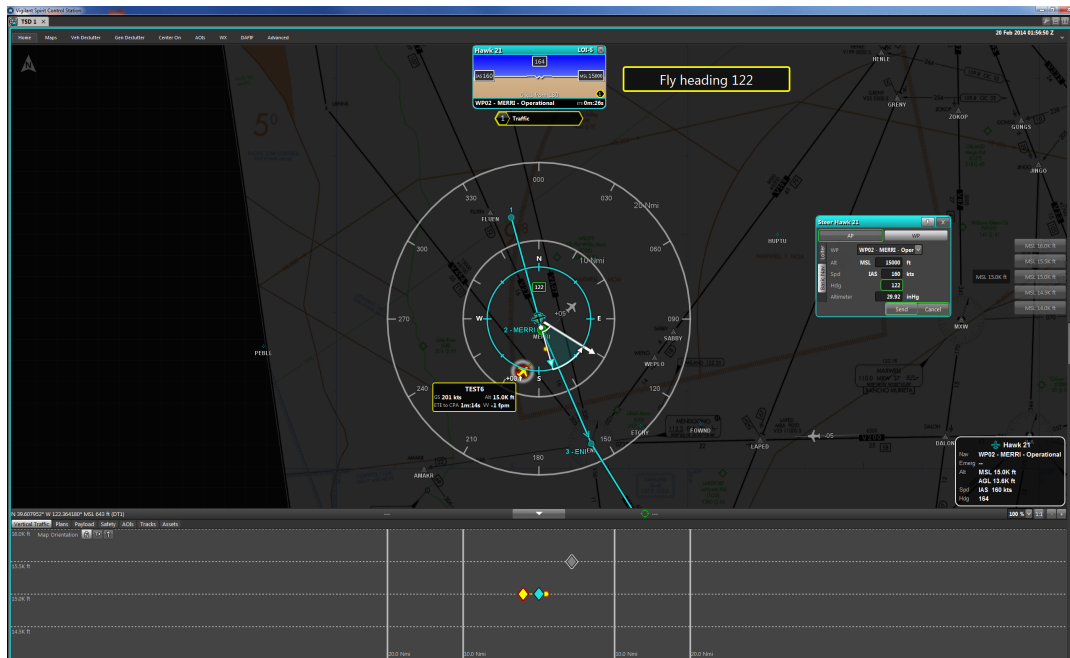


Figure 14. Screenshot of the Advanced Integrated DAA display condition with an active predicted collision avoidance threat. The maneuver recommended by Autoresolver-AD is shown in the upper box (“Fly Heading 122”), the lateral trial planner is indicated by the arrow pointing to heading 122 off the nose of the ownship icon (center), the vertical trial planner is located on the far right side of the TSD, and the vertical situation display is shown in the lower quarter of the display.

Auto-Resolutions. At the onset of a self separation alert, Autoresolver-AD provided pilots with a recommended maneuver. The text-based recommendation was displayed in the upper-right hand corner of the GCS primary display. If Autoresolver-AD computed a more effective maneuver, a “Refresh” button appeared directly below the recommended maneuver text. When pressed, the new resolution maneuver replaced the previous one.

Trial Planning Tools. As with the Advanced Standalone condition, two separate trial planning tools, lateral and vertical, were included in the display. While the tools’

overall function was still to allow pilots to test various heading and altitude vectors, their implementation had to be modified to allow for integration with the vehicle controls. Once again, the trial planning tools were automatically engaged at the onset of a self separation alert and populated with the maneuver recommended by Autoresolver-AD. The tools could also be launched manually in the absence of an active self separation alert.

The lateral trial planning tool was integrated into the GCS's autopilot interface. At the onset of a self separation alert the compass rose automatically opened on the primary GCS display and a vector arrow appeared, extending from ownship to the edge of the compass rose in the direction of the recommended Autoresolver-AD maneuver. As pilots dragged the heading bug, the collocated arrow gave instantaneous feedback as to the quality of the proposed heading. As with the proposed route line in the Advanced Standalone condition, the vector arrow was color-coded based on the safety level of the heading being probed. Proposed headings that were predicted to lead to only proximal alerts turned the vector arrow green, while headings predicted to lead to at least one preventive alert turned the arrow white. Proposed headings that were predicted to lead to at least one self separation alert or at least one collision avoidance alert turned the vector arrow yellow or red, respectively. Since the vector arrow was integrated into the autopilot interface, pilots could directly send any proposed heading holds up to the aircraft by pressing "Send" in the autopilot's steering command window.

The vertical trial planning tool utilized an altitude table permanently displayed on the far right side of the TSD. The altitude table consisted of five discrete altitude options:

one at the current altitude, two 1,000 ft above the current altitude (in 500 ft increments), and two 1,000 ft below the current altitude (also in 500 ft increments). The same color-coding scheme as described for the lateral trial planning tool was used for the altitude table to indicate the predicted safety level of each altitude option. Each altitude option was a selectable button and tied to the GCS's autopilot interface. When pressed, the selected altitude was pushed to the steering command window. Pilots could then upload the new altitude to the aircraft by pressing "Send" within the steering command window. While trial planning with the lateral or vertical tools, halos appeared around intruders that were predicted to become a proximal, self separation or collision avoidance alerts. The halos were color-coded according to predicted alert

Method

Participants

Twelve active duty RQ-4 pilots ($M = 39$ years of age) were recruited for this experiment. Participants had an average of 216 hours of experience flying UAS in combat and non-combat military operations. Eight of the participants had prior experience flying UAS in civil airspace, each with an average of 60 hours. A single retired air traffic controller served as a confederate for the study.

Simulation Environment

The simulation environment utilized in Experiment 1 contained a number of software and hardware components located at NASA Ames Research Center, all of which were networked using NASA's Live Virtual Constructive (LVC) gateway (Murphy & Kim, 2013).

Ground Control Station

Participants were situated at a UAS GCS containing two pieces of software distributed across four separate monitors, the Vigilant Spirit Control Station (VSCS) and the Cockpit Situation Display (CSD). Figure 15 shows the UAS GCS display set up. VSCS, developed by the Air Force Research Laboratory, is a mature GCS operator interface designed to support the control of UAS and their associated payloads (Feitshans, Rowe, Davis, Holland & Berger, 2008). In this study, VSCS generated three separate pilot displays in the ground station: a Tactical Situation Display (TSD), a health and status panel, and a simulated out-the-window nose camera display.

The TSD (shown in Figure 16) served as the pilot's primary display, providing ownership and route information, a moving map, and navigation and control interfaces. The TSD supported two separate vehicle control interfaces. The first, a waypoint-editing interface, allowed pilots to modify the assigned altitude or location of any waypoint on their mission route. The second, an autopilot editing interface, allowed pilots to enter into altitude, speed, and heading holds without modifying their mission route. Heading holds could be executed through numerical inputs to a steering command window or through interaction with a graphical compass rose interface (shown in Figure 17). The compass rose interface allowed pilots to drag a heading bug to their desired direction rather than enter the value manually. Altitude and speed holds, however, could only be executed using numerical inputs to the steering command window. Pilots uploaded changes to the aircraft by pressing a "Send" button within the steering command window. The TSD also hosted the DAA display in the integrated display configurations detailed above.



Figure 15. UAS GCS. CSD (not displayed), bottom left; TSD, bottom center; out-the-window view, top center; health and status panel, bottom right.

A health and status panel was positioned to the right of the TSD and contained subsystem information, telemetry data, a chat client, and an electronic checklist. The third VSCS display was a simulated out-the-window nose-camera view positioned directly above the TSD. This monitor included synthetic terrain and an integrated head up display with current airspeed, altitude and heading information. All interaction with the three VSCS components occurred through standard mouse and keyboard inputs.

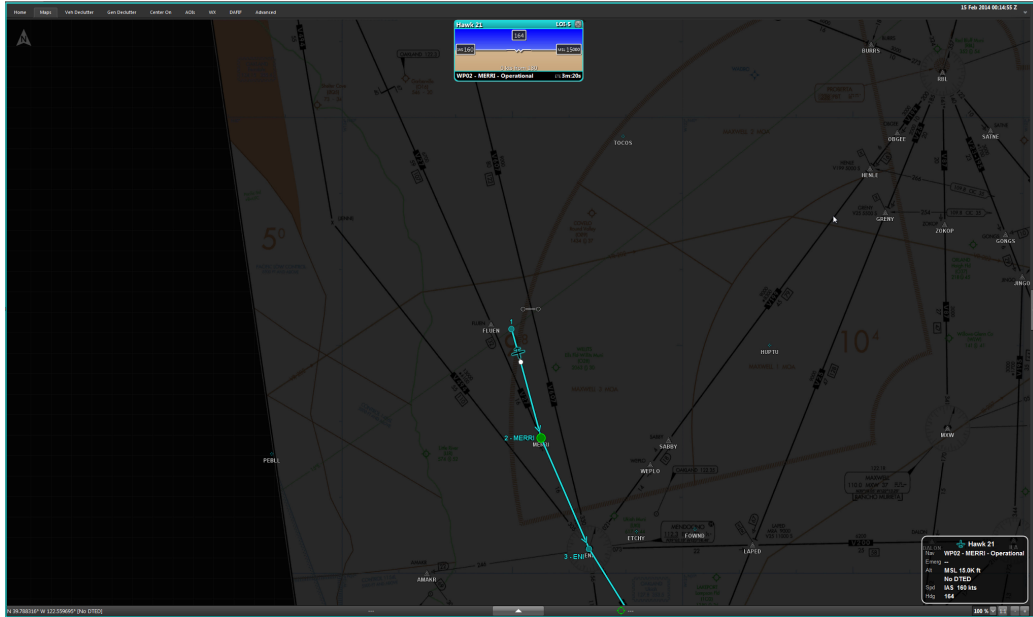


Figure 16. Vigilant Spirit Control Station tactical situation display (AFRL/RH).
 Distribution A: Approved for public release; distribution unlimited, 3/18/2013; 88ABW-2013-1303.



Figure 17. Vigilant Spirit Control Station's compass rose interface.

A fourth monitor was populated by the CSD (shown in Figure 18), a 3D volumetric cockpit display of traffic information developed by the Flight Deck Display

Research Laboratory at NASA Ames Research Center (Johnson, Battiste & Bochow, 1999). The CSD, positioned directly to the left of the TSD on a separate monitor, and as described above, was configured to display ownship information, surrounding traffic and alerting information, and, in select conditions, maneuver tools. Pilot interaction with the CSD, while limited, was enabled through standard keyboard and mouse inputs. Pilots were able to adjust the display range of the CSD by using the mouse's scroll wheel within the boundary of the volumetric display or by using a range dial on an external menu. The CSD had a minimum display range of 10 nm and a maximum display range of 640 nm. Along with intruder and ownship information, the CSD provided pilots with range rings and heading markers along the edge of the display. For the purposes of this experiment, the CSD was restricted to a 2D top-down view and an ego-centric, north-up orientation.

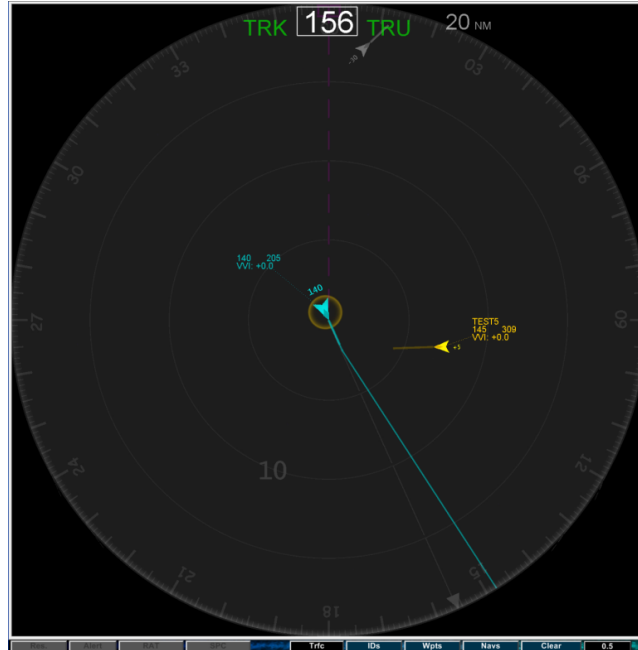


Figure 18. The Cockpit Situation Display (CSD) developed by the Flight Deck Display Research Laboratory.

Detect and Avoid System

The DAA system utilized in this experiment was simulated using a software architecture called the Java Architecture for DAA Modeling and Extensibility (JADEM; Santiago, Abramson, Refai, Mueller, Johnson, & Snow, unpublished). JADEM is capable of modeling various components of the DAA system, including surveillance functions, like detect and track; alerting functions, such as evaluate, prioritize, and declare; and maneuver recommendation and determination functions. These DAA system functions provide the information elements, alerting, and maneuver guidance that is presented on the DAA traffic display. For a more complete description of the capabilities provided by JADEM see Santiago et al., unpublished.

Surveillance

The software was configured to replicate Automatic Dependent Surveillance – Broadcast (ADS-B) surveillance, with a lateral sensor range of 80 nm and a vertical sensor range of +/- 5000 ft. Aircraft outside of this range were not displayed to the pilot.

Alerting

JADEM contained conflict detection logic that evaluated and prioritized surrounding traffic according to their predicted threat level with ownship. To calculate the predicted threat level, JADEM compared ownship's known intent to surrounding traffic, extrapolating the intruder's future position assuming their constant velocity. The spatial and temporal thresholds used to determine threats, along with their associated threat alerting levels as detailed in Table 5.

Maneuver Guidance

JADEM provided pilots in the advanced display conditions with text-based directive guidance (i.e., auto-resolutions) through its conflict resolution algorithm, Autoresolver-AD (Santiago et al., unpublished). Autoresolver-AD presented pilots with a recommended maneuver that was calculated as having the least amount of delay for resolving the active threat of all possible maneuvers (vertical or horizontal). If Autoresolver-AD was unable to generate a threat-free maneuver, the algorithm presented pilots with a maneuver that maximized the horizontal miss distance between ownship and the intruder. If at any point Autoresolver-AD computed a more effective (i.e., safer or more efficient) maneuver, pilots were given the ability to replace the previous

recommendation with the latest solution. Autoresolver-AD was disabled in the basic display conditions.

Traffic Simulation

The Multi-Aircraft Control Station (MACS) provided the air traffic simulation environment for this study (Prevot, 2002). MACS was used to generate simulated traffic targets, the confederate controller's display, and two pseudo pilot stations. An en route air traffic control (ATC) display provided the confederate controller with the ability to realistically manage all traffic within the designated experimental sector (Oakland Center ZOA 40/41). Two pseudo pilot stations enabled confederate pilots to take control of, and respond as, any manned aircraft in the simulated airspace. All experimental participants communicated over a single voice IP communication application.

Pilot Task

Pilots were tasked with operating a simulated MQ-9 Reaper, "HAWK21," along one of two pre-filed flight paths within Oakland Center airspace (ZOA 40/41). Pilots flew under instrument flight rules (IFR) and were responsible for navigating the aircraft and responding to a variety of scripted health and status tasks. These secondary tasks included responding to requests for status information (e.g., current fuel level) in a chat client and completing electronic checklists in response to aircraft system malfunctions. Pilots were also instructed to monitor their traffic display for potential safety of flight concerns. If a safety of flight threat was discovered, they were instructed to coordinate a maneuver with ATC (time permitting) and to upload the maneuver to the aircraft. They were encouraged

to minimize their deviation from the flight plan and to coordinate a return to their mission route and/or altitude with air traffic control as soon as practical.

Scenarios

Pilots flew two different mission routes, a “Fire Line” mission and a “Coastal Watch” mission. Each route started with HAWK21 already at mission altitude, flying towards its second programmed waypoint. There were no scripted altitude changes for either of the mission routes. The Fire Line mission route was level at 12,000 ft and Coastal Watch mission route remained level at 14,000 ft. Two different manned traffic scenarios were scripted to run alongside the two mission routes. Both traffic scenarios were developed by an ATC subject matter expert and designed to provide equivalent pilot workload. In each scenario, eight intruders were scripted to progress to a self separation then collision avoidance alert, absent of pilot action, while four different intruders were scripted to progress only to a preventive alert. All encounters were built for a single intruder, however, dynamic changes to the surrounding traffic made it possible for multiple intruders to occur simultaneously. Live traffic data was referenced in order to help simulate traffic patterns and densities that were representative of a busy day at Oakland Center.

Procedure

Training

Participants first completed an informed consent form and a demographics form, which elicited information regarding their experience in manned and unmanned aviation. This was followed by a brief overview of the day’s schedule and an introduction to the

pilot tasks. Pilots then underwent extensive training on the basic functionality of VSCS. This included practice on how to use the TSD's vehicle control interfaces as well as how to perform the various health and status tasks that would be present during the experimental trials. Pilots concluded this portion of the training with a 20 min practice scenario. Pilots received hands-on training and completed additional 20 min practice scenarios prior to all subsequent display conditions.

Experimental Trials

Participants completed four, 37 min experimental trials. All participants received the four different DAA display conditions described above: Basic Standalone, Basic Integrated, Advanced Standalone, and Advanced Integrated. The presentation of the display conditions was counterbalanced across participants to account for order and learning effects. Following each experimental trial, participants completed the NASA Task Load Index (TLX; Hart & Staveland, 1988) and a post-trial subjective questionnaire, which focused on the unique display elements of the preceding condition. A post-simulation questionnaire and debrief followed the final experimental trial.

Measures

All measured response and loss of well clear metrics detailed in Chapter 3 were collected.

Results

Measured Response

Each of the five measured response metrics listed above were analyzed utilizing a 2 (information level: basic, advanced) X 2 (display location: standalone, integrated) repeated measures Analysis of Variance (ANOVA) with Bonferroni pairwise corrections

for the main effect post hoc comparisons. An alpha level of .05 was used for all analyses. Significant interactions were analyzed using T test comparisons. Pilots responded to a total of 261 discrete alerts, 251 alerts appeared initially at the self separation threshold level, the remaining ten alerts appeared initially at the collision avoidance threat level. The means for each metric by display configuration is shown in Table 7.

Table 7. Measured response means by display configuration for Experiment 1.

	Notification Time	Clearance Approval Time	Proportion Uploads with Clearance	Initial Response Time	Initial Edit Time	Total Edit Time	Total Response Time
<i>Basic Standalone</i>	31.91s	5.24s	0.49	15.26s	13.02s	21.60s	38.68s
<i>Basic Integrated</i>	32.41s	5.04s	0.56	21.62s	17.11s	22.65s	44.86s
<i>Advanced Standalone</i>	25.91s	5.01s	0.56	18.30s	11.43s	16.28s	35.60s
<i>Advanced Integrated</i>	26.68s	5.21s	0.50	22.08s	5.51s	10.08s	32.35s
Grand Mean	29.23s	5.12s	0.53	19.32s	11.77	17.65s	37.87s

Notification Time

The main effect of information level on pilot notification time approached statistical significance, $F(1, 11) = 4.432, p = .059$. The advanced displays ($M = 25.99; SE = 2.78$) appeared to support faster pilot notification times than the basic displays ($M = 32.16; SE = 4.85$). There was no significant main effect of display location, $F(1,11) = .140, p > .05$, nor was there a significant interaction, $F(1, 11) = .018, p > .05$ (Figure 19).

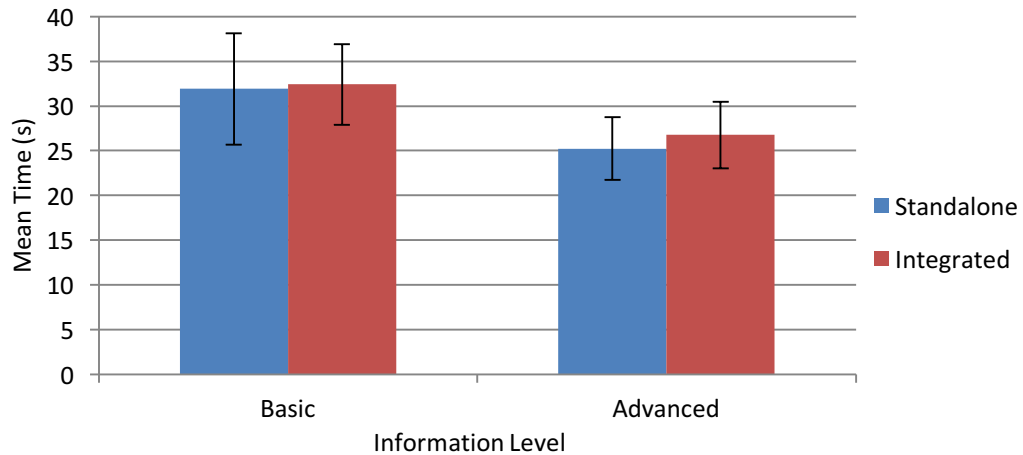


Figure 19. Notification time by information level and display location.

Clearance Approval Time

There was not a significant interaction of information level by display location for clearance approval time, nor were either of the main effects significant, $p > .05$. It took, on average, 5.12 sec for the confederate controller to begin transmission of a clearance approval after the beginning of pilots' request for one.

Proportion of Uploads with Clearance Approval

There was not a significant interaction of information level by display location for the proportion of maneuvers that received prior approval, nor were either of the main effects significant, $p > .05$. Across all displays, 53% of all maneuvers had clearance approval prior to the initiation of the maneuver.

Initial Response Time

There was not a main effect of information level on initial response time, $F(1, 11) = .459, p > .05$. However, the main effect of display location on initial response time

approached significance, $F(1, 11) = 4.635, p = .054$. Pilots' initial response time appeared faster for the standalone displays ($M = 16.78; SE = 1.80$) compared to the integrated displays ($M = 21.85; SE = 2.20$). The interaction between information level and display location was not significant, $F(1, 11) = .239, p > .05$ (Figure 20).




Figure 20. Initial response time by information level and display location.

Initial Edit Time

There was a significant interaction of information level and display location on pilots' initial edit time, $F(1, 11) = 13.851, p < .01$ (Figure 21). Post hoc pairwise T tests revealed a significant difference in initial edit time between the basic ($M = 17.11; SE = 3.11$) and advanced ($M = 5.51; SE = 3.42$) displays for the integrated display condition, $t(11) = 3.449, p < .01$. However, the difference between basic and advanced was not significant for the standalone condition displays, $t(11) = 1.126, p > .05$. The difference between basic and advanced for the integrated displays was large enough to result in a

significant main effect of information level on initial edit time, $F(1, 11) = 8.972, p < .05$. There was not a significant main effect of display location on initial edit time, $F(1, 11) = .139, p > .05$.

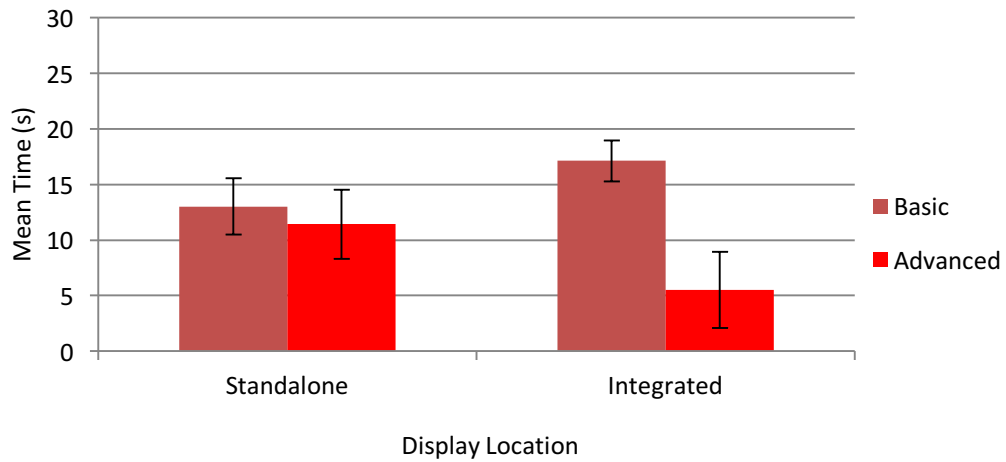


Figure 21. Initial edit time by display location and information level.

Total Edit Time

Information level had a significant effect on total edit time, $F(1, 11) = 11.821, p < .01$ (Figure 22). The advanced displays ($M = 13.18; SE = 2.48$) had significantly shorter edit times than the basic displays ($M = 22.12; SE = 3.14$). There was not a significant main effect of display location on total edit time, $F(1, 11) = 1.192, p > .05$, nor was there a significant interaction between information level and display location, $F(1, 11) = 2.804, p > .05$.




Figure 22. Total edit time by information level.

Total Response Time

There was a significant main effect of information level on total response time, $F(1, 11) = 6.619, p < .05$ (Figure 23). On average, the advanced display condition ($M = 33.98; SE = 3.34$) was 13.79 sec faster than the basic display condition ($M = 41.77; SE = 3.53$). The main effect for display location and the interaction were not significant, $F(1, 11) = .633$ and $F(1, 11) = 2.472$, respectively, $p > .05$.




Figure 23. Total response time by information level.

Loss of Well Clear

Due to a low number of losses of well clear across all display configurations, the proportion of losses of well clear were analyzed utilizing separate paired-samples t-tests for each independent variable (information level and display location). An alpha level of .05 was used for all analyses. Descriptive statistics and distributions are provided for the severity metric.

Proportion

The proportion of losses of well clear by information level and display location is shown in Table 8. There was not a significant difference in the proportion of losses of well clear by information level, $t(11) = 1.185, p > .05$ (Figure 24). There also was not a significant difference in the proportion of losses of well clear by display location, $t(11) = .225, p > .05$ (Figure 24). The proportion of encounters that resulted in a loss of well clear across all displays was 0.440, or 44%.

Table 8. Proportion of encounters that resulted in losses of well clear by information level and display location.

	<i>Basic</i>	<i>Advanced</i>	<i>Grand Mean</i>
<i>Standalone</i>	0.55	0.37	0.46
<i>Integrated</i>	0.49	0.28	0.42
<i>Grand Mean</i>	0.51	0.36	0.44




Figure 24. Proportion of encounters that became losses of well clear by information level (left) and display location (right). Reprinted with permission from Santiago & Muller (2015), *Pilot Evaluation of a UAS Detect-And-Avoid System’s Effectiveness in Remaining Well Clear*.

Severity

The means of the separation index across all four display configurations are shown in Table 9. The distribution (minimum, maximum, median, and 25th and 75th

percentiles) of severity indices by information level and display location is shown using box plots in Figure 25. Overall, Advanced Standalone had the highest separation index ($M = 1.42$; $SE = .256$) while Basic Integrated had the lowest separation index ($M = 1.08$; $SE = .054$). The mean separation index across all displays was 1.23.

Table 9. Mean separation index across all display configurations.

	<i>Basic</i>	<i>Advanced</i>	<i>Grand Mean</i>
<i>Standalone</i>	1.11	1.42	<i>1.17</i>
<i>Integrated</i>	1.08	1.19	<i>1.14</i>
<i>Grand Mean</i>	<i>1.10</i>	<i>1.41</i>	<i>1.23</i>

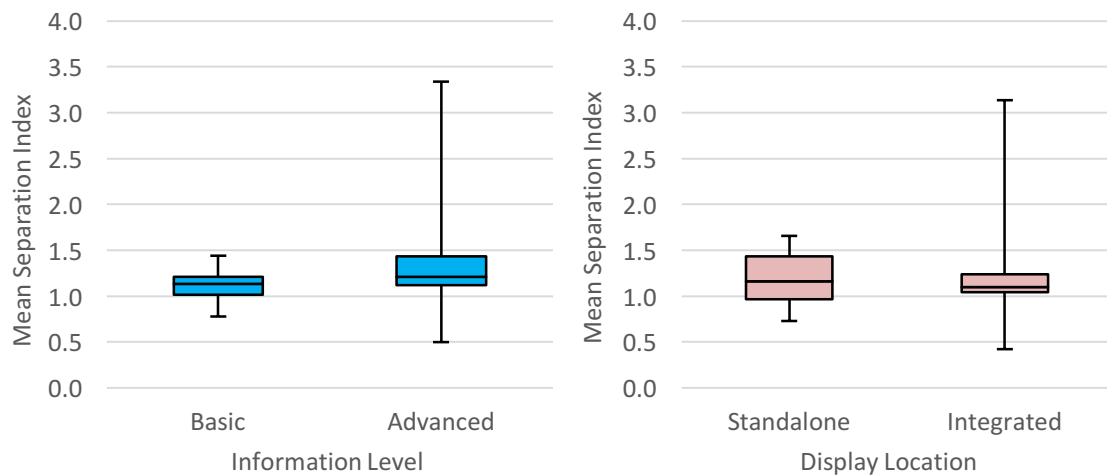


Figure 25. Separation index distribution by information level (left) and display location (right). The lower and upper boxes represent the bottom 25th and 75th percentiles, respectively, the median is indicated by the line intersection the boxes, and the minimum and maximum values are shown by the lower and upper whiskers.

Results Summary

The advanced information displays showed an advantage in terms of pilot performance over the basic, or minimum information, displays. Both the standalone and integrated versions of the advanced display condition showed significantly faster response times than the basic display conditions for two of the reported metrics: total edit time and total response time. In addition, the Advanced Integrated display supported significantly faster initial edit times than the Basic Integrated display. Finally, there was a trend toward lower proportions of losses of well clear for the advanced information level compared to based, which, while not statistically significant, provide converging evidence of the better performance of the two advanced displays compared to their basic counterparts.

Surprisingly however, there were no significant findings for display location; of the six response time metrics reported only one, initial response time, showed a near significant difference between standalone and integrated. In terms of the loss of well clear results, proportion and separation index appeared to be nearly identical across the two levels of this independent variable.

Chapter 5: Experiment Two

Overview

Experiment 2 followed immediately on the heels of Experiment 1, with the main goal of determining which of the Advanced Integrated display features contributed to better pilot performance. No major changes had been made within the SC-228 committee between Experiment 1 and 2; the committee was still working from the original detect and avoid (DAA) draft whitepaper published in November 2013 (RTCA, 2013a). At the time, the well clear separation threshold was being defined separately from SC-228 by the Sense and Avoid Science Research Panel (SARP), a group of key experts from government organizations (Cook, Brooks, Cole, Hackenberg & Rask, 2015), although no final decision had yet been made. Experiment 2 was a distributed simulation, utilizing facilities at both NASA Armstrong and NASA Ames to highlight the capabilities of NASA's UAS Integration into the NAS (UAS-NAS) project.

Purpose

Experiment 1 presented pilots with basic traffic display and advanced traffic display configurations. The advanced display configurations contained a suite of maneuver guidance tools that were meant to aid pilots in their determination of an appropriate resolution maneuver, whereas the basic displays lacked any maneuver guidance. The location of the traffic display was also manipulated, either integrating the

traffic information within the UAS moving map interface or situating the traffic information outside of the moving map in a standalone capacity. This resulted in a total of four different conditions for pilots: Basic Standalone, Basic Integrated, Advanced Standalone and Advanced Integrated.

Experiment 1 concluded that the advanced display configurations lead to the best performance overall with significantly shorter initial and total edit times, as well as significantly faster total response times, compared to the basic display configurations. Although the only statistically significant difference between the Advanced Standalone and Advanced Integrated displays was in initial edit time, the Advanced Integrated display trended toward shorter response times total edit times and total response times compared to the Advanced Standalone, and also had the lowest proportion of losses of well clear across all for displays compared. Thus, the Advanced Integrated was identified as being associated with the most acceptable pilot performance.

However, since the Advanced Integrated display configuration contained a suite of advanced information features and maneuver guidance, Experiment 1 could not provide conclusive input to specific minimum human machine interface (HMI) requirements for a DAA system. The goal of Experiment 2 was to systematically study the features of the Advanced Integrated display from Experiment 1 with the hope of determining the specific tool(s) that provide the greatest benefit to pilots, and therefore which HMI elements might constitute the minimum requirements for DAA.

Experimental Design and HMI Configurations

This study utilized a within-subjects, repeated measures factorial design to assess pilots' ability to maintain well clear across four different traffic display configurations: D1) Information Only (Info Only), D2) Information + Vector Planner Tools (Info + Vector), D3) Information + Auto-Resolutions (Info + AR), and D4) Information + Vector Planner Tools + Auto-Resolutions (Info + Vector + AR; Figure 26). The four displays provided pilots with different levels of maneuver guidance, with the Info + Vector + AR configuration being derived from the Advanced Integrated display in Experiment 1. In conditions where it was present, the maneuver guidance was integrated into the GCS command and control interfaces. Table 10 shows the alerting, information elements, and maneuver guidance for each display configuration.

Table 10. Alerting, information elements and display location for the four display configurations in Experiment 1.

	Info Only	Info + Vector	Info + AR	Info + Vector + AR
Alerting				
<i>Self Separation Alert</i>	✓	✓	✓	✓
<i>Predicted Collision Avoidance Alert</i>	✓	✓	✓	✓
<i>Collision Avoidance Alert</i>	✓	✓	✓	✓
Information Elements				
<i>Minimum Information Elements (Table 12)</i>	✓	✓	✓	✓
<i>CPA Location</i>	✓	✓	✓	✓
<i>Well Clear Ring</i>	✓	✓	✓	✓
Maneuver Guidance				
<i>Recommended Maneuvers</i>			✓	✓
<i>Trial Planner Tools</i>		✓		✓



Figure 26. The four DAA displays that were evaluated in Experiment 2: Info Only (top left), Info + Vector (top right), Info + AR (bottom left), and Info + Vector + AR (bottom right).

Alerting

Thresholds. The same alerting and well clear thresholds from Experiment 1 were used in Experiment 2. The well clear threshold was 0.8nm lateral, 400ft vertical, and 40 seconds to CPA. The DAA threshold for this experiment was 110 seconds, the time to

CPA of the self separation, and the predicted collision avoidance alerts, which alert to predicted losses of well clear. The same alerting structure was used for all four display conditions and included the predicted collision avoidance alert level.

Alert Levels. Table 11 presents the multi-level alerting structure that was used in Experiment 2. Alert level was based on the location of and time to CPA between the ownship and an intruder aircraft. Both the location and time thresholds had to be met for an intruder to be assigned a threat level. Location of the CPA was measured by both the lateral distance, in nautical miles (nm), and vertical distance, in feet (ft), to the predicted CPA location. The time to CPA was measured in seconds (sec). As in Experiment 1, the collision avoidance and well clear thresholds were treated as equivalent.

Table 11. The multi-level alerting structure used in Experiment 2 for all display configurations.

Alert/Threat Level	CPA Distance from Ownship		Time to CPA	Color
	Lateral	Vertical		
<i>Proximal</i>	> 2 nm	> 900 ft	N/A	Grey
<i>Preventative</i>	< 2 nm	< 900 ft	< 120 sec	White
<i>Self Separation</i>	< 1.2 nm	< 900 ft	< 110 sec	Yellow
<i>Predicted Collision Avoidance</i>	< 0.8 nm	< 400 ft	< 110 sec	Yellow, Red Border
<i>Collision Avoidance/Well Clear</i>	< 0.8 nm	< 400 ft	< 40 sec	Red

Information Elements

All display conditions in Experiment 2 contained the same baseline information elements, which was derived from the Advanced Integrated configuration in Experiment 1. In addition to the “minimum information” set of requirements (Table 12) from Experiment 1, the baseline information configuration included predicted location of CPA and well clear ring. Time to CPA and the vertical situation display were removed based on subjective feedback from Experiment 1; time to CPA was rated low on a subjective usefulness scale, and while the vertical situation display had some very high subjective ratings, it also had some of the lowest subjective ratings and was considered to add to the perception of clutter on the display. The Display 1 (Information Only) configuration was limited to this set of information elements.

Table 12. The minimum information set for the basic display condition.

Intruder Information	Visibility
Location	Always visible
Range	Always visible
Bearing	Always visible
Heading	Always visible
Relative Altitude	Always visible
Vertical Trend	Always visible
Heading Predictor	Always visible
Vertical Velocity	Within data tag
Absolute Altitude	Within data tag
Ground Speed	Within data tag
Aircraft ID	Within data tag

CPA Location. At the onset of a self separation alert, the predicted physical location of the intruder's CPA to ownship was depicted by a color-coded circle. In the Advanced Integrated condition, the ownship's predicted location at CPA could also be displayed by hovering over the aircraft icon of the potential intruder. The CPA location remained on the display as long as a threat was active and automatically disappeared once an alert was cleared.

Well Clear Ring: The presence of a self separation alert enabled the appearance of a "well clear ring" around ownship. With a radius of 0.8 nm, the ring gave pilots a visual reference to the lateral collision avoidance threshold.

Maneuver Guidance

The other three display configurations utilized in Experiment 2 contained suggestive and/or directive maneuver guidance in addition to the baseline information set described above. Suggestive maneuver guidance presents pilots with a range of potential solutions to a potential conflict, while directive guidance recommends a single maneuver. Suggestive guidance was provided in the form of vector planner tools and directive guidance was provided via auto-resolutions.

Vector Planner Tools

The Info + Vector and Info + Vector + AR displays contained the vector planner tools, which included two separate vector planning tools, lateral and vertical. The tools allowed pilots to determine the predicted threat level associated with various heading and altitude vectors by querying a conflict resolution algorithm prior to the upload of a maneuver (Autoresolver-AD; Santiago & Mueller, 2015). The lateral vector planning tool

was integrated into the TSD's autopilot interface (Figure 27). At the onset of a self separation alert the compass rose automatically opened on the TSD and a vector arrow appeared, extending from ownship to the edge of the compass rose in the direction of the recommended Autoresolver-AD maneuver. To use the lateral vector planner, pilots repositioned the vector arrow that changed color according to the threat level associated with the heading being tested.

The vertical vector planner tool resided within an altitude table that was permanently displayed on the far right side of the TSD. The altitude table consisted of five discrete altitude options displayed as buttons: one at the current altitude, two 1,000 feet above the current altitude (in 500 foot increments), and two 1,000 feet below the current altitude (also in 500 foot increments). When the vertical vector tool was engaged by the pilot (by hovering over it with the mouse), each altitude option in the table received a color-coded border according to its predicted threat level (Figure 27). Headings and altitudes that were predicted to lead to a collision avoidance threat were colored red, while those that were predicted to lead to a self separation or predicted collision avoidance threat were colored yellow. The tools turned white and green when the proposed vector was predicted to lead to preventive and proximal targets, respectively.

As the heading bug and vector arrow were moved around the compass rose, the selected associated heading would be automatically populated in the GCS's command and control auto-pilot interface. Similarly, when pilots clicked on an altitude button, that altitude populated the auto-pilot interface. It is important to note that the tools would only

allow a change to either the heading or the altitude at one time; if the lateral vector planner tool was engaged, the altitude field in the auto-pilot interface remained at current altitude and if the vertical vector planner tool was engaged, the heading field in the auto-pilot interface remained at the current heading. When satisfied with a lateral or vertical maneuver, the pilot could upload the maneuver by pressing “Send” within the auto-pilot interface. While the tools automatically appeared at the onset of a self separation alert or higher, the tools could also be launched manually in the absence of an active self separation alert.

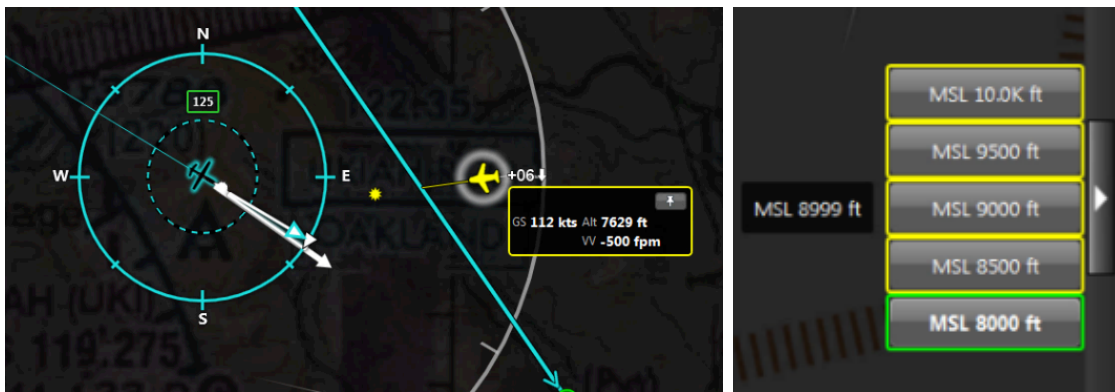


Figure 27. The lateral (left) and vertical (right) vector planner tools. The white arrow on the lateral vector planner tool indicates that a small right turn would change current conflict from a self separation threat to preventive threat. The green-bordered altitude box indicates that a descent to 8000 ft MSL would result in a conflict free trajectory; the yellow-bordered altitudes indicate that altitudes from 8500 to 10,000 ft MSL would result in a self separation alert level.

Auto-Resolutions

Info + AR and Info + Vector + AR added direct maneuver guidance, generated by Autoresolver-AD, to the feature set provided in the Information Only condition. Info + Vector + AR provided auto-resolutions in addition to the vector planner tools described above, while Info + AR provided only the information set and auto-resolutions. The directive maneuver guidance was presented via a text box in the upper right hand corner of the TSD, providing the pilot with a specific heading or altitude to fly in order to resolve a conflict (e.g., “Fly heading 343”, Figure 28). The recommended maneuver was also loaded into the auto-pilot interface. If pilots were comfortable with a maneuver suggested by Autoresolver-AD, they simply had to accept the maneuver to upload it to the aircraft.



Figure 28. The auto-resolution text box providing directive guidance, here, “Fly heading 343.”

Method

Participants

Nine active duty UAS pilots ($M = 46$ years of age) participated in the study. All pilots had military UAS experience (average of 1182 hours of military combat & non-

combat experience) and experience flying UAS in civilian airspace (average of 153 hours). A single retired air traffic controller served as a confederate.

Simulation Environment

The simulation components for Experiment 1 and Experiment 2 were largely the same, with few notable modifications. First, Experiment 2 was a distributed simulation that took advantage of the Live Virtual Constructive (LVC) Gateway (Murphy & Kim, 2013). In Experiment 2, the UAS GCS was located at NASA Armstrong Flight Research Center, along with simulated DAA system; the remaining simulation components were located at NASA Ames Research Center. Since Experiment 2 utilized different GCS hardware than was used in Experiment 1, the physical layout of the GCS displays was slightly different also, although the number of displays (four) remained the same. Second, since Experiment 2 utilized only an integrated DAA display, the health and status panels were allocated to both the lower left and right displays (in Experiment 1 the lower left display was reserved for the standalone display conditions). Finally, Experiment 2 simulated two types of aircraft based on different surveillance capabilities, cooperative and non-cooperative aircraft; in Experiment 1, only the former was simulated by the DAA system.

Ground Control Station

As in Experiment 1, participants in Experiment 2 interacted with the simulation software using desktop PCs and standard keyboard and mouse inputs. The UAS pilot participants were situated at a UAS GCS containing the Vigilant Spirit Control Station software (VSCS; Feitshans, Rowe, Davis, Holland & Berger, 2008). Figure 29 shows the

UAS GCS display set up. VSCS generated four separate displays: the Tactical Situation Display (TSD), an out-the-window view, and two side panels. The TSD served as the pilot's primary display, providing ownship and route information, traffic information, a moving map, and navigation/control interfaces. The out-the-window view provided pilots with synthetic terrain information and an integrated head-up display. The two side panels included health and status windows, a chat client, and an electronic checklist.



Figure 29. UAS GCS set up for Experiment 2; TSD, bottom center; out-the-window view, top center; health and status panel, bottom right and bottom left.

Detect and Avoid System

The Java Architecture for DAA Modeling and Extensibility (JADEM) was utilized in this experiment as the traffic surveillance, threat detection, and maneuver

guidance calculation system (Santiago, Abramson, Refai, Mueller, Johnson, & Snow, unpublished).

Surveillance

JADEM referenced a quantified definition of well clear in order to determine the associated threat level of any aircraft within range of a two simulated sensors, cooperative and non-cooperative, which had different ranges. The simulated cooperative sensor was based on Automatic Dependent Surveillance – Broadcast (ADS-B) and with a lateral range of 80 nm and a vertical range of +/- 5000 ft. The non-cooperative sensor range was based on the current state of the art onboard active radar with a lateral range of 6 nm, an azimuth of +/- 110 degrees (from the nose of the ownship), and an elevation of +/- 20 degrees (from horizontal).

Alerting

The multi-level alert structure (shown in Table 11) used the predicted lateral and vertical distance at CPA between ownship and the intruder and the time to CPA to determine the threat level for nearby aircraft. Predicted distance was calculated using ownship intent information and intruders' state information.

Maneuver Guidance

JADEM also contained Autoresolver-AD, a conflict resolution algorithm that supported the vector planner tools and auto-resolutions described above.

Traffic Simulation

The Multi-Aircraft Control Station (MACS) software suite was used to generate simulated traffic targets, controller displays, and two pseudo pilot displays (Prevot,

2002). The confederate controller managed all simulation traffic, including the simulated UAS, from their controller station. Pseudo pilots were likewise able to take control of, and respond as, any manned aircraft in the simulated airspace at their respective MACS stations.

Pilot Task

Pilots were tasked with operating a simulated UAS in Oakland Center airspace (ZOA 40/41) under Instrument Flight Rules (IFR). Participants were instructed to prioritize their responsibilities in the following way: 1) comply with ATC clearances and traffic display alerts to maintain well clear, 2) maintain the pre-approved course and altitude as much as practical, and 3) monitor and respond to secondary chat and health and status tasks. Pilots were told to coordinate any maneuvers around traffic with ATC, time permitting. VSCS emulated the performance characteristics of an MQ-9 (Reaper).

Scenarios

The “Fire Line” mission was retained from Experiment 1. Two versions of the “Fire Line” mission route were created with different altitudes. Both mission routes were contained entirely within Class E airspace. Each flight plan had an associated traffic scenario, which resulted in pilots receiving each route and its traffic scenario twice. The simulated traffic scenarios were scripted by an ATC subject matter expert and designed to reflect a busy, current day at Oakland Center. The traffic scenarios provided the manned background traffic that populated the experimental sector and included eight scripted encounters with the UAS with the intention of testing the pilot’s ability to respond to DAA alerts.

Procedure

Training

Participants completed an informed consent for minimal risk form and a demographic survey that elicited information about their manned and unmanned flight experience. Participants began with extensive training on the basic functionality of VSCS, receiving instruction on the TSD's vehicle control interfaces and how to successfully complete the secondary tasks. Pilots received additional, dedicated hands-on training, along with 20-minute practice sessions, on each display configuration prior to its associated experimental trial.

Experimental Trials

Participants completed a 37-minute experimental trial for each of the four display conditions. Over the course of the experimental trial, pilots were required to maintain well clear from scripted conflicts utilizing the tools available in the given display condition. The order of presentation of the four displays was counterbalanced across participants to account for order effects.

Measures

All measured response and loss of well clear metrics detailed in Chapter 3 were collected.

Results

Measured Response

The results that follow compare pilots' ability to respond to self separation and collision avoidance alerts with the four different display configurations: Info Only, Info + Vector Planner, Info + AR, and Info + Vector + AR. Each of the response time metrics

(except stage execution time) described above were analyzed using a one-way repeated measures Analysis of Variance (ANOVA). An alpha level of .05 was used for all analyses, with Bonferroni corrections made for pairwise comparisons. Means for all four displays across all measured response metrics (except stage execution time) are provided in Table 13.

Table 13. Measured response means by display configuration for Experiment 2.

	Notification Time	Clearance Approval Time	Proportion Uploads with Clearance	Initial Response Time	Initial Edit Time	Total Edit Time	Total Response Time
<i>Info Only</i>	13.68s	5.68s	0.42	7.83s	8.10s	8.91s	16.75s
<i>Info + Vector</i>	13.61s	5.65s	0.47	7.14s	9.04s	10.15s	17.29s
<i>Info + AR</i>	12.01s	5.67s	0.40	9.45s	2.76s	4.70s	14.15s
<i>Info + Vector + AR</i>	14.89s	5.26s	0.34	6.87s	4.34s	4.44s	11.52s
Grand Mean	13.54s	5.56s	0.41	7.98s	6.07s	6.94s	14.92s

Notification Time

Display configuration did not have a significant impact on notification times, $F(3,24) = 0.37, p > .05$. The mean notification time across all displays was 13.42 sec.

Clearance Approval Time

There was not a significant effect of display configuration on clearance approval time, $p > .05$. It took, on average, 5.56 sec for the confederate controller to begin transmission of a clearance approval after the beginning of pilots' request for one.

Proportion of Uploads with Clearance Approval

There was not a significant effect of display configuration on the proportion of maneuvers that received prior approval, $p > .05$. Across all displays, 41% of all maneuvers had clearance approval prior to the initiation of the maneuver.

Initial Response Time

Display configuration also failed to have a significant effect on pilots' Initial Response Times, $F(3,24) = 0.52, p > .05$. The mean initial response time across all display configurations was 7.98 sec.

Initial Edit Time

Display configuration did have a significant main on initial edit times, $F(3,24) = 11.51, p < .001$ (Figure 30). Info + Vector ($M = 9.04, SE = 0.97$) resulted in significantly longer initial edit times than both Info + AR ($M = 2.76, SE = 0.50$) and Info + Vector + AR ($M = 4.34, SE = 1.23$), p 's $< .05$. Initial edit times for Info Only ($M = 8.10, SE = 1.19$) were also found to be significantly longer than those seen for Info + AR ($p < .05$), while the difference between Info Only and Info + Vector + AR conditions approached significance ($p = 0.69$).




Figure 30. Initial edit times by display configuration.

Total Edit Time

Display configuration was also found to have a significant effect on total edit times, $F(3,24) = 4.50, p < .05$ (Figure 31). Info + Vector + AR condition resulted in significantly shorter total edit times ($M = 4.65, SE = 1.13$) than Info + Vector ($M = 10.15, SE = 1.19$), $p < .05$.




Figure 31. Mean total edit times by display configuration.

Total Response Time

There was not a significant effect of display configuration on total response times, although it did approach significance, $F(3,24) = 2.37, p = 0.10$ (Figure 32). Info + Vector resulted in the longest total response times ($M = 17.29, SE = 1.32$), while Info + Vector + AR resulted in the shortest ($M = 11.52, SE = 1.87$).

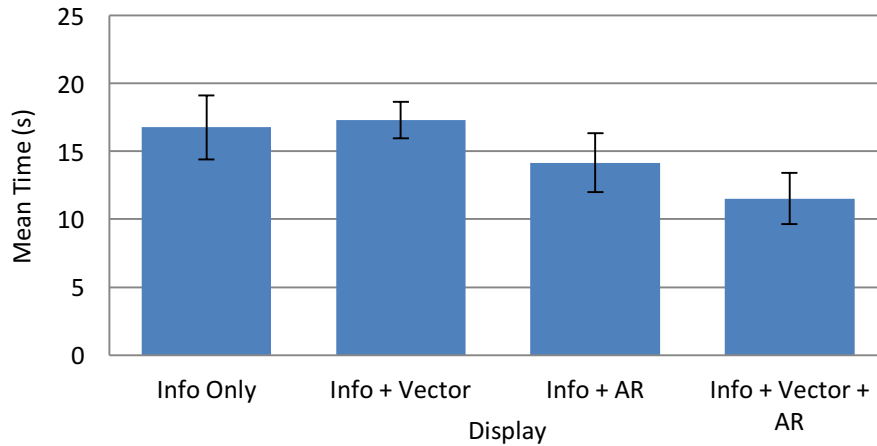


Figure 32. Mean total response times by display configuration.

Loss of Well Clear

The proportion of losses of well clear was analyzed using a one-way repeated measures Analysis of Variance (ANOVA) for display configuration (Info Only, Info + Vector Planner, Info + AR, and Info + Vector + AR). An alpha level of .05 was used for all analyses, with Bonferroni corrections made for pairwise comparisons. Due to the low numbers of losses of well clear, descriptive statistics only are provided for the severity metric.

Proportion

There was not a significant difference in the proportion of losses of well clear by display configuration, $F(3, 27) = 1.266, p > .05$ (Figure 33). The proportion of encounters that resulted in a loss of well clear across all displays was 0.130, or 13%. Table 14 shows the number and proportion of losses of well clear by display configuration.

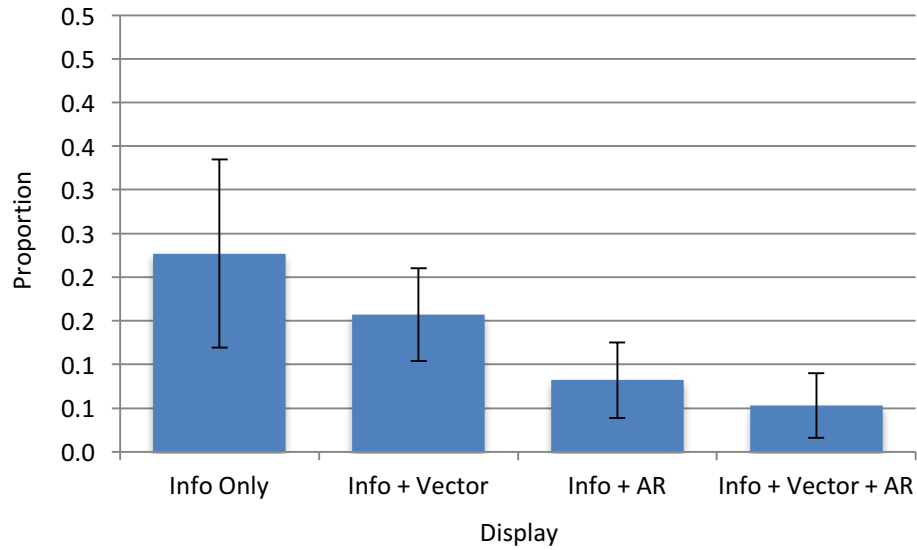


Figure 33. Proportion of encounters that became losses of well clear by display condition. *Reprinted with permission from Santiago & Muller (2015), Pilot Evaluation of a UAS Detect-And-Avoid System’s Effectiveness in Remaining Well Clear.*

Table 14. Number, proportion, and mean separation index of losses of well clear by display condition.

	<i>Info Only</i>	<i>Info + Vector</i>	<i>Info +AR</i>	<i>Info + Vector + AR</i>	<i>Grand Total</i>
Number of Losses of Well Clear	8	6	3	2	19
Proportion of Losses of Well Clear	0.227	0.157	0.082	0.053	0.130
Mean Separation Index	1.13	1.24	1.43	2.06	1.30

Severity

The means of the separation index across all four display configurations are shown in Table 14. The distribution (minimum, maximum, median, and 25th and 75th

percentiles) of separation indices by display type is shown using box plots in Figure 34. Overall, the Info + Vector+ AR had the highest separation index ($M = 2.06$; $SE = .149$) while Info Only had the lowest separation index ($M = 1.13$; $SE = .066$). The mean separation index across all displays was 1.30.

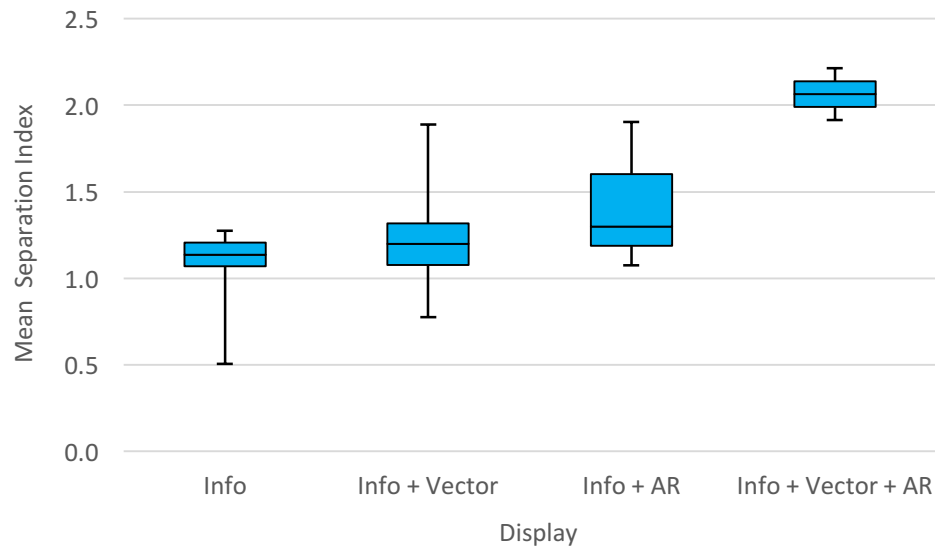


Figure 34. Separation index distribution by display configuration. The lower and upper boxes represent the bottom 25th and 75th percentiles, respectively, the median is indicated by the line intersection the boxes, and the minimum and maximum values are shown by the lower and upper whiskers.

Results Summary

This experiment failed to reveal significant differences in pilot performance across most of the measures analyzed. There as a significant difference in edit times between the two display configurations that included the auto-resolutions tool (Info + AR and Info + Vector + AR) compared to the two display configurations that did not contain this tool, however, these differences did not correlate with significantly faster total

response times overall. Despite a lack of significant differences, however, the measured response and loss of well clear results showed converging trends toward better pilot performance in maintaining well clear with the Info + AR and Info + Vector + AR displays, with the best performance observed for the latter display.

Chapter 6: Experiment 3

Overview

Data collection for Experiment 3 began in March 2015. A well clear separation threshold had been recommended by the Sense and Avoid Science and Research Panel (SARP) in August 2014, right after the start of planning for Experiment 3 (SARP; Cook, Brooks, Cole, Hackenberg & Rask, 2015). The quantitative well clear definition recommended by the SARP consisted of a modified tau³ value of 35 sec (to CPA) in the horizontal dimension with a horizontal miss distance (HMD) and a distance modifier (DMOD), both of 4000 ft. The distance threshold applies to the predicted location of CPA. The distance modifier provides a minimum distance from ownship for well clear given a slow closure rate. Thus, in order to cross the horizontal threshold of this well clear definition, an intruder must be less than 35 sec to CPA with the predicted CPA inside the 4000 ft threshold, or, the intruder itself is within the 4000 ft distance modifier. The SARP well clear definition also includes a fixed vertical separation threshold (for intruder location) of 700 ft. With this well clear formulation, depicted in Figure 35, a loss

³ TAU is time-based concept calculated as the slant range between aircraft divided by the rate of closure or range rate.

of well clear occurs when another aircraft penetrates both the horizontal and vertical thresholds.




Figure 35. *SARP well clear threshold recommendation. Reprinted with permission from: RTCA (2015c), DRAFT Minimum Operational Performance Standards (MOPS) for Unmanned Aircraft Systems (UAS) Detect and Avoid (DAA) Systems, Appendix C, Development of DAA Well Clear*

Subsequent to the SARP’s recommendation, the Federal Aviation Administration (FAA) issued a whitepaper expressing operational concerns about the 700 ft vertical separation requirement of the SARP’s proposed well clear definition, since the existing accepted separation between instrument flight rules (IFR) and visual flight rules (VFR) aircraft is 500 ft (Walker, 2014). In order to accommodate operational concerns, but also to maintain the integrity and vertical protection of the selected well clear threshold, the FAA whitepaper proposed to change the well clear definition to 450 ft vertical separation, but to add 250 ft protection in the vertical dimension with the use of a traffic alert or advisory in the DAA alert structure. This modified well clear definition from the FAA whitepaper is depicted in Figure 36. It is the definition that was adopted by SC-228 and became the basis for a new alert structure.




Figure 36. *FAA whitepaper well clear threshold recommendation. Reprinted with permission from: RTCA (2015c), DRAFT Minimum Operational Performance Standards (MOPS) for Unmanned Aircraft Systems (UAS) Detect and Avoid (DAA) Systems, Appendix C, Development of DAA Well Clear*






In addition to the new well clear definition adopted by SC-228, a number of other changes were made to the proposed alert structure (Table 15). The first major change to the alert structure was in the use of yellow and red intruder symbology. The use of a red outline around a yellow intruder icon, used for the predicted collision avoidance alert level in Experiments 1 and 2, was deemed non-standard and out of compliance with existing FAA guidelines that require the use of a yellow *or* red icon, e.g., Advisory Circular: Flightcrew Alerting [Department of Transportation (DOT), 2010]. According to these existing regulations, the use of yellow (caution) and red (warning) provide urgency information to pilots as well as the expected immediacy of actions. Advisory Circular 24.1322-1 (2010) defines a caution as “the level or category of alert for conditions that require immediate flightcrew awareness and a less urgent subsequent flightcrew response than a warning alert,” and a warning alert as “the level or category of alert for conditions that require immediate flightcrew awareness and immediate flightcrew response.” In order to comply with existing flightcrew alerting requirements for manned aviation, the

alert structure was updated to reflect expected pilot actions for each alert level (shown in the third column of Table 15).

The second major change to the alert structure was the addition of an alert level with the sole purpose of notifying the pilot of a need to take immediate action to avoid a loss of well clear. The addition of this alert level (DAA warning alert, level 4 in Table 15) was based on the comparison of total response time for maneuvers with and without an attempt to obtain a prior clearance from air traffic control (ATC). This data revealed that pilots took on average 11 sec to upload a final maneuver in response to a DAA alert when ATC was not contacted versus 19 sec when an attempt to contact ATC was made. The purpose of the DAA warning alert is to indicate to the pilot that they should maneuver immediately and then contact ATC as soon as possible afterward. The alert is triggered when there is a predicted loss of well clear with 25 sec or less until penetration. Thus, the DAA warning alert uses the same “predicted a loss of well clear” criteria, however, it shortens the time required to 25 sec to loss of well clear (roughly 60 sec to CPA using a 35 sec modified tau). This alert essentially allocates roughly 11 sec to pilot response time and 14 sec to aircraft maneuver time in order to avoid a loss of well clear. The alert symbology is red, which in compliance with the flightcrew guidelines, indicates that immediate response by the pilot is necessary. The corrective DAA alert also uses the same “predicted loss of well clear” criteria as the DAA warning alert but the time criteria is extended out further to 75 sec to loss of well clear, or approximately 110 sec to CPA, making it essentially equivalent to the predicted collision avoidance alert from

Experiments 1 and 2, with the clear pilot action to contact ATC to obtain a clearance prior to maneuvering.

Table 15. Proposed SC-228 MOPS alert structure.

Alert Level	Name	Pilot Action	Alert Time (time until penetrating separation criteria)	Symbology	Aural Alert Verbiage
4	DAA Warning Alert	<ul style="list-style-type: none"> • Immediate action required • Notify ATC as soon as practicable after taking action 	25 sec (tCPA approximate: 60 sec)		“Traffic, Maneuver Now”
3	Corrective DAA Alert	<ul style="list-style-type: none"> • On current course, corrective action required • Coordinate with ATC to determine an appropriate maneuver 	75 sec (tCPA approximate: 110 sec)		“Traffic, Separate”
2	Preventive DAA Alert	<ul style="list-style-type: none"> • On current course, corrective action should not be required • Monitor for intruder course changes • Talk with ATC if desired 	75 sec (tCPA approximate: 110 sec)		“Traffic, Monitor”
1	DAA Proximate Alert	<ul style="list-style-type: none"> • Monitor target for potential increase in threat level 	85 sec (tCPA approximate: 120 sec)		N/A
0	None (Target)	<ul style="list-style-type: none"> • No action expected 	N/A		N/A

A third change to the alert structure is the introduction of the preventive DAA alert (alert level 2). This alert level provides the vertical protection to the well clear threshold as described by the FAA whitepaper on well clear (Walker, 2014) and depicted in Figure 36. The preventive DAA alert has a vertical threshold of 700 ft, and it’s primary

purpose is to alert the pilot to aircraft that currently have achieved the acceptable 500 ft separation, but for which a sudden change in altitude by either the ownship or the intruder could make a loss of well clear imminent. A 0.25 nm buffer was also added to the horizontal miss distance so that the preventive DAA alert provides both a horizontal and vertical buffer around a predicted loss of well clear. This alert level informs the pilot that an intruder is currently not predicted to result in a loss of well clear, but that it is close enough to the well clear thresholds that it should be monitored for changes in trajectory.

Another change to the alert structure was the removal of the collision avoidance alert level from Experiments 1 and 2. The purpose of this alert level in Experiments 1 and 2 was to indicate to the pilot that a loss of well clear had occurred, and to emulate when an independent collision avoidance system would take over. The use of this alert level was deemed unnecessary for two reasons. One being that a pilot should be acting to avoid a loss of well clear *prior* to the loss happening, and that an actual loss of well clear does not change the expected pilot action to maneuver immediately. The second reason the alert is considered unnecessary for DAA, is that in the first phase of the DAA MOPS, the installation of a separate, certified, collision avoidance system, like the Traffic alert and Conflict Avoidance System II (TCAS II), is considered optional. The functional purpose of the DAA system is to remain well clear of other aircraft; the addition of a separate collision avoidance system is anticipated to provide an additional layer of protection on top of the DAA system. As such, the DAA alerting should not include collision avoidance alert levels, however it will be required to accommodate, or interoperate, with the current collision avoidance system standard (i.e., TCAS II). If/when a collision

avoidance system is integrated with the DAA system, it will provide its own alert level(s) in addition to the DAA alert levels.

Finally, aural alerts were added to the alert structure for the preventive, corrective and warning alert levels. The verbiage was based on subject matter expertise within SC-228, particularly from human factors professionals from the FAA, as well as relevant literature and regulatory guidance and recommendations regarding aural alerts, e.g., National Transportation Safety Board, 1992.

Purpose

Experiment 1 manipulated the location of the DAA display (standalone and integrated) and the level of DAA information (basic and advanced) provided to the pilot. The result of Experiment 1 revealed both faster pilot response times and lower rates of losses of well clear with the advanced information displays, with the best performance observed when UAS pilots were presented with the Advanced Integrated display configuration. However, because of the mix of an additional alert level, multiple information elements, and both directive (auto-resolutions) and suggestive (vector planner tools) guidance, the source(s) of the improved pilot performance could not be identified. There were no significant performance differences found in Experiment 1 between the standalone and integrated display conditions.

The goal of Experiment 2 was to isolate each of the additional display elements in the Advanced Integrated configuration from Experiment 1 in order to be able to assess their contribution to pilot performance in maintaining well clear from other aircraft. In order to do this, four displays were compared: Info Only, Info + Vector (i.e., vector

planner tools), Info + AR (i.e., auto-resolutions), and Info + Vector + AR. The baseline Info Only condition contained three “advanced” features in addition to the minimum information set utilized in Experiment 1’s basic display conditions: the location of the intruder’s predicted CPA, the well clear ring (0.8 nm around ownship), and the predicted collision avoidance alert level. The Info + Vector display added suggestive maneuver guidance in the form of the vector planner tools that were utilized in the Advanced Integrated display from Experiment 1. The Info + AR display added directive maneuver guidance in the form of auto-resolutions, also utilized in the Advanced Integrated display configuration. Finally, the Info + Vector + AR display contained both the vector planning tools and the auto-resolutions, and was roughly the same as the Advanced Integrated display configuration. The results of Experiment 2 suggested that the directive guidance (i.e., auto-resolutions) contributed to better pilot performance. Though only a couple of response time metrics showed statistical significance, other response time metrics and the proportion of losses of well clear showed strong converging trends to better performance by the displays equipped with the auto-resolution features. Pilots spent roughly half the amount of time interacting with the GCS in these two and the loss of well clear rate dropped to up to half of what it was without the auto-resolutions. This improved performance was likely the result of the auto-resolution maneuver being auto-loaded into the vehicle’s control interfaces, which substantially reduced the amount of interaction time required by pilots to upload the maneuver to the aircraft in the cases where they were satisfied with the recommendation. This stands in contrast to the Info + Vector display condition, where pilots had to determine their own resolution and take additional

time (if so desired) to use the tools to ensure that the maneuver was safe, and then interact with the command and control interface to upload the desired maneuver to the aircraft. One of the conclusions from Experiment 2 was that forcing pilots to interact with maneuver guidance tools could erase the benefit of providing the in the first place. A conclusion based on the finding that the Info + Vector display had the slowest response times of all four display conditions as well as the highest rates of losses of well clear. The integrated nature of the directive maneuver guidance with its reduced GCS interaction time, along with its ability to circumvent the need for the pilot to determine their own maneuver, can therefore be understood as the primary factors behind the improved performance associated with the Advanced displays in Experiment 1.

Despite the improved pilot performance seen with the integrated directive maneuver guidance tools in Experiments 1 and 2, their contribution to the DAA was not as straightforward. One issue was whether integration of guidance tools with a GCS command and control interface would be considered a minimum requirement that manufacturers would have to meet. The goal of the MOPS is to establish a minimum performance standard that has been validated to meet a specified level of safety, not to prescribe design solutions on manufacturers. The integration of required guidance is likely to be at the level of manufacturer design discretion, and not at the level of specification in the MOPS. Given that the MOPS are unlikely to specify that maneuver guidance be integrated, it is unclear if any benefits would be found with guidance tools that are not integrated into the command and control interface of the GCS. Secondly, as was mentioned at the end of the previous chapter, concerns about the certification

requirements for an algorithm that generates a single maneuver recommendation that pilots are required, or highly likely (if not required), to follow, falls outside of the scope and the resources of the phase 1 DAA MOPS. Finally, although pilot performance with directive guidance was roughly 50% faster and resulted in 50% less losses of well clear than without the directive guidance, performance differences overall were not found to be statistically significant, nor possibly practically significant (total response time difference were less than 6 s between the worst and best performing displays). In addition, if directive guidance is definitely taken out as an option for the DAA system, at least in the near term, the suggestive guidance in the form of vector planner tools in Experiment showed worse performance than the Info Only display condition, though the difference was not statistically significant. Thus, the data from Experiments 1 and 2 do not definitively show that maneuver guidance constitutes a minimum requirement for the DAA human machine interface (HMI).

One of the main purposes of the current experiment was to evaluate various potential suggestive maneuver guidance tools or displays. In order to move beyond the concept of requiring pilot interaction with the tools, an alternate display form, “banding”, was tested. Banding is term for a display or algorithm approach in which bands of color are used to indicate safe and unsafe values or regions. In terms of cockpit displays of traffic information, the banding display that pilots are most familiar with is the resolution advisory display that is part of TCAS II, which was originally developed in the 1980s (DOT, 2011). This display uses colored bands to indicate on an instantaneous vertical speed indicator the vertical speed s to be avoided (red) and achieved (green). The original

“non-glass cockpit” resolution advisory display utilized the round-dial instantaneous vertical speed indicator display as shown in Figure 37. Newer implementations in glass cockpits use bands the vertical speed tape, sometimes with pitch cues as well (Figure 38).

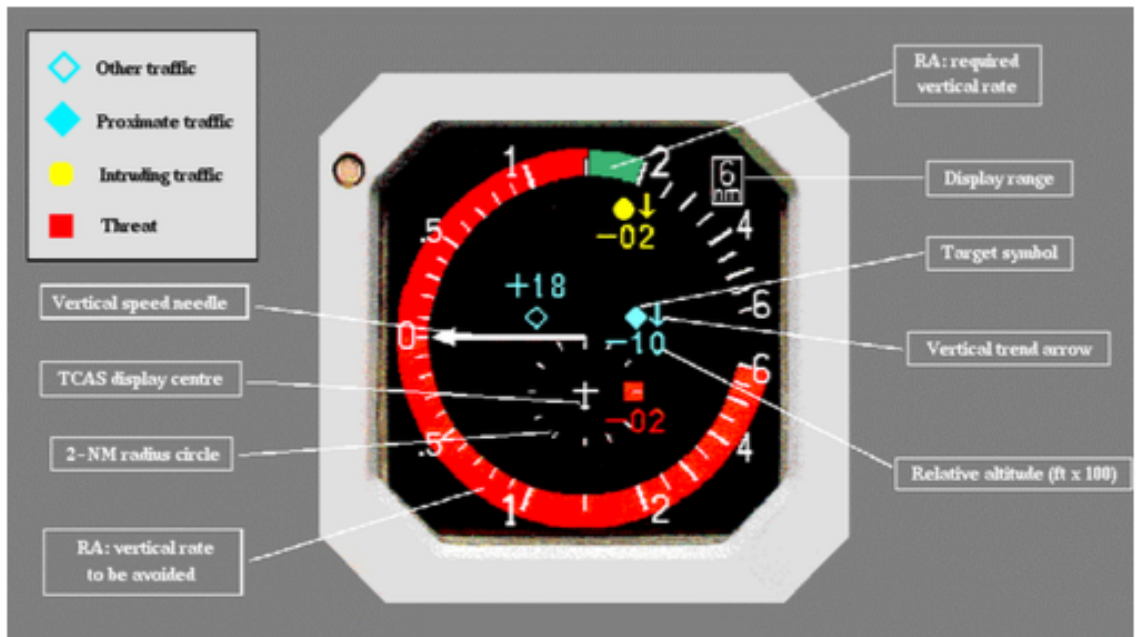



Figure 37. TCAS II resolution advisory display implemented on a round-dial instantaneous vertical speed indicator, with the required vertical rate indicated by a green “band” and the vertical rate to be avoided indicated by a red “band”. Reprinted from *DOT (2011), Introduction to TCAS II Version 7.1*.



Pitch Cue Implementation



Vertical Speed Tape Implementation

Figure 38. TCAS II resolution advisory display implemented on the primary flight display. The right shows the implementation of the red and green bands on the vertical speed tape. The left shows the pitch cue implementation (red pitch lines in the center of the attitude direction indicator). Reprinted from *DOT (2011), Introduction to TCAS II Version 7.1*.

In the early 2000s, researchers at different laboratories began working on various “free flight” concepts whereby separation assurance responsibility is moved from ATC to the pilots. Free flight concepts require an airborne separation assurance system that provides detection, prevention and resolution tools that can assist pilots in maintaining safe separation from other aircraft. Two laboratories in particular, one at NASA Langley Research Center in Hampton, VA, and one at the Dutch national aerospace laboratory (NLR), developed banding algorithms and displays as part of the suite of tools in their airborne separation assurance testbeds/prototypes. Figure 39 shows the conflict prevention heading and vertical speed bands that were developed at NASA Langley

Research Center as part of their prototype autonomous operations planner for Distributed Air/Ground Traffic Management project (Mondoloni, Palmer & Wing, 2002). Figure 40 depicts the conflict prediction bands developed at NLR as part of their predictive airborne separation assurance testbed (Hoekstra, van Gent & Ruigrok, 2002).



Figure 39. Illustration of NASA Langley Research Center's conflict prevention heading (right) and vertical speed (left) bands. Conflict prevention bands were generated by the prototype autonomous operations planner. Images reprinted from Mondoloni, Palmer & Wing (2002), Development of a Prototype Airborne Conflict Detection and Resolution Simulation Capability.

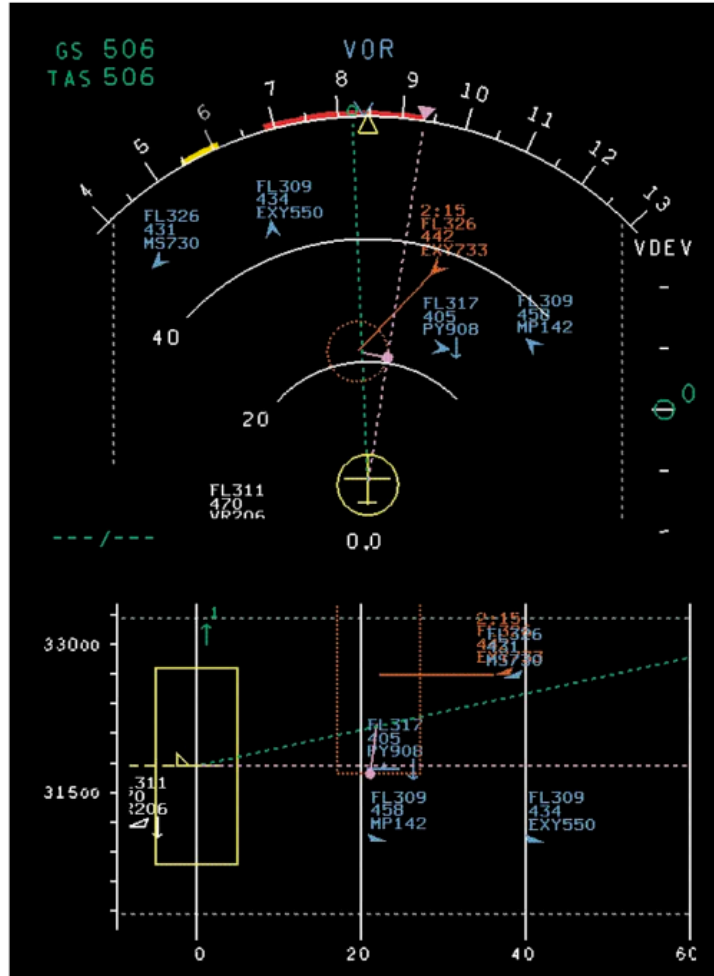


Figure 40. Illustration of NLR's conflict heading bands generated by the predictive airborne separation assurance system. Reprinted from *Hoekstra, van Gent & Ruigrok (2002), Designing for safety: the 'free flight' air traffic management concept.*

Not surprisingly, both NASA Langley and Dutch researchers have transitioned their banding concepts, and associated research, for use in UAS DAA systems, where pilots will be responsible for maintain separation assurance (i.e., well clear) under certain conditions. Recent research from both research teams on the application of conflict prediction and prevention bands can be found here: Chamberlain, Consiglio, Comstock,

Ghatas & Munoz, 2015; Munoz, Narkawicz, Hagen, Upchurch, Dutle, Consiglio, Chamberlain, 2015; Theunissen, Suarez & de Haag, 2013; Suarez, Kirk & Theunissen, 2012. The major advantage of these banding displays, especially when compared to the vector and trial planner tools evaluated in Experiments 1 and 2, is that the maneuver guidance information is readily available to pilots without requiring interaction with the display.

In this experiment three suggestive maneuver guidance displays were compared to a baseline information only display. Two of the suggestive maneuver guidance displays were based on the banding concept described above. The vector planner tools from Experiment 1 and 2 were also evaluated again, however they were modified so that they were decoupled from the vehicle's control and navigation interfaces. Experiment 3 also removed some of the advanced information features from all of the display concepts, such as the well clear ring and the intruder's CPA location, to better approximate minimum information set. The directive maneuver guidance was also removed from this display evaluation.

Experimental Design and HMI Configurations

This study utilized a within-subjects, repeated measures factorial design to analyze pilots' performance in maintaining well clear across four different DAA display configurations. The four displays with different forms of suggestive maneuver guidance were evaluated for their effect on pilot performance in maintaining well clear:

Information Only (Info Only), No Fly Bands, Omni Bands, and Vector Planning Tools.

The displays are shown in Figure 41. Unlike the suggestive and directive maneuver

guidance utilized in Experiments 1 and 2, the maneuver guidance utilized in the current experiment was de-coupled from GCS command and control interfaces.

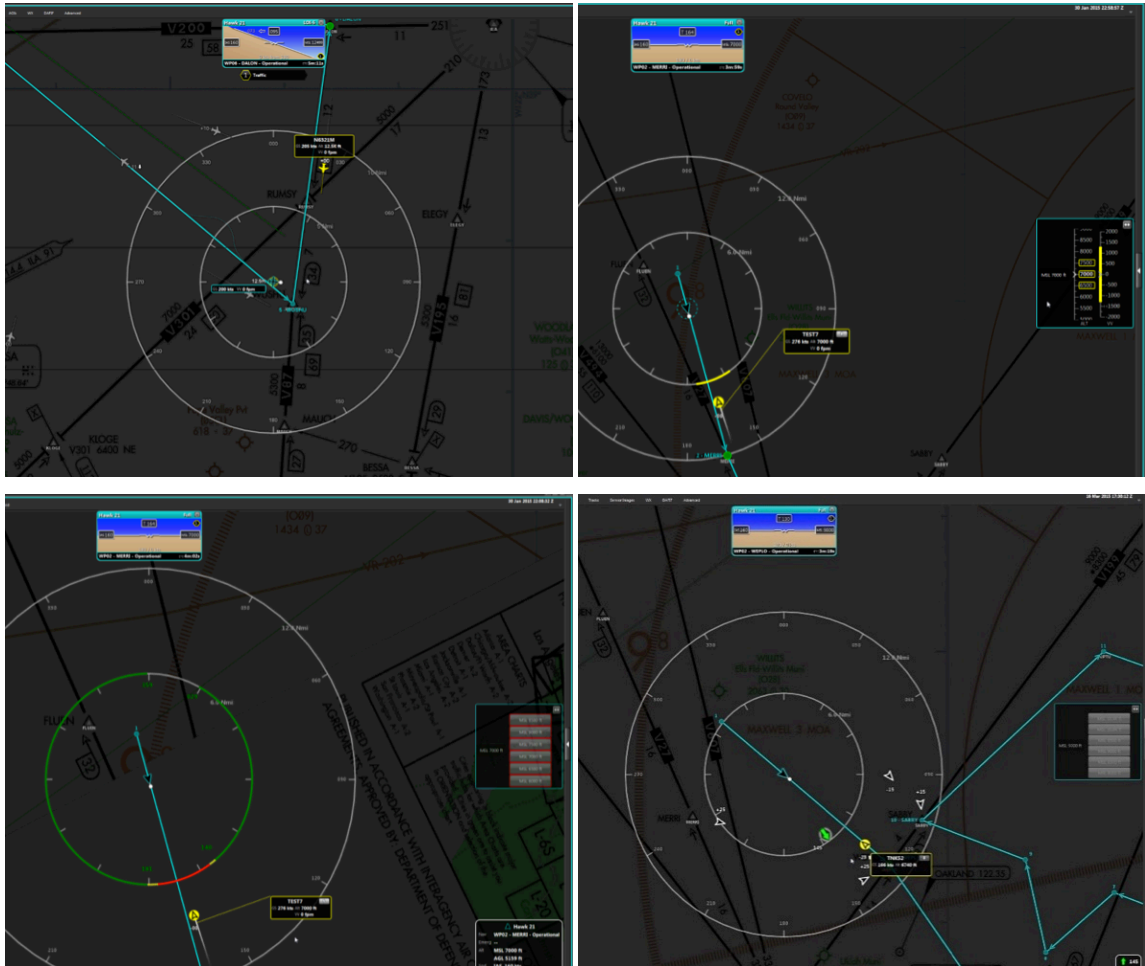


Figure 41. The four DAA displays that were evaluated in Experiment 3: Info Only (top left), No-Fly Bands (top right), Omni Bands (bottom left), and Vector Planning Tool (bottom right).

Alerting

Thresholds. The well clear definition adopted by SC-228 was used in Experiment 3 (Figure 36). Well clear was defined as 35 s to CPA modified tau in the horizontal






dimension with a 0.75 nm distance modifier (DMOD), 0.75 nm horizontal miss distance (HMD), and a vertical separation of 450 ft (ZTHR). The DAA threshold, employed for both the corrective and preventive DAA alerts, was 75 s to loss of well clear (approximately 110 s to CPA).

Alert Levels.

Table 16 presents the multi-level alerting structure that was used in Experiment 3. Alert level was based on predicted time to loss of well clear or a buffered well clear volume. The DAA warning and corrective DAA alerts were based on time to penetration of the well clear volume. The preventive DAA and DAA proximate alerts were based on time to penetration of a buffered well clear volume. Both the time and spatial thresholds had to be met for an intruder to be assigned a threat level. Location of the CPA was measured by both the lateral distance, in nautical miles (nm), and vertical distance, in feet (ft), to the predicted CPA location. The time to CPA was measured in seconds (s).

For three of the display configurations, Info Only, Omni Bands, and Vector Planner Tools, the alerting logic was generated by the JADEM DAA system, which utilized intent information for the ownship and state information for intruders in its trajectory predictions. The alerting logic in the No-Fly Bands display configuration was generated by the Stratway+ DAA system which utilized state information for both the ownship and intruders in its trajectory predictions.

Table 16. Multi-level alerting structure and separation criteria used in Experiment 3 for all display configurations.

Alert Level	Separation Criteria	Time to Loss of Well Clear	Symbology	Aural Alert Verbiage
DAA Warning Alert	modTau = 35 sec HMD = 0.75 nm DMOD = 0.75 nm ZTHR = 450 ft	25 sec		“Traffic, Maneuver Now”
Corrective DAA Alert	modTau = 35 sec HMD = 0.75 nm DMOD = 0.75 nm ZTHR = 450 ft	75 sec		“Traffic, Separate”
Preventive DAA Alert	modTau = 35 sec HMD = 1.0 nm DMOD = 0.75 nm ZTHR = 700 ft	75 sec		“Traffic, Monitor”
DAA Proximate Alert	modTau = 35 sec HMD = 1.5 nm DMOD = 0.75 nm ZTHR = 1200 ft	85 sec		N/A
None (Target)	Within surveillance field of regard	N/A		N/A

Information Elements

All of the display conditions in Experiment 3 contained the same baseline set of minimum information elements. This information set was reduced down from the ones used in the previous two experiments in order to better approximate a minimum set of information. The display requirements from the “Minimum Operational Performance Standards (MOPS) for Aircraft Surveillance Applications (ASA) System” (RTCA, 2011) was used as a reference to determine the minimum set of information elements. Table 17 shows a comparison of the information elements used across all three experiments. The

Info Only display configuration contained only these information elements and alerting, no maneuver guidance was provided.

Table 17. Comparison of the information elements used in Experiments 1, 2 and 3.

Intruder Information	Experiment 1 Basic Displays	Experiment 1 Advanced Displays	Experiment 2 All Displays	Experiment 3 All Displays
Location	✓	✓	✓	✓
Range	✓	✓	✓	✓
Bearing	✓	✓	✓	✓
Heading	✓	✓	✓	✓
Vertical Trend	✓	✓	✓	✓
Relative Altitude	✓	✓	✓	✓
Absolute Altitude	✓	✓	✓	✓
Aircraft ID	✓	✓	✓	✓
Ground Speed	✓	✓	✓	✓
Heading Predictor	✓	✓	✓	
Vertical Velocity	✓	✓	✓	
CPA Location		✓	✓	
Time to CPA		✓	✓	
Well Clear Ring		✓ (Integrated only)		
Vertical Situation Display		✓ (Integrated only)		

Maneuver Guidance

The three suggestive maneuver guidance display configurations differed only in the way that the guidance that was provided. Each is described in detail below.

No-Fly Bands

The No-Fly Bands display configuration provided pilots with vertical and horizontal suggestive maneuver guidance generated by the Stratway+ DAA system (Figure 42). While the Stratway+ software has been used previously to provide UAS pilots with DAA guidance information (Chamberlain et al., 2015), it had never before been integrated into the GCS software used in this experiment or assigned the alerting structure detailed in Table 16. Headings and vertical speeds that were predicted to lead to preventive DAA, corrective DAA or DAA warning alerts received amber bands, whereas headings and vertical speeds that would lead to no such alerts received no banding. The heading bands were presented within the inner range ring of the primary GCS display, while the vertical speed bands were presented within a vertical speed indicator located on the far right of the primary display. (The vertical speed indicator was not present in the other three configurations.) The No-Fly Bands were additive, allowing the guidance information to address multiple intruders or instances where a single intruder resulted in separate No-Fly regions. In situations where a loss of well clear could no longer be avoided, this display configuration added dashed-green “recovery bands”, which provided guidance that was calculated to result in the quickest resolution of the threat, despite the inability to avoid a loss of well clear.

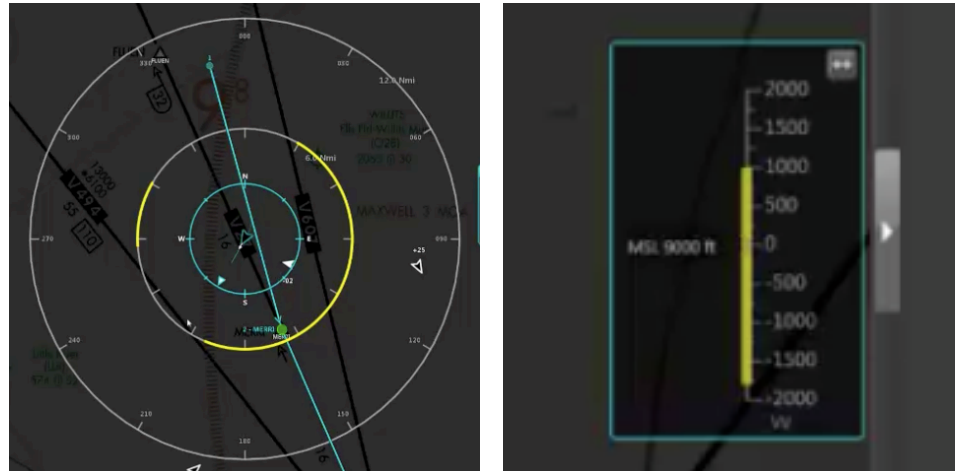


Figure 42. The No-Fly Bands suggestive maneuver guidance shows the pilot which headings and vertical speeds are to be avoided with yellow bands. The heading bands are shown on the inner range ring of the primary GCS display (left) and the vertical speed bands are down on the vertical speed indicator located on the right side of the primary GCS display (right).

Omni Bands

The Omni Bands display configuration also provided pilots with maneuver guidance bands in the horizontal and vertical dimensions, however they differed from the No-Fly Bands in several respects. First, the Java Architecture for DAA Modeling and Extensibility (JADEM; Santiago, Abramson, Refai, Mueller, Johnson & Snow, unpublished) DAA system instead of Stratway+ generated the Omni Bands. JADEM was used in Experiments 1 and 2 to generate the vector planner tools and auto-resolutions, and had been modified to output banding information for Experiment 3. Second, the vertical guidance was applied to specific altitude options, or “blocks,” instead of a vertical speed indicator. Altitude blocks were generated in 500 ft increments, and the altitude table displayed 1000 ft below current altitude and 1500 ft above current altitude.

Finally, the bands used a multi-color coding scheme modeled after the multi-level alerting structure. Whereas the No-Fly Bands used a single band color (amber) to indicate all threats predicted to be a preventive DAA alert or higher, the Omni Bands included a ‘safe’ banding color (green), as well as three other banding types to differentiate between preventive DAA (dashed yellow), corrective DAA (solid yellow) and DAA warning alerts. Figure 43 shows the Omni Bands display with the multi-colored bands. Like the No-Fly Bands, the Omni Bands were additive, allowing the guidance information to account for multiple or complex encounters.

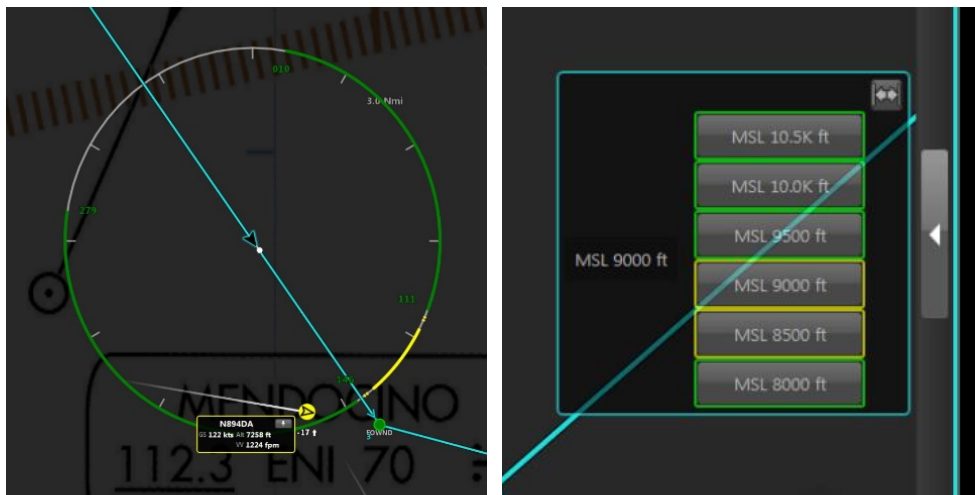


Figure 43. The Omni Bands maneuver guidance shows pilots which heading and altitudes are predicted be safe versus those that are predicted to lead to a conflict. The heading bands are shown on the inner range ring of the primary GCS display (left) and the altitude bands are shown on altitude blocks on the right of the primary display (right). The green bands indicate a range of heading and altitude options (e.g., 8000 ft) that are not predicted to lead to a preventive DAA alert level or higher.

Vector Planner Tools

The Vector Planner Tools display configuration departed from the No-Fly Bands and Omni Bands displays in its presentation of vertical and horizontal suggestive maneuver guidance. Similar to the vector planning tools used in Experiments 1 and 2, this configuration required the pilot to engage a vertical or horizontal planning tool in order to test individual heading or altitude options (Figure 44). To test heading options, pilots had to click and drag a dedicated arrow attached to the inner range ring. As soon as pilots moved the arrow, its color changed to reflect the predicted alert level of the probed heading. A text readout off the tip of the arrow displayed the precise heading being probed. Similarly, pilots had to select, by clicking, an individual altitude block from the altitude table to receive vertical guidance information. Individual altitude blocks were generated in 500 ft increments, and the altitude table displayed 1000 ft below current altitude and 1500 ft above current altitude. The Vector Planning tools were generated by the JADEM DAA system, and thus, the same color-coding scheme used in the Omni Bands condition was applied; the main difference between these two conditions was whether the maneuver guidance information was always available to the pilot on the display or whether the pilot had to engage the display manually to retrieve the information. When not in use, the arrow and altitude table were colored dull grey.

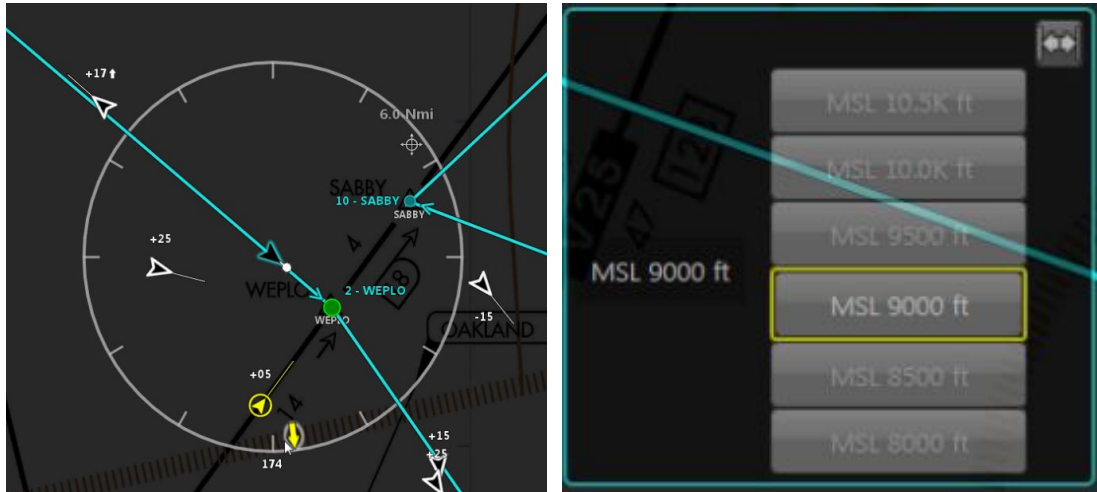


Figure 44. The Vector Planner Tools maneuver guidance shows the predicted safety level of individual heading and altitude options when probed by the pilot. The heading options are probed by dragging an arrow along the inner range ring of the primary GCS display (right). The altitude options are probed by clicking on individual 500 ft blocks in the altitude table located on the far right side of the primary display (left).

Method

Participants

Sixteen active duty UAS pilots ($M = 37$ years of age) were recruited for this experiment. Participants had an average of 1100 hours of experience flying UAS in military operations, and an average of 30 hours operating in civil airspace. Two active air traffic controllers from the National Air Traffic Controllers Association (NATCA) served as confederates for the study.

Simulation Environment

The simulation environment for Experiment 3 was very similar to that utilized in Experiment 2, except that the entire simulation was run at NASA Ames Research Center.

In addition, a new software package containing the Stratway+ DAA system was integrated into the Live Virtual Constructive (LVC) environment.

Ground Control Station

As in Experiments 1 and 2, Vigilant Spirit Control Station (VSCS), developed by the Air Force Research Laboratory (AFRL), served as the GCS for this study (Feitshans, Rowe, Davis, Holland & Berger, 2008). VSCS was configured to provide four different pilot displays: a Tactical Situation Display (TSD), two health and status panels, and a simulated out-the-window nose camera display (Figure 29). The TSD functioned as the pilot's primary display, providing all navigation and control interfaces, along with general ownership, routing and airspace information. The TSD also contained the traffic and DAA information, however, the amount and type of information it provided depended on the particular display configuration under test. Directly to the left and right of the TSD were two health and status panels. The left panel included a chat client and a history of system events, while the right panel included system information (e.g., subsystem status, telemetry data) and an electronic checklist. A fourth display positioned directly above the TSD provided video of a simulated nose camera, populated with synthetic terrain and an integrated head up display. Pilots interacted with VSCS using standard mouse and keyboard inputs.



Figure 45. Experimental ground control station. TSD, bottom-center; health and status panels, bottom-left and right; out-the-window view, top center.

Detect and Avoid System

Two different DAA systems were utilized in this study: JADEM and Stratway+⁴ (Santiago et al., unpublished; Chamberlain et al., 2015). JADEM was the DAA system utilized for the Info Only, Omni Bands and Vector Planner Tools display configurations. Stratway+ provided the DAA system for the No-Fly Bands display configuration. Both

⁴ The Stratway+ DAA software application has since been updated to the Detect and Avoid Alerting Logic for Unmanned Systems (DAIDALUS; Munoz, Narkawicz, Hagen, Upchurch, Dutle, Consiglio & Chamberlain, 2015).

DAA systems provided surveillance emulations as well as alerting and maneuver guidance logic.

Surveillance

Both systems were configured to emulate two different sensors onboard the UAS, one to detect cooperative traffic, and another to detect non-cooperative traffic. The simulated cooperative sensor used automatic dependent surveillance-rebroadcast (ADS-R)/TCAS-like ranges of 15 nm laterally and +/- 5000 ft vertically. (The lateral range was reduced from 80 nm in Experiments 1 and 2). The simulated non-cooperative sensor emulated a state-of-the-art onboard RADAR with a lateral range of 8 nm, an azimuth of +/- 110 deg and elevation of +/- 20 deg was simulated. (The lateral range was increased from 6 nm in Experiment 2). Traffic within range of these two sensor models was evaluated with respect to its predicted proximity to ownship's well clear volume.

Alerting

Traffic predicted to enter the well clear volume, according to an extrapolation of their current state assuming their constant velocity, received dedicated visual and aural alerts (Table 16). JADEM compared the extrapolated position of intruders to ownship's intent/trajectory, while Stratway+ compared intruders' future position to ownship's future state assuming its constant velocity. Additional alert levels were included to indicate proximate traffic.

Maneuver Guidance

Both JADEM and Stratway+ were capable of providing DAA suggestive maneuver guidance, as explained in the previous section. For detailed descriptions of how

the suggestive maneuver guidance is generated for displays see Santiago et al., 2015, Santiago et al., unpublished, Chamberlain et al., 2015, and Munoz et al., 2015.

Traffic Simulation

The Multi-Aircraft Control Station (MACS) software suite was used to generate simulated traffic targets, controller displays, and two pseudo pilot displays (Prevot, 2002). The confederate controller managed all simulation traffic, including the simulated UAS, from their controller station. Pseudo pilots were likewise able to take control of, and respond as, any manned aircraft in the simulated airspace at their respective MACS stations.

Pilot Task

Pilots were responsible for navigating a simulated MQ-9 Reaper, “HAWK21,” under instrument flight rules (IFR) within Oakland Center airspace (ZOA 40/41). Pilots flew one of two different mission routes while responding to a variety of scripted health and status tasks. These tasks included responding to requests for status information (e.g., current fuel level) in a chat client and completing electronic checklists in response to aircraft system malfunctions. Simultaneously, pilots were tasked with monitoring the DAA system for potential threats to well clear. Instructions were to coordinate a maneuver off of route with ATC (time permitting) in the event of a conflict requiring corrective action. While of lesser importance, pilots were encouraged to minimize their deviation from the flight plan and to coordinate a return to their mission route and/or altitude with ATC as soon as practical.

Scenarios

Two different mission routes were developed for this experiment, “Fire Line Low” and “Fire Line High.” Fire Line Low started the aircraft at 6,000 feet and included a temporary climb to 9,000 feet, with a subsequent return to 6,000 feet. Fire Line High inverted this pattern, starting the aircraft at 9,000 feet and then had the pilot descend to 6,000 feet before returning to 9,000 feet. Each mission route was paired with its own manned traffic scenario, both of which were developed by an ATC subject matter expert using real sector data and designed to provide equivalent pilot workload. Pilots saw each mission route, and thus traffic scenario, twice to account for all four display configurations. Each traffic scenario included six intruders that were scripted to lose well clear with ownship. Since the encounters occurred in a dynamic environment (the pilots’ response to a given encounter could not be predicted ahead of time with certainty), encounters occasionally had to be modified during runtime to ensure they were captured. While the scripted encounters only accounted for single-intruder conflicts, this activity during runtime meant multiple-intruder encounters occasionally occurred.

Procedure

Training

Participants first completed an informed consent form and a demographics form, which elicited information regarding their experience in manned and unmanned aviation. There were six distinct training sessions. The first training session focused on basic VSCS functionality, such as how to use the vehicle control interfaces, respond to chat messages, and complete the electronic checklists. A second training session introduced

the DAA system and focused on how to interpret traffic and alerting information. Pilots had to successfully complete a ‘training checklist’ before they could move on. The four display configurations received their own dedicated training session, which were administered immediately preceding the experimental trial. Display configuration training consisted of a hands-on demonstration as well as a 20-minute practice session where the participant used the configuration under test to fly a practice mission.

Experimental Trials

Participants completed a single 37-minute experimental trial for each of the four display configurations: Info Only, No-Fly Bands, Omni Bands and Vector Planning tools. The presentation of the display conditions was counterbalanced across participants to account for order and learning effects. Participants completed the NASA Task Load Index (TLX; Hart & Staveland, 1988) and a post-trial subjective questionnaire after each experimental trial. At the conclusion of the experiment, a post-simulation questionnaire was issued to participants and a debrief was conducted.

Measures

All measured response and loss of well clear metrics detailed in Chapter 3 were collected. An additional loss of well clear metric was collected: loss of well clear category. Losses of well clear were assigned one of four categories: ineffective maneuver, too slow, and too early return. An ineffective maneuver occurred when the pilot did not make a sufficient maneuver to avoid the loss of well clear, even though the guidance was correct (in the conditions that guidance was available). A loss of well clear was categorized as too slow when the pilot did not initiate the maneuver quickly enough

to avoid the loss of well clear. A too early return loss of well clear occurred when the pilot returned to course too soon and caused a loss of well clear with an aircraft that had been successfully avoided. The number of each type of loss of well clear was calculated.

Results

Measured Response

All of the measured response were analyzed across the four display configurations (Info Only, No-Fly Bands, Omni Bands, and Vector Planning Tools) utilizing a one-way repeated measures Analysis of Variance (ANOVA) with Bonferroni pairwise corrections for the main effect post hoc comparisons. Greenhouse-Geisser corrections are reported in cases where Mauchly's sphericity test was significant. An alpha level of .05 was used for these analyses. Descriptive statistics are provided for total response time by display configuration by alert type. Measured response means for all four displays are provided in Table 13.

Table 18. Measured response means by display configuration for Experiment 3.

	Notification Time	Clearance Approval Time	Proportion Uploads with Clearance	Initial Response Time	Initial Edit Time	Total Edit Time	Total Response Time
<i>Info Only</i>	18.46s	5.38s	0.56	12.45s	9.06s	14.72s	27.05s
<i>No-Fly Bands</i>	12.14s	5.10s	0.59	9.32s	8.42s	9.46s	18.78s
<i>Omni Bands</i>	16.02s	5.09s	0.59	11.28s	8.78s	9.61s	20.90s
<i>Vector Planner Tools</i>	17.97s	5.13s	0.48	16.35s	6.33s	8.12s	24.47s
Grand Mean	16.15s	5.17s	0.55	12.35s	7.94s	10.70s	22.80s

Notification Time

No significant main effect of display configuration was found on pilot's ATC notification times, $F(1.77, 26.51) = 1.355, p > .05$. Across all four displays, pilots took an average of 16.15 sec to notify ATC in response to a corrective DAA or DAA warning alert.

Clearance Approval Time

There was not a significant effect of display configuration on clearance approval time, $p > .05$. It took, on average, 5.17 sec for the confederate controller to begin transmission of a clearance approval after the beginning of pilots' request for one.

Proportion of Uploads with Clearance Approval

There was not a significant effect of display configuration on the proportion of maneuvers that received prior approval, $p > .05$. Across all displays, 55% of all maneuvers had clearance approval prior to the initiation of the maneuver. However, the inclusion of a warning alert prior to a loss of well clear in Experiment 3 allows for an additional look at pilots' rate of clearance approval in response to caution alerts separately from warning alerts. Table 19 and Figure 46 show the proportion of uploads with prior approval by alert type. While pilots only received prior approval 55% of the time overall, pilots received approval prior to maneuvering 72% of the time when only a corrective DAA alert was issued for an intruder (i.e., it did not progress to a DAA warning alert), compared to 15% of the time when only a DAA warning alert was issued for an intruder. This pattern is consistent across all display configurations.

Table 19. Proportion of uploads with prior clearance approval by alert type and display.

Alert Type	Info Only	No-Fly Bands	Omni Bands	Vector Planner Tools	Grand Mean
All	0.56	0.59	0.59	0.48	0.55
Only Corrective DAA (CORR)	0.72	0.70	0.75	0.72	0.72
Only DAA Warning (WARN)	0.17	0.25	0.00	0.20	0.15

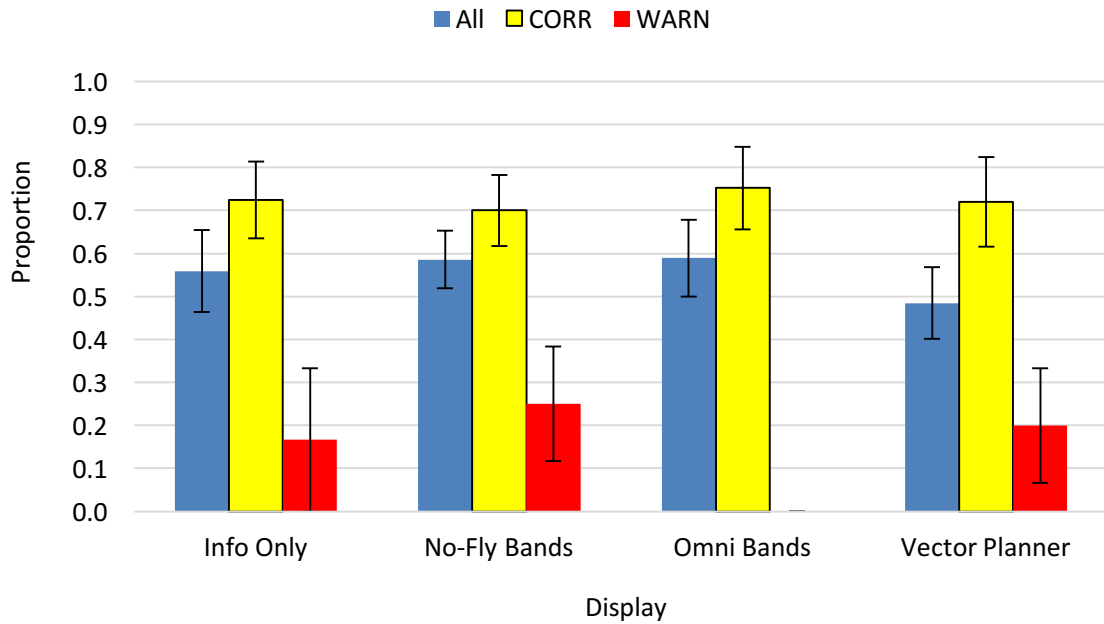


Figure 46. Proportion of uploads with prior clearance approval by alert type and display.

Initial Response Time

Display configuration was found to have a significant effect on pilots' initial response times, $F(3, 45) = 5.408, p < .01$ (Figure 40). Pilots' initial response times were fastest, on average, in the No-Fly Bands condition ($M = 9.32; SE = 1.47$), followed by the Omni Bands ($M = 11.29; SE = 1.69$) and Info Only ($M = 12.45; SE = 1.55$) displays. The Vector Planner Tools display ($M = 16.35; SE = 1.83$) resulted in the longest initial response times, which post-hoc analyses revealed to be significantly longer than those seen in the No-Fly Bands and Omni Bands displays (p 's $< .05$).

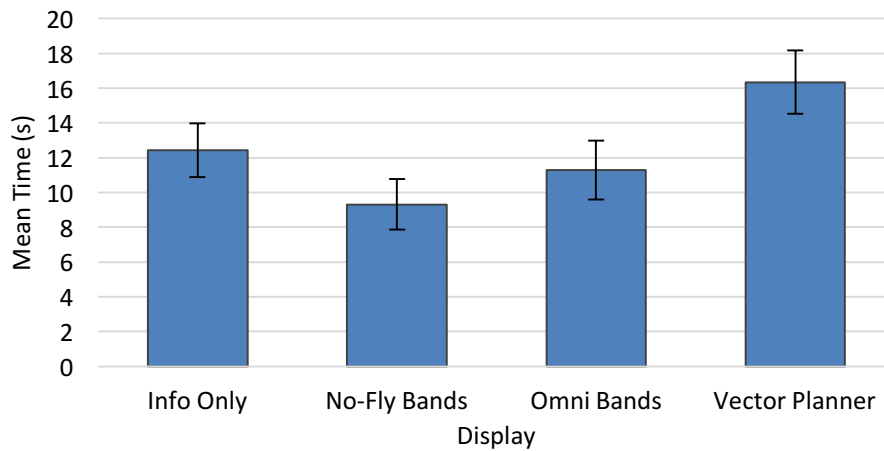


Figure 47. Initial response time by display configuration.

Initial Edit Time

No significant effect of display configuration was found on pilot's initial edit times, $F(2.13, 31.91) = 2.898, p > .05$ (Figure 31). Averaging across all four displays, pilots required 7.94 sec to upload their first maneuver to the aircraft.

Total Edit Time

Display configuration was found to have a significant effect on pilots' total edit times, $F(3, 45) = 5.713, p < .01$ (Figure 31). Pilots' total edit times were fastest, on average, in the Vector Planning Tools condition ($M = 8.12; SE = 0.88$), followed by the No-Fly Bands ($M = 9.46; SE = 0.98$) and Omni Bands ($M = 9.61; SE = 1.13$) displays. The longest total edit times were seen in the Info Only display condition ($M = 14.72; SE = 2.16$), which post-hoc analyses showed to be significantly longer than those seen in the Vector Planning Tools condition configuration ($p < .05$).

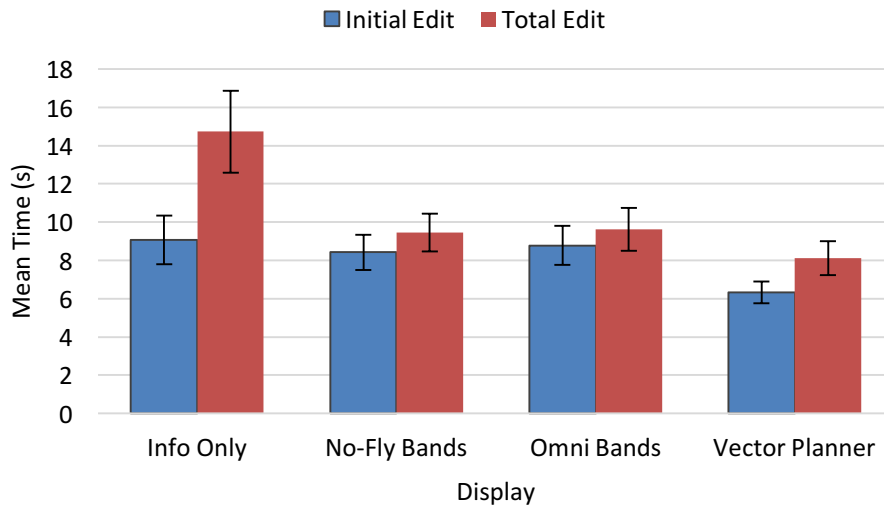


Figure 48. Mean initial and total edit times by display configuration.

Total Response Time

There was a significant effect of display configuration on total response times, $F(3, 45) = 3.744, p < .05$ (Figure 32). On average, the No-Fly Bands display ($M = 18.78$; $SE = 2.11$) resulted in the shortest overall response times, followed by the Omni Bands display ($M = 20.90$; $SE = 2.08$). The Vector Planner display ($M = 24.47$; $SE = 2.24$) and the Info Only display ($M = 27.05$; $SE = 3.09$) resulted in the slowest times. Post-hoc analyses revealed that the difference between the No-Fly Bands and the Info Only displays was the only comparison that approached significance ($p = .064$).

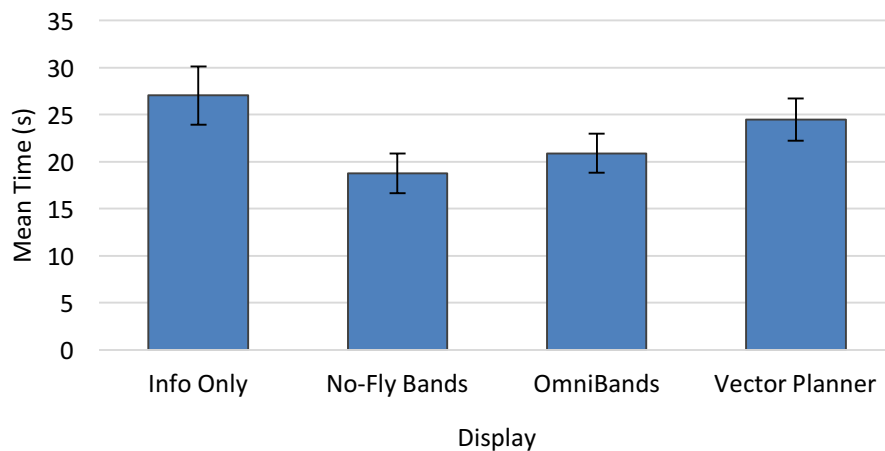


Figure 49. Mean total response times by display configuration.

The total response time means for alert type at first appearance (corrective DAA or DAA warning alerts) are shown in Table 20 and Figure 50. Pilots are expected to contact ATC for prior maneuver clearance for a corrective DAA alert, but not for a DAA warning alert.

Table 20. Total response time means and medians (in brackets) by display configuration for corrective DAA and DAA warning alerts.

Alert Type at First Appearance	Info Only	No-Fly Bands	Omni Bands	Vector Planner Tools	Grand Mean
Corrective DAA Alert	30.52s (22s)	19.35s (17.5s)	22.86s (20s)	28.18s (22.5s)	23.88s (21s)
DAA Warning Alert	12.57s (9s)	11.50s (10s)	13.24s (12s)	14.83s (10s)	13.29s (10s)

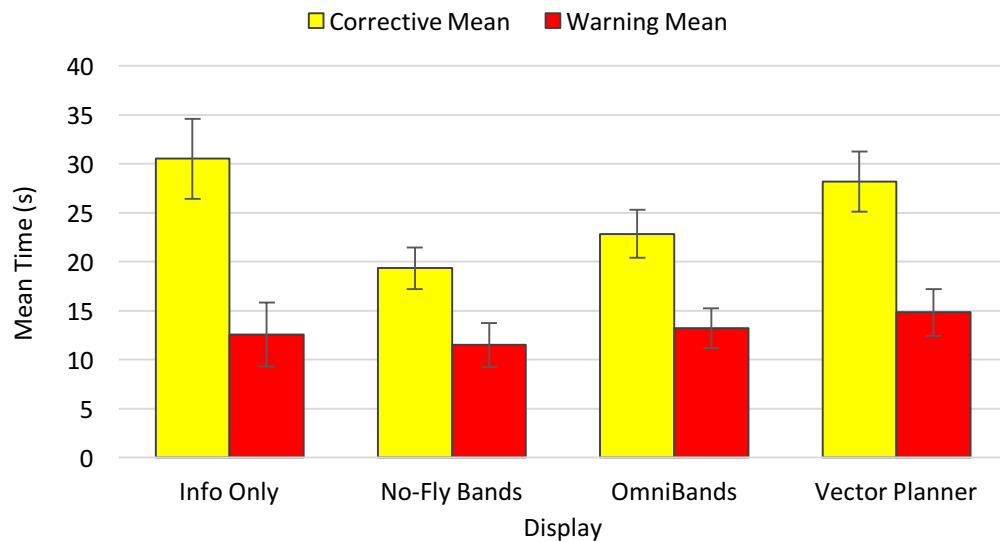


Figure 50. Total response time by alert type at first appearance, by display.

Loss of Well Clear

The proportion of loss of well clear results was analyzed across the four display configurations (Info Only, No-Fly Bands, Omni Bands, and Vector Planning Tools) utilizing a one-way repeated measures Analysis of Variance (ANOVA) with Bonferroni

pairwise corrections for the main effect post hoc comparisons. An alpha level of .05 was used for these analyses. Descriptive statistics are provided for the severity of losses of well clear by display configuration by alert type. The number of each category of losses of well clear are reported.

Proportion

The proportion of all encounters that were predicted to lose well clear and became losses of well clear by display condition is presented in Figure 51. There was a significant effect of display configuration on the proportion of predicted to actual losses of well clear, $F(3, 45) = 4.927, p < .01$. The Info Only display condition ($M = 0.087; SE = .023$) had a significantly higher proportion of losses of well clear compared to the Omni Bands display ($M = 0.000; SE = .000$), $p < .01$. No other differences between displays were significant.

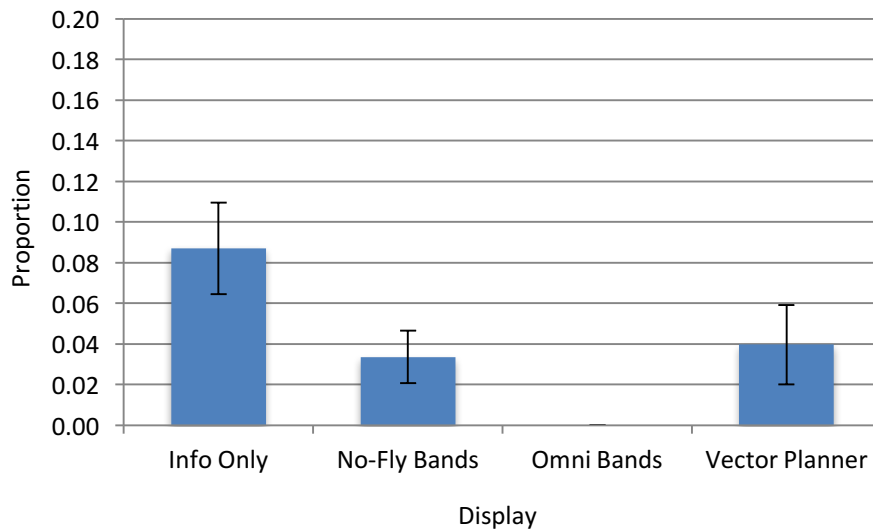


Figure 51. Proportion of encounters that were predicted to lose well clear that became actual losses of well clear by display condition.

Category

The number of each category of well clear was calculated by display condition (Table 21).

Table 21. Number of each category of loss of well clear by display condition.

	Ineffective Maneuver	Too Slow	Too Early Return	Total
Info Only	8	0	4	12
No-Fly Bands	2	1	2	5
Omni Bands	0	0	0	0
Vector Planner	3	2	1	6
Total	13	3	7	23

Severity

Separation index means and medians across all four display configurations are shown in Table 9. The distribution (minimum, maximum, median, and 25th and 75th percentiles) of separation indices by display type is shown using box plots in Figure 52. Of the three display configurations that had losses of well clear, the Vector Planner had the highest separation index ($M = 1.21$; $SE = .156$) while No-Fly Bands had the lowest average separation index ($M = 0.98$; $SE = .084$). The distributions show that the Info Only display had the overall minimum separation index of 0.24. The mean separation index across all displays was 1.08.

Table 22. Mean and median separation index across all display configurations.

	Info Only	No-Fly Bands	Omni Bands	Vector Planner	ALL
Mean	1.09	0.98	N/A	1.21	1.08
Median	1.13	1.00	N/A	1.33	1.06

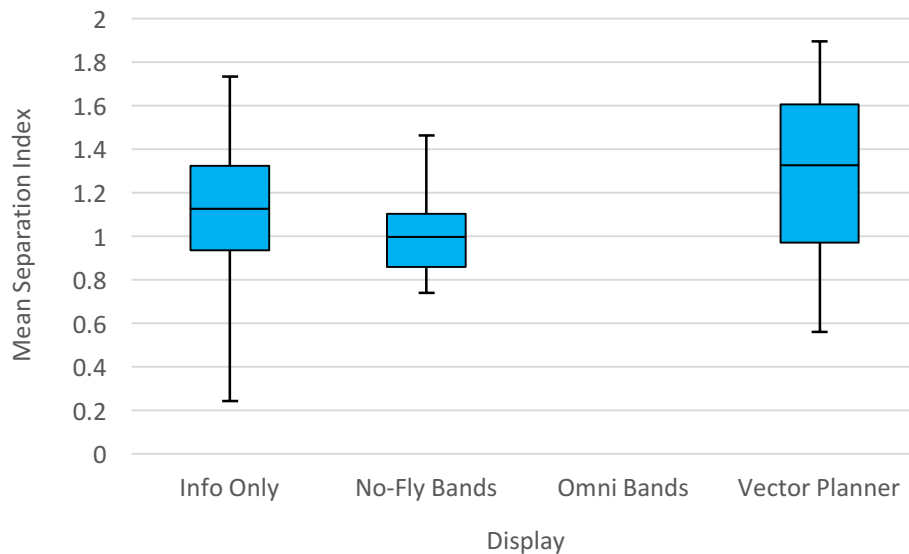


Figure 52. Separation index distribution by display configuration. The lower and upper boxes represent the bottom 25th and 75th percentiles, respectively, the median is indicated by the line intersection the boxes, and the minimum and maximum values are shown by the lower and upper whiskers.

Results Summary

The results of Experiment 3 showed significantly better performance for both of the two banding displays (No-Fly Bands and Omni Bands) compared to the Info Only display. The No-Fly Bands had significantly shorter overall total response times

compared to Info Only, while the Omni Bands had a significantly lower proportion of losses of well clear. Pilot performance with the Vector Planner display appeared to fall between the performance observed with the banding displays and the Info Only display. The only metric on which the Vector Planner display performed better a banding display was severity: the mean and median separation index for the Vector Planner appears to be higher than those for the No-Fly Bands (no losses of well clear occurred for the Omni Bands configuration).

Chapter 7: Research Results and Implications for HMI MOPS

The previous three chapters detailed the experimental design and results of three studies designed to evaluate the effect how various DAA HMI configurations affect the pilot performance on maintaining well clear of other aircraft. The goal of the larger research program within which these experiments took place, was to provide empirical data on which to establish the configuration of HMI components (traffic information, alerting, and guidance) that comprise the minimum requirement for DAA displays. Two categories of measures, measured response and loss of well clear, provide the basis from which to compare and evaluate the different prototype display configurations presented in the three experiments. The full set of results across all three experiments and the implications of those results for the DAA MOPS are presented in this chapter. The results are discussed first within each category of measures first, followed by an examination of how the two categories relate to, and complement each other. Finally, the results are linked to actual requirements in the current draft of the DAA MOPS.

Measured Response

The primary purpose of a DAA display for UAS is to enable pilots to carry out the three primary well clear functions: detect potential threats, determine appropriate resolution maneuvers, and execute those maneuvers via the GCS control and navigation interface. In addition, pilots must coordinate those well clear maneuvers with ATC prior

to execution under most circumstances. A number of discrete stages mark key pilot activities along the pilot-DAA timeline (Table 3). The eight stages map roughly, though not perfectly, to the three DAA functions a pilot is responsible for. The execute function has the clearest mapping to the final upload (T_{4b}). The detection function can be measured as some time after the first alert, T_0 , and before the next stage that the pilot executes, for example, notifying ATC of the need to maneuver (T_1) or initiating a maneuver response in the GCS (T_3 ; pilots did not always request a clearance prior to maneuvering). The determine function is the most difficult to map since pilots can be actively determining an appropriate maneuver any time between when the alert appears and the final maneuver is executed, essentially the entire timeline that the pilot is engaged in the DAA task. In fact, the stage execution time from the three reported experiments shows that pilots tend to initiate a maneuver change in the GCS prior to notifying ATC of their desire to maneuver, providing evidence that pilot tasks in support of executing the well clear function are being carried out concurrently rather than sequentially. Despite the challenge of directly mapping the eight timeline stages to the pilot's DAA functions, a number of pilot performance metrics can be extracted and provide insight to pilots' activities during these stages. The results from the three experiments detailed in the previous chapters show that some of these metrics are sensitive to various display configurations, especially differences in information across configurations. Each metric is discussed with the results from each experiment, below. The means for each for all three experiments are shown by display configuration in Table 23.

Table 23. Measured response means by display configuration for Experiments 1, 2 and 3.

		Notification Time	Clearance Approval Time	Proportion Uploads with Clearance	Initial Response Time	Initial Edit Time	Total Edit Time	Total Response Time
Exp. 1	<i>Basic Standalone</i>	31.91s	5.24s	0.49	15.26s	13.02s	21.60s	38.68s
	<i>Basic Integrated</i>	32.41s	5.04s	0.56	21.62s	17.11s	22.65s	44.86s
	<i>Advanced Standalone</i>	25.91s	5.01s	0.56	18.30s	11.43s	16.28s	35.60s
	<i>Advanced Integrated</i>	26.68s	5.21s	0.50	22.08s	5.51s	10.08s	32.35s
	Grand Mean	29.23s	5.12s	0.53	19.32s	11.77	17.65s	37.87s
Exp. 2	<i>Info Only</i>	13.68s	5.68s	0.42	7.83s	8.10s	8.91s	16.75s
	<i>Info + Vector</i>	13.61s	5.65s	0.47	7.14s	9.04s	10.15s	17.29s
	<i>Info + AR</i>	12.01s	5.67s	0.40	9.45s	2.76s	4.70s	14.15s
	<i>Info + Vector + AR</i>	14.89s	5.26s	0.34	6.87s	4.34s	4.44s	11.52s
	Grand Mean	13.54s	5.56s	0.41	7.98s	6.07s	6.94s	14.92s
Exp. 3	<i>Info Only</i>	18.46s	5.38s	0.56	12.45s	9.06s	14.72s	27.05s
	<i>No-Fly Bands</i>	12.14s	5.10s	0.59	9.32s	8.42s	9.46s	18.78s
	<i>Omni Bands</i>	16.02s	5.09s	0.59	11.28s	8.78s	9.61s	20.90s
	<i>Vector Planner</i>	17.97s	5.13s	0.48	16.35s	6.33s	8.12s	24.47s
	Grand Mean	16.15s	5.17s	0.55	12.35s	8.00s	10.70s	22.80s

Notification Time

Notification time is a measure of the time it takes for a pilot to notify ATC of a potential conflict and the need to execute a resolution maneuver. A pilot could notify the controller and request a maneuver *prior* to maneuvering, or notify the controller *after* maneuvering. Notification time was not significantly impacted by changes in DAA display configurations for any of the experiments, however, notification times were roughly twice as long in Experiment 1 compared to the latter two experiments. For all three experiments, notification time was, on average, longer than initial response time meaning that pilots typically began interacting with the GCS in order to initiate a resolution maneuver prior to notifying ATC. The likely reason for this is that pilots wanted to determine an appropriate maneuver before requesting a clearance. Other possible reasons are that pilots were too busy determining a maneuver to notify ATC, or that the communication channel was busy, causing delays in pilots' requests for clearances. In addition, in Experiment 3 pilots were instructed to maneuver immediately for a DAA warning alert and contact ATC as soon as practicable once the conflict was cleared, which would also contribute to longer notification times compared to initial response times.

Clearance Approval Time

Clearance approval time is a measure of the time that it takes for ATC to approve a pilot's request for clearance to maneuver in response to a potential loss of well clear. There was surprisingly very little variance in clearance approval time across display configurations within each experiments, or across the three experiments. Clearance

approval times ranged from a mean of 5.1 sec in Experiment 1 to 5.6 sec in Experiment 2. There were also no significant effects of HMI configuration on clearance approval time for any experiment. The lack of variability is likely due to the confederate role of the air traffic controllers who were instructed to approve all clearance requests unless they would impact the safety of the surrounding airspace. In real life operations with a busy controller and a busy sector, clearance approval is unlikely to be as consistent or as short.

Proportion of Uploads with Clearance Approval

The proportion of uploads with clearance approval tests a major operational assumption for UAS operations in the NAS. UAS pilots will be expected to obtain a clearance to maneuver off of their approved IFR flight plan except in the case of an emergency, such as a Traffic alert and Collision Avoidance System (TCAS II) resolution advisory (14CFR Part 91, §91.123; Code of Federal Regulations, 2004). There were no significant effects of display configurations on the proportion of time that pilots received a clearance prior to maneuvering. The percentage of maneuvers executed by pilots that had prior clearance by ATC ranged from a low of 41% (Experiment 1) to a high of 55% (Experiment 3). Although there is not objective standard these numbers can be compared against, they are likely below what would be considered acceptable in real life operations.

Though not analyzed for statistical significance, the expected pilot actions corresponding to corrective versus warning alerts appeared to have an impact on the proportion of maneuvers that had prior clearance in Experiment 3. While pilots only received prior approval 55% of the time overall, pilots received approval prior to maneuvering 72% of the time when only a caution alert was issued, compared to 15% of

the time when only a warning alert was issued. Thus in Experiment 3, the overall proportion of prior clearance approval is brought down by the expected low proportion for warning alerts, where pilots are trained to maneuver immediately and to wait until after maneuvering to contact ATC. This trend of higher proportions of clearances for correctives versus warnings persisted across all display configurations, which strongly supports the hypothesis that it is a result of the alerting structure, not differences in HMIs. The 72% clearance approval for the caution alerts is still not as high as would be desirable in normal operations, since the caution alert is designed with time built in for pilots to contact ATC prior to maneuvering, however, this could simply be a training and experience issue which would improve over time.

These results from Experiment 3 point to a potential explanation for why the proportion of maneuvers with prior clearance was low in the first two experiments: the ambiguity in the time until a loss of well clear in the alert levels. Both the self separation and predicted collision avoidance alerts thresholds used in Experiments 1 and 2 were active from 0 sec to up to 70 sec from the well clear thresholds, but pilots did not have a clear indication of exactly how much time they had to maneuver to avoid a loss of well clear (in the advanced display conditions they had a time to CPA indicator, which requires extrapolation of the well clear threshold to find time to loss of well clear). This lack of clear indication of the imminence of a loss of well clear – 0 sec versus 70 sec is a relatively large range – coupled with the newness of the DAA task and system for UAS pilots, may have caused pilots to be reluctant to wait for a clearance approval from ATC and to maneuver earlier than they needed to. In addition, Experiments 1 and 2 lacked

clear operational training in when pilots should contact ATC, relying on pilots to make a judgment as to whether a safety of flight concern existed (i.e., how much time they had until a loss of well clear). Thus, the ambiguity in the time to loss of well clear in the alert levels coupled with a lack of clear operational instructions likely created a situation where pilots where pilots traded off operational interoperability with ATC in order to be more conservative by maneuvering sooner than if they'd obtained a clearance first.

Initial Response Time

Initial response time is a measure of the time it takes for a pilot to initiate a maneuver in the GCS control and navigation interface after the appearance of an alert on the DAA display. This time includes the time it takes the pilot to detect a potential threat. In addition, it may include some time where the pilot is determining an appropriate resolution maneuver. Differences in initial response time could be due to differences in the time for pilots to detect a potential threat, differences in the amount of time spent determining a resolution maneuver, or differences in control and navigation interfaces (see Rorie & Fern, 2014 for a discussion of the effect of different control mode interfaces on pilot's measured response). Since each study utilized the same control and navigation interface for all four display conditions, potential differences would likely be a result of variation in the DAA HMIs.

The only significant effect of display configuration found across the three experiments in Experiment 3 where the two banding displays (No-Fly Bands and Omni Bands) both had significantly shorter initial response times compared to the Vector Planner display. The Vector Planner tool display led to an increase in initial response

times of 43% compared to the No-Fly Bands display and 31% compared to the Omni Bands display. This difference was likely due to the fact that pilots had to engage the vector planner tool in order to determine a resolution maneuver prior to initiating a maneuver into the control and navigation interface. In Experiments 1 and 2 pilots were able to use the control and navigation interface while simultaneously engaging the Vector Planner tools, thus the use of the guidance tools coincided with initiating a maneuver response. In Experiment 3, that mirroring was de-coupled, forcing pilots to first interact with the tools and then enter their desired maneuver into the GCS manually, leading to longer initial response times. This explains the observed difference in initial response time between the two vector planner display configurations in Experiment 2 and 3 (7 sec and 16 sec, respectively). The other three displays in Experiment 3, and particularly the No-Fly Bands and Omni Bands displays, which provided similar maneuver guidance, required no such manual interaction prior to making inputs into the vehicle control interfaces and subsequently led to shorter initial response times. Differences in initial response times in Experiment 3, however, did not translate to significant differences in total response times between the banding and Vector Planner displays (see Total Response Times discussion).

Overall initial response times decreased substantially between Experiment 1 and Experiment 2, and then increased slightly in Experiment 3 (grand means of 19, 8 and 12, sec respectively). It is not clear why there was a slight increase in initial response times between the latter two experiments, however, the decrease from Experiment 1 to 2 is a

trend seen across almost all of the measured response metrics which will be discussed in more detail in a later section.

Initial and Total Edit Times

Initial and total edit times reflect the amount of time it takes for a pilot to input and upload a first and final maneuver to the aircraft from the initiation of that input. In other words, it is the amount of time that pilots interact directly with the GCS control and navigation interface. Edit times capture most (if not all) of the time spent determining an appropriate resolution maneuver; however, the general lack of significant differences in initial response times across all experiments compared to frequently large observed significant differences in edit times suggests that pilots execute much of the determine function within the edit time (i.e., pilots begin interacting with the display consistently, regardless of configuration, but configuration affects the execution of the DAA tasks and therefore edit time). Initial and final upload times are captured because pilots sometimes make multiple uploads to the aircraft, using subsequent uploads to tweak or correct their initial upload.

In Experiment 1, pilots operating with the advanced displays had, on average, 9 sec shorter initial edit times, although the difference between the advanced and basic information display conditions was only significant for the integrated displays at the first upload, (5 and 17 sec, respectively – a 68% improvement for the Advanced Integrated display). The advanced displays were also roughly 9 sec, or 40% faster than the basic displays on total edit time, a difference that was significant for both display locations. Thus, it appears that the additional information and tools provided by the advanced

display conditions assisted pilots in more quickly determining and uploading a maneuver in response to displayed self separation and collision avoidance alerts. There were no significant differences in either initial or total edit between the standalone and integrated displays.

In Experiment 2, the HMI configurations that included auto-resolutions (Info + AR and Info + Vector + AR) consistently resulted in the fastest initial and total edit times, while the displays that lacked the tool (Info Only and Info + Vector) often led to the slowest edit times responses. The improvements were significant, with reductions in edit times ranging from 46% to 70%.

Unlike Experiments 1 and 2, there was not a significant difference in initial edit times. This finding suggests that pilots were consistent in uploading an initial maneuver regardless of the presence or type of DAA suggestive maneuver guidance provided by the GCS. The lack of a significant effect here is likely due to the removal of the integrated directive maneuver guidance (i.e., auto-resolutions) provided in the first two studies, which drastically reduced the amount of time pilots spent interacting with the GCS before making an upload. In those studies, pilots only had to accept an upload recommended by the DAA system in the cases where they were satisfied by the system's recommendation. By requiring manual input by the pilot in all of the display configurations in Experiment 3, the differences between pilots' initial edit times leveled off at an average of 8 seconds across displays.

While display configuration in Experiment 3 did not affect initial edit times, it did affect total edit times. The Info Only display lead to considerably longer total edit times

than were seen for the other three displays. The total edit times for the Information Only display were 36% longer than those for the No-Fly Bands display, 35% longer than the Omni Bands display, and 45% longer than observed for Vector Planner tools display (the only significant difference was with Vector Planner). Furthermore, the total edit times seen for the three suggestive maneuver guidance displays were very similar to their initial edit times, differing by between 1 and 2 sec. The total edit times for the Information Only display, conversely, were 6 sec longer than the initial edit times for that display, suggesting that pilots spent much more time modifying their initial upload to the aircraft when no suggestive maneuver guidance was present.

Advantages in edit times, however, did not necessarily equate to overall faster response times, as it only captures one stage of a pilot's functions to determine and then execute a resolution maneuver. Specifically, edit times only capture the time that pilots interact directly with the GCS, and not the time prior to interacting with the GCS which may be spent just visually examining the DAA display. A slower response earlier in the timeline could negate any benefits found in the edit times for a particular display configuration since total response time captures initial response time in addition to total edit time.

Total Response Time

Total response time captures the time that it takes for pilots to upload their final maneuver in response to a display alert, from when the alert first appears. Total response time captures all three well clear functions (detect, determine, execute) in addition to pilot interaction with ATC; thus total response time encapsulates the entire pilot-in-the-

loop activity of the DAA system. By its definition, total response times consist of all of the aforementioned measured response metrics, with the exception of notification times, which in some instances were not completed until after a maneuver had been made. Total response times therefore provide the best overall picture of the effect of display configuration on pilots' ability to comply with DAA display alerts in a timely manner. Thus, total response time is the response time performance metric most critical to the DAA MOPS development.

In Experiment 1 Pilots completed their DAA functions about 14 sec, or 33% faster, with the advanced information displays than with the basic information displays. The total response time results in this experiment followed the same pattern of results for initial and total edit times, whereby the advanced displays showed faster response times than the basic displays, by 9 sec on average. A 14 sec difference in total response time could have a significant impact on the potential for a loss of well clear, depending on the rate of closure between the ownship and another aircraft. As mentioned in previous chapters, pilots will potentially face constrained timeframes for responding.

Conversely, there were no significant differences in Experiment 2 in total response time by display configuration, although the results approached significance with shorter observed response times with Info + AR (14 sec) and Info + Vector + AR (12 sec) compared to 17 sec for both of the other two displays. The lack of significance in the total response time results is surprising given that the initial and total edit times did significantly differ by display configuration. The two displays that contained the auto-resolutions tool were roughly 4 to 6 sec faster on average compared to the Info Only and

Info + Vector displays. That difference was reduced to 3 sec for Info + AR and 5 sec for Info + Vector + AR for total response time.

For both of the first two experiments, the well clear threshold was set at 40 sec to CPA and the alerting threshold at 110 sec, leaving a maximum 70 sec for the pilot and aircraft to respond and maneuver in time to avoid a potential threshold violation. There are essentially two buffers on each side of a pilot's total response time prior to a well clear violation: 1) a buffer for when a potential threat is detected (i.e., some allowance for a later detection due to sensor performance, encounter geometry or pilot distraction), and 2) a buffer for aircraft performance (and lower restrictions on aircraft performance reduce costs on manufacturers and allow a wider variety of aircraft to meet defined standards and regulations). Shorter total response times for a pilot to respond to a DAA alert could allot for greater buffers for late detections and more limited aircraft performance.

Figure 53 and Figure 54 show the relationship between the total response time for each of the four display configurations in Experiment 1 and 2 relative to the entire DAA timeline. For these simplified depictions, all of the buffer time that is leftover from total response time is allocated to aircraft maneuver time. There is a clear and substantial reduction in total response times between Experiment 1 and Experiment 2, even for the two displays that were nearly identical between them [Advanced Integrated from Experiment 1 (32 s) and Info + Vector + AR from Experiment 2 (12 s)].

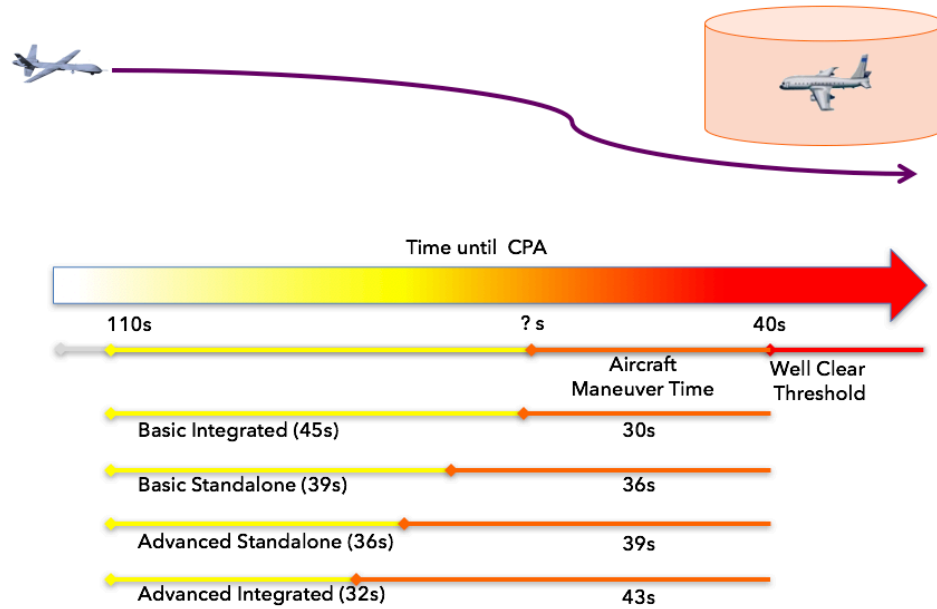


Figure 53. Relationship between total response time for all four displays evaluated in Experiment 1, and the entire DAA timeline.

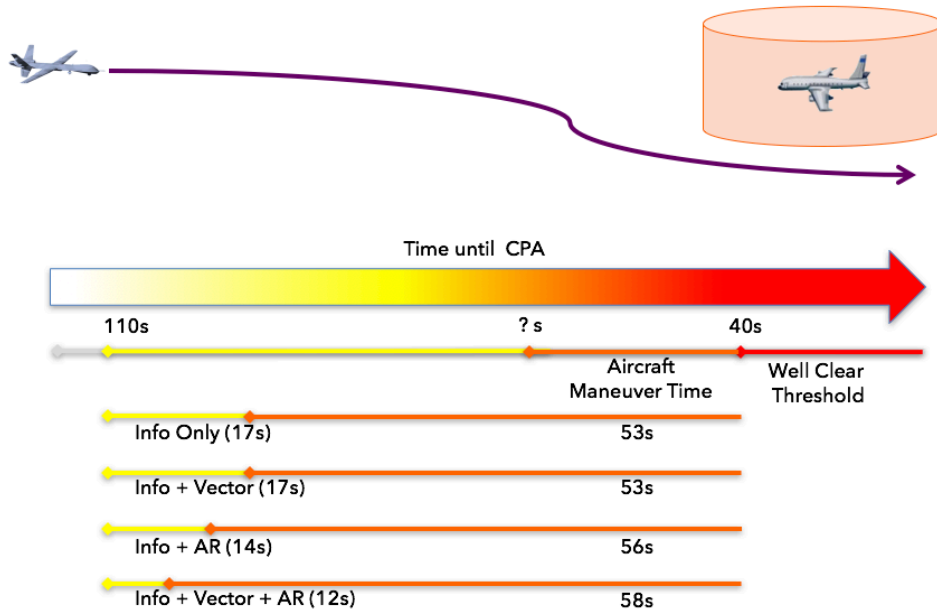


Figure 54. Relationship between total response time for all four displays evaluated in Experiment 2, and the entire DAA timeline.

In Experiment 3 DAA display configuration did have a significant effect on total response times for all encounters, with the Info Only display resulting in the longest total response times. The total response times associated with the Information Only display configuration were 31% longer than those for the No-Fly Bands condition, 23% for the Omni Bands display configuration, and 10% longer than the Vector Planning tools condition (only the difference between the Information Only display and No-Fly Bands display was statistically significant). Total response time in Experiment 3 was not analyzed by alert type, however the descriptive statistics show that pilots were remarkably consistent in responding to DAA warning alerts: means ranged from 12 to 15 sec, and medians ranged from 9 – 12 sec (Table 20). The total response time for DAA warning alerts represents the fastest times that pilots can respond to a DAA alert, since the warning provides an indication of urgency to pilots, and the need to maneuver quickly. In addition, pilots are not required to contact ATC until after they maneuver. Compared to total response times for DAA warning alerts, those for corrective DAA alerts are both longer and much more variable. Although not analyzed for statistical significant, it is unlikely that there are significant differences between display configurations for total response time when the DAA warning alerts are looked by themselves (although the trend for Vector Planner to be slower than No-Fly Bands appears to persist). This suggests that under time pressure, pilots can upload a maneuver within approximately 10 – 15 sec, regardless of display configuration. The results of the maneuvers (i.e., whether a loss of well clear occurred or not) would need to be analyzed by alert type also to determine display effects for this type of time-pressured alert. Given

the overall differences in the proportion of losses of well clear by display type, however, it is likely that the same trend would follow whereby the banding displays result in less violations of the well clear threshold.

Figure 55 shows the relationship between the total response time for each of the display configurations in Experiment 3 and the entire DAA timeline, with an updated well clear time to CPA of 35 s. The total response time means for Experiment 3 are slightly longer than those observed in Experiment 2 (Figure 54). The two better performing banding displays from Experiment 3 had average total response times ranging from 19 – 21 sec, compared to 12 – 14 sec average total response times with the two better performing displays from Experiment 2. This is an expected result, however, since the better performance seen with the those displays from Experiment 2 was likely the result of the directive guidance coupled with the command and control interface. Assuming neither of these features would constitute a minimum requirement for DAA systems, it is unlikely that the same low response times could be achieved with suggestive maneuver guidance displays. Further, although the total response times observed in Experiment 3 were slightly higher than those in Experiment 2, the proportion of losses of well clear observed in Experiment 3 (discussed in more detail in the next section) for the two banding displays were lower to those observed for the Info + Vector + AR display, which had the lowest observed rates for all four displays in Experiment.

Overall, despite the removal of both the directive guidance and the command and control interface coupling, the measured response averages observed in Experiment 3 were not substantially different than those observed in Experiment 2, providing some

evidence as to the reliability and repeatability of the results across the two experiments. Grand mean response time differences ranged between 2 s (initial response time) and 8 s (total response time), with Experiment 2 consistently showing the faster times. In addition, the results suggest that similar performance can be achieved with suggestive guidance displays as with directive guidance displays. The total response time means for all three experiments is shown in Figure 56.




Figure 55. Relationship between total response time for all four displays evaluated in Experiment 3, and the entire DAA timeline.




Figure 56. Total response time means by display configuration for Experiments 1, 2 and 3.

Measured Response Summary

Overall, the measured response metrics are useful for characterizing and evaluating the performance of pilots operating with various DAA HMI configurations. All of the metrics help to understand the general stages and order of tasks that pilots perform in executing the well clear functions. The measured response data also provides some insight into the pilot-DAA timeline with respect to ATC interaction. Despite operational rules requiring pilots to obtain a clearance from ATC prior to maneuvering, overall it appears that pilots only do that half of the time. However, based on the results of Experiment 3, this percentage can improve with an HMI design that provides clear

indication to the pilots when there is or is not sufficient time to contact ATC prior to maneuvering. In addition, Experiment 3 provided a rough estimate of how long ATC interaction takes by comparing the difference in total response times for corrective DAA alerts and DAA warning alerts, which differ in whether pilots are expected to contact ATC prior to, or after, maneuvering. When pilots are not required to coordinate with ATC prior to maneuvering, total response time is roughly 10 sec faster than when they are required to. Thus, ATC interaction adds roughly 10 sec to the overall DAA timeline.

In terms of utility in evaluating effects of various HMI configurations, the initial and total edit times as well as the total response times appear to be the most sensitive to detecting changes in pilot performance. Initial and total edit times, however, are limited in that they only capture one portion of the pilot-DAA time (i.e., the time spent inputting maneuvers into the the GCS control and navigation interface). When the DAA HMI is decoupled from this interface, these times mostly become a reflection of how long it takes pilot to interact with a specific instantiation of a GCS. Although in some cases, total edit time can capture when certain display configurations cause pilots to make multiple uploads before they have fully resolved the conflict.

Total response time appears to be the most useful metric for evaluating the effect of different DAA display configuration performance. This is likely because it captures the pilot's execution of all three well clear functions, and therefore provides more of a global performance. Relying solely on metrics that capture only one portion of pilot performance, such as initial response time or edit times, introduces the risk that tradeoffs that pilots make in performing the well clear functions could be missed. For example, one

HMI configuration may support the pilot in achieving slightly shorter initial response times but much longer total edit times compared to another HMI configuration. If these metrics are not evaluated together in total response time, a false conclusion about the relative performance of the displays could be made. An important point to note here, however, is that faster response times overall do not automatically equate to better system performance. Within the context of evaluating the performance of a DAA system, the proportion and severity of losses of well clear are equally important. Faster response times are not indicative of better performance if they are correlated with higher rates or severity of losses of well clear. The loss of well clear results will be discussed in the next section.

One notable feature of the measured response is the substantial reduction in response times from Experiment 1 to Experiment 2 across all measures except clearance approval time (which showed little variability across all experiments). One hypothesis for the large differences in measured response results between these two experiments is that the display configurations in Experiment 1 posed a training challenge. Both experiments conducted training and data collection on the same day for each pilots. In Experiment 1, the four evaluated displays had large differences that required substantial training effort. Between two different display platforms [the standalone Cockpit Situation Display (CSD) and the integrated display in VSCS] and two information levels (of which the advanced information conditions were instantiated differently), it is possible that pilots did not achieve proficiency with the displays prior to data collection. In contrast, Experiment 2 had four display configurations were largely the same, starting with the baseline

information and adding the vector planner and/or maneuver guidance; from the Info + Vector + AR display configuration all of the three other displays could be derived. Thus, the substantially faster measured response times may simply indicate better display proficiency.

A second hypothesis is that the DAA task was simply more clearly trained in Experiment 2 compared to Experiment 1. Experiment 1 was the first attempt by the researchers to define an experimental test around the DAA task of remaining well clear. At the time of this first experiment, even the operational application of well clear to the DAA task was not clear, and many questions remained an open question even in the SC-228 community, such as: how would pilots be trained on well clear, what would the operational rules around maintain well clear be, when is it ok for pilots to maneuver without a clearance from ATC. Given that the pilots that were recruited for the experiments lacked any experience with maintaining well clear with a UAS, or even avoiding traffic with during UAS operations (UAS are currently segregated in airspace from other aircraft in both military and civil operations), they relied solely on the researcher training on how to execute the DAA task. The long initial response times in Experiment 1 (19 sec, Table 23) could point to pilots being unsure about how to respond to a DAA alert, rather than a delay in alert detection. In Experiment 2 initial response times averaged 8 sec, showing that pilots were more ready to respond immediately to a DAA alert, potentially due to having received clearer training and instructions about how to respond. Overall, it is likely that a combination of both hypotheses explain the significant difference in measured response results between the two experiments.

As noted, there was a small increase in response times between Experiment 2 and Experiment 3. A close inspection of the results in Table 23 shows that the differences can be attributed to the two displays in Experiment 2 that contained the auto-resolution tool and their associated edit times. The auto-resolution tool substantially reduced the time required for pilots to determine and execute a resolution maneuver, which reduced all of the response times that capture the well clear functions (initial response time, initial edit time, total edit time and total response time). When these two displays are removed from the analysis, the measured response results from Experiments 2 and 3 are fairly comparable.

Loss of Well Clear

Since the primary purpose of the DAA system is to provide a UAS pilot located at a ground based control station the means of maintaining well a direct measure of how well a system does that is needed. The metrics chosen for the three reported experiments was the proportion of losses of well clear and the severity of losses of well clear. The first is a measure of the actual losses of well clear that occurred out of all of the encounters that were predicted to result in a loss of well clear. The inverse of this measure as the proportion of time that pilots were able to successfully avoid a predicted loss of well clear. The second is a measure of “how bad” the losses of well clear that did occur were, using a separation index which indicates the proportion of the spatial well clear threshold that was penetrated. The results for each of these metrics across all three experiments is discussed below. The proportion and severity of losses of well clear for each experiment, by display configuration, is shown in Table 24.

Table 24. Loss of well clear metrics for Experiment 1 and Experiment 2.

		Proportion Loss of Well Clear	Separation Index
Exp. 1	<i>Basic Standalone</i>	0.55	1.11
	<i>Basic Integrated</i>	0.49	1.08
	<i>Advanced Standalone</i>	0.37	1.42
	<i>Advanced Integrated</i>	0.28	1.19
	Grand Mean	0.44	1.23
Exp. 2	<i>Info Only</i>	0.28	1.13
	<i>Info + Vector</i>	0.16	1.24
	<i>Info + AR</i>	0.08	1.43
	<i>Info + Vector + AR</i>	0.05	2.06
	Grand Mean	0.13	1.30
Exp. 3	<i>Info Only</i>	0.09	1.09
	<i>No-Fly Bands</i>	0.03	0.98
	<i>Omni Bands</i>	0.00	N/A
	<i>Vector Planner</i>	0.04	1.21
	Grand Mean	0.04	1.08

Proportion

The proportion of encounters that are predicted to lose well clear and do result in a loss of well clear provide the most direct measure of how well a given DAA system is at performing its primary function of maintaining well clear of other aircraft. It is important to note that the proportions observed in a single human-in-the-loop experiment do not provide a very good indication of the expected rate of losses of well clear that will be observed in an operational setting. However, the proportion of actual to predicted losses of well clear across different display configurations provide a good measure of the relative efficacy of a particular human machine interface implementation. The proportion

of losses of well clear across all display configurations in Experiment 1 through 3 are shown in Figure 57.

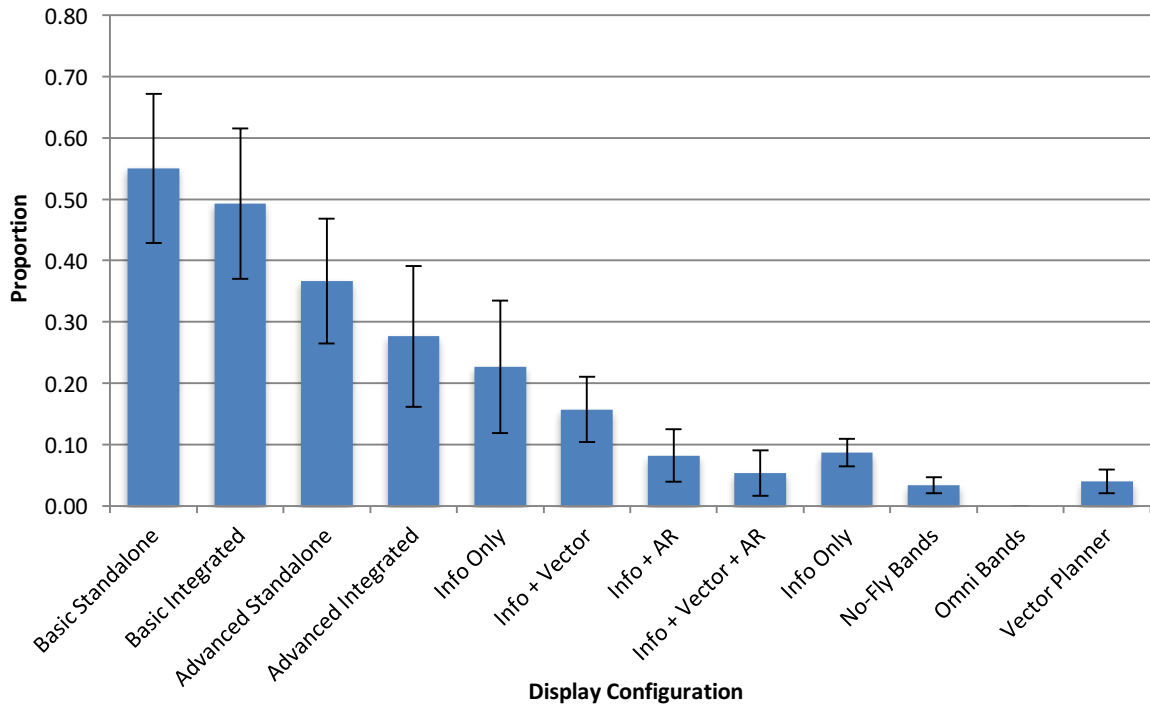


Figure 57. Proportion of actual to predicted losses of well clear by display configuration for Experiments 1, 2 and 3.

In Experiment 1, the proportion of losses of well clear results show a trend toward better performance for the advanced information displays compared to the basic information displays; roughly 36% of encounters in the advanced information display resulted in a loss of well clear compared to half of all encounters in the basic displays. This trend did not occur within the display location manipulation. The percentage of losses of well clear for standalone and integrated appeared to be nearly the same, 46%

and 42%, respectively. The Advanced Integrated had the lowest percentage of losses of well clear overall (28%) and the Basic Standalone had the highest (55%).

There was also an observed trend in proportions of losses of well clear in Experiment 2. Here, proportions trended in favor of the display configurations that contained the auto-resolution tool. Info + Vector + AR had the lowest observed proportion of losses of well clear – 75% lower than the Info Only configuration (0.053 versus 0.227, respectively). The proportion of losses of well clear was also much shorter for Info + AR (0.082) – 64% lower than the Info Only configuration. Info + Vector performed in between the Info Only display and the two auto-resolution displays.

Experiment 3 was the only experiment that revealed significant differences in the proportions of losses of well clear by display configuration. On average, 9% of encounters predicted to lose well clear actually did in the Info Only display which was significantly higher than the percentage found for the Omni Bands display, which had no losses of well clear across all trials (i.e., 0%). While there were no statistically significant differences between the three suggestive guidance displays, the No-Fly and Vector Planner displays performed worse than Omni Bands with 3-4% of all encounters resulting in a loss of well clear. Both displays had losses of well clears due to ineffective maneuvers, too slow response, and too early return to course.

The poorer performance of the Vector Planner configuration compared to Omni Bands can most likely be attributed to the HMI presentation of the suggestive guidance, since the underlying algorithm is the same. The number of each type of loss of well clear provides some insights as to the cause of this performance difference. Out of the six

losses of well clear that occurred in the Vector Planner condition, half were due to an ineffective maneuver. Omni Bands, on the other hand, had no losses of well clear. These results show that when pilots are forced to interact with the display in an attempt to uncover the best maneuver they did not always choose a suitable one, probably because all potential maneuvers were not available at once. Interestingly, the Vector Planner also had two losses of well clear because pilots returned to course too soon, compared to none in the Omni Bands condition doing this. The Vector Planner can be used to test whether the vector back to route is clear of conflict, however it appears that pilots either forgot or chose not to engage the tool. Conversely, the information was readily available for pilots to use when returning to route with the Omni Bands.

The higher rate of losses of well clear in the No-Fly Bands display compared to the Omni Bands display is more perplexing since the approach to the display of suggestive guidance is nearly identical, with the major difference being the use of a single, or multiple colored, bands. A surprising finding from the loss of well clear results was the types of losses of well clears that occurred in the No-Fly bands: two ineffective maneuvers, one too slow response, and two too early returns. The wide array of reasons for losses of well clear do not point to any one explanation. The loss of well clear results indicate that pilots did not make as effective maneuvers based on the available guidance – although the rates of losses of well clear was not significantly different than Omni Bands. It is not clear why losses of well clear occurred to pilots maneuvering where there would theoretically be amber bands (i.e., ineffective maneuvers and too early returns) or why these types of losses of well clear occurred for the No-Fly Bands display but not the

Omni Bands display. One hypothesis is that the multi-color banding concept provided more salient information both for indicating a maneuver that would lead to an imminent loss of well clear (red) versus a less imminent loss of well clear (yellow) and for maneuvers that would result in no conflict (green), compared to amber bands for both a predicted and imminent loss of well clear and no bands for no predicted conflict.

Severity

Loss of well clear severity is a measure of the spatial separation with another aircraft when a loss of well clear does occur. In other words, it is a measure of “how bad” the loss of well clear was. Because of the time component of the well clear definition, in encounters with fast closure rates, losses of well clear can happen even when the aircraft have a large spatial separation. This differs from the way that separation is maintained by ATC in controlled airspace, which is a spatial threshold only. A loss of well clear event that is outside of the spatial well clear thresholds is relatively safer than a loss of well clear event that deeply penetrates the spatial threshold. For comparison, an NMAC is typically measured at 500 ft horizontal and 100 ft vertical, which in Experiments 1 and 2 would be a separation index of 0.25 (a 100 ft vertical separation is 75% penetration of the 400 ft well clear vertical threshold). In experiment 3, an NMAC would be a separation index of 0.22, or 78% penetration of the 450 ft well clear vertical threshold. Because of the low numbers of actual losses of well clear, statistical analyses cannot be done on the separation index results. Due to the wide variability in separation indices across all of the display configurations, distributions are useful for interpreting the separation index

means. The distribution of separation indices for all display configurations is provided in (Figure 58).




Figure 58. Distribution of separation index for all display configurations in Experiments 1, 2, and 3. The lower and upper boxes represent the bottom 25th and 75th percentiles, respectively, the median is indicated by the line intersection the boxes, and the minimum and maximum values are shown by the lower and upper whiskers.

In Experiment 1 the severity of losses of well clear provide mixed results. Overall, the advanced displays showed roughly 30% greater separation compared to the basic displays on average (1.4 versus 1.1, respectively). However, both levels of the information display had a mean separation index greater than 1, which indicates that on average, when a loss of well clear did happen, only the time threshold was violated. In addition, when looking at the distribution of the severity indices, the advanced display configurations actually resulted in the lowest separation index overall, despite the mean and median being higher compared to basic display. The standalone and integrated displays showed roughly the same severity of losses of well clear (1.2 and 1.1 respectively). Overall, the Advanced Standalone saw the largest separation index (1.4) and the Basic Integrated contained the smallest (1.1). When looking at their distributions, the integrated displays had a much wider variability of separation indices, including both the lowest and the highest indices over all. Thus the severity index results do not provide any conclusive trend information favoring any of the display configurations.

The separation index results from Experiment 2 indicate that the Info + Vector + AR had the greatest separation when losses of well clear did occur. The average separation for this display configuration was twice the spatial separation of the well clear spatial threshold, compared to just over one for Info and Info + Vector. Info + AR performed in the middle, with nearly 1.5 times the well clear spatial separation when losses did occur.

There was very little variation in the separation index for the three displays that saw losses of well clear in Experiment 3. The actual to threshold spatial separation at

CPA ranged from a low of 1.0 for No-Fly Bands to 1.2 for Vector Planner. The lack of variation here is likely the result of a very low number of losses of well clear overall. However, it is important to note that the Info Only condition had a loss of well clear with a separation index of 0.24, just outside of an NMAC.

Loss of Well Clear Summary

The loss of well clear metrics provide a means for evaluating the overall performance of a DAA system maintaining well clear of other aircraft. The proportion of predicted to actual losses of well clear is the most direct measure of how well the pilot and DAA system are executing the well clear function. Again, it is important to keep in mind that the observed proportion in a single experiment is not a direct correlation to the rate of losses of well clear that is likely to be observed in real operations. This is because the proportion observed in a human-in-the-loop experiment is dependent on the types and frequency of encounters that are generated. It would be nearly impossible to replicate the expected types of encounters that would be observed in the NAS over a long period of time. Despite that, this metric showed great utility in comparing the performance of pilots executing the well clear functions with difference DAA HMI configurations. In most cases, the proportion of losses of well clear provided converging support to the measured response data. In the first experiment, the advanced display configuration had significantly faster edit and total response times and the proportion of losses of well clear trended in the same direction with lower losses of well clear for the advanced displays. The second experiment showed significantly faster edit times for the two displays that contained the auto-resolutions tools, and again, the proportion of losses of well clear

trended in the same direction. Finally, the third experiment showed significantly faster initial and total response times for the banding displays, and there was a significant difference in the proportion of losses of well clear favoring Omni Bands.

This trend whereby faster response times are correlated with lower losses of well clear makes logical sense. Because of the temporal nature of loss of well clear events, delays in responding to alerts that indicate a predicted loss of well clear are more likely to result in an actual loss of well clear since the delays in time correspond to moving closer to the predicted threat. The later a resolution maneuver is uploaded, the closer you are to the threat aircraft when you begin the maneuver and the less time you have to avoid the well clear threshold. However, the loss of well clear results from Experiment 3 show that late maneuvers are not the only causes of losses of well clear (Table 21). Pilots also make ineffective maneuvers away from other aircraft in addition to turning back to course too soon. In fact, only 13% of all of the losses of well clear in Experiment 3 were due to pilots being too slow whereas 57% were due to ineffective maneuvers and 30% were due to too early return. Thus, the main cause of losses of well clear appears to be more attributable to failures in determining resolution (as well as return to course) maneuvers rather than delays in execution. This means that the proportion of losses of well clear results are not simply a reflection or repetition of the measured response results, but instead appear to measure how well pilots are determining resolution maneuvers. Together, better performance on measured response along with better performance on the proportion of losses of well clear are needed to provide converging support for the efficacy of a DAA system.

One similarity in the proportion of losses of well clear results compared to the measured response results, is the general improvement between Experiment 1 and Experiment 2. The percentage of encounters that resulted in a loss of well clear in Experiment 1 ranged from 28% to 55% with an overall mean of 44%. In Experiment 2, the range of percentages of losses of well clear was 5% - 28% with an overall mean of 13%. That's a roughly 70% reduction in the proportion of losses of well clear overall. There was also a slight decrease in the proportion of losses of well clear between Experiments 2 and 3, with the latter having a range of 0% - 9% losses of well clear. Even the Info Only displays saw a reduction in Experiment 3 compared to Experiment 2 (9% versus 28%, respectively) despite there being less information in the former configuration compared to the latter. This difference may be largely attributable to the new alert structure used in Experiment 3, which provided unambiguous information to the pilot about when they no longer had time to coordinate with ATC in order to avoid a loss of well clear. Although overall response times were longer in Experiment 3 compared to Experiment 2, when response times are broken down by alert type the average response time for DAA warning alerts drops by approximately 10 seconds (Figure 50). These 10 seconds may be the critical difference between maintaining or violating the well clear threshold. Given the consistently fast total response times across all displays for the DAA warning alert, it is reasonable to assume that the inclusion of this alert contributed to an overall reduction in the number of losses of well clear. One way to verify this hypothesis would be to try to go back to the Experiment 2 data and categorize the losses of well clear

the same way it was done in Experiment 3 to see if there is a reduction in the percentage of losses of well clear that were the result of a too late maneuver.

The severity metrics did not provide any conclusive results favoring one HMI configuration over another, with perhaps two exceptions. First, Table 24 and Figure 58 show that the Info + Vector + AR display in Experiment 2 had a much higher mean and median separation index compared to all of the other displays that were evaluated. It also had the least amount of variance in responses, which is likely attributable to this display also having the shortest and least variable response times of all displays (Figure 56). The second clear performance indicator can be seen in the minimum separation index values in Figure 58. The Info Only display in Experiment 3 had an encounter with a separation index of 0.24, which is nearly an NMAC. In general, however, there does not appear to be much correlation between the measured response and proportion of losses of well clear results with the severity results, and there was no real variation across the three experiments. The utility of the severity metric may be limited, except in extreme cases, given the low number of losses of well clear that are expected overall.

Research Summary

The three experiments documented in the previous three chapters sought to determine the minimum HMI requirements for a DAA system. All three experiments utilized an experimental design that compared the well clear performance effects of a “minimum” information-only display (i.e., informative, does not contain maneuver guidance or decision-aiding) to suggestive and directive maneuver guidance displays.

Table 25 provides a quick look “cheat sheet” on the significant differences and trends from the three experiments.

Table 25. Quick look overview of the significant differences and trends from the three reported experiments. The green checkmarks indicate displays that were found to have significantly better performance compared to another display(s), which are marked with red exes. The black checkmarks and exes denote the better and worse performing displays given strong trends in the results.

		Initial Response Time	Initial Edit Time	Total Edit Time	Total Response Time	Proportion of Losses of Well Clear
Exp. 1	<i>Basic Standalone</i>			✗	✗	
	<i>Basic Integrated</i>		✗	✗	✗	
	<i>Advanced Standalone</i>			✓	✓	✓
	<i>Advanced Integrated</i>		✓	✓	✓	✓
Exp. 2	<i>Info Only</i>					
	<i>Info + Vector</i>		✗	✗		✗
	<i>Info + AR</i>		✓			✓
	<i>Info + Vector + AR</i>		✓	✓		✓
Exp. 3	<i>Info Only</i>			✗	✗	✗
	<i>No-Fly Bands</i>	✓			✓	
	<i>Omni Bands</i>	✓				✓
	<i>Vector Planner</i>	✗		✓		

Experiment 1 examined the performance effects of a “minimum” information display compared to an advanced information display with both suggestive and direction maneuver guidance tools. The results showed a clear performance benefit for the

advanced information displays which contained the suggestive and directive maneuver guidance tools in addition to additional (i.e., “advanced”) information element. However, due to the experimental design, it was impossible to determine whether it was the additional information elements or one of the two maneuver guidance tools (or both) that was the source of the performance improvement. Experiment 2 sought to answer this question by separating the “advanced” information only set from the two guidance tools: the vector planner tool (suggestive guidance) and the auto-resolutions tool (directive guidance). Experiment 2 also re-tested the display configuration with both the suggestive and directive guidance. The results indicated that the two displays with the auto-resolutions tool performed significantly better than the two without.

The disproportionate impact of the auto-resolutions was likely the result of two main factors. First, the direct maneuver guidance likely reduced the amount of time it took pilots to determine their avoidance maneuver. In the Info Only and Info + Vector configurations, pilots were responsible for determining their own maneuver. In Info Only this determination was made without any automation-assistance, while in Info + Vector, pilots had the option to use the vector planner tools to inform their decision. Second, the coupling of the auto-resolutions with the GCS’s control and navigation interface meant pilots did not need to make edits to the interface if they were satisfied with what was proposed and pre-loaded by the auto-resolutions algorithm; pilots simply had to upload the recommendation to the aircraft. Analysis of pilots’ compliance rate with the auto-resolutions revealed that pilots followed the algorithm’s suggestion about 70% of the

time it was available, meaning that, more often than not, pilots had no need to interact with the vehicle control interfaces in the Info + AR and Info + Vector + AR conditions.

When the vector planner tools were presented on their own, by contrast, they tended to increase pilot response times, even as compared to the Info Only configuration. This suggests that any benefit provided by maneuver guidance tools in determining an appropriate maneuver is offset by the time it takes to use them. This effect, while not surprising, is counterproductive: by slowing down pilot responses you increase the likelihood that they are not able to upload a maneuver to the aircraft before losing well clear.

Unfortunately, based on discussions within the SC-228 advisory committee, a directive algorithm was unlikely to be accepted as a minimum requirement for the DAA MOPS. In addition, requiring manufacturers to integrate the DAA HMI with the GCS's control and navigation interface was likely to not be accepted as a minimum requirement either. Therefore, Experiment 3 was designed to overcome what was seen as three limitations of Experiment 2: 1) the use of directive maneuver guidance, 2) the integration of the DAA HMI with the GCS's control and navigation interface, and 3) requiring pilots to interact with the suggestive guidance tools. The first issue was addressed by only comparing difference suggestive guidance displays to an information only baseline. The second was addressed by decoupling all of the DAA HMIs from the control and navigation interface. And finally, the third was addressed by introducing a new suggestive guidance display concept, banding. The experimental design for this last experiment compared an information only display to two different banding displays as

well as an updated vector planner display which had been decoupled from the GCS control and navigation interface.

The results of experiment three revealed that the two banding displays resulted in significantly better performance compared to the Info Only display. Both banding displays resulted in faster response times than the Info Only and Vector Planner displays. More importantly, the Omni Bands display had no losses of well clear, a significantly lower proportion of losses of well clear compared to the Info Only display. While the Vector Planner display did not differ statistically from the banding displays, it trended toward slower response times and higher losses of well clear, a trend that could become significant in real operations. Further, a design approach that requires pilots to have to engage the suggestive guidance tool and that essentially hides the guidance information is more likely to be brittle and result in poorer performance, especially under stricter time constraints.

The likeliest culprit of the poor performance seen in the Info Only displays condition is the well clear definition that includes both a spatial and a temporal threshold. While maintenance of a spatial-only well clear threshold, with fixed dimensions like a hockey puck, could likely be supported with an informative or information-only display, the addition of the temporal threshold adds a dynamism to the well clear threshold that is difficult for pilots to judge or visualized, and is also difficult to depict on a two-dimensional display, since the well clear volume changes even during a single encounter as a function of speed. Although Experiment 3 implemented a bare bones minimum information display, given the challenge presented by the temporally based well clear

definition, it is unlikely that performance would be improved by adding additional informative display elements. However, the temporal aspect of the well clear definition is deemed important in order to account for speed differences in perceptions of well clear of manned pilots; a well clear definition for UAS that causes manned pilots to perceive a loss of well clear will not be interoperable with the NAS.

Another major finding of Experiment 3 was the performance difference associated with different alert levels. The observed difference in the total response times of pilots responding to the DAA Warning alert compared to the Corrective DAA alert, especially in terms of shorter and less variable total response times, provides a strong basis for establishing HMI requirements that are closely tied to operational roles. The reduction in the proportion of losses of well clear from Experiment 2 to Experiment 3, despite a slight increase in total response times, also show strong support for the efficacy of the new alert structure tested in the latter experiment. In conclusion, the three reported experiments generated sufficient empirical results to support and inform substantial portions of the DAA MOPS.

Implications for the MOPS

The results of three experiments provide significant objective input into the development of the alerting, guidance and display MOPS for future DAA systems. Most of these inputs were finalized with Experiment 3, however it is important to recognize that the HMI configurations in Experiment 3 were the result of the iterative testing and improvements made after Experiments 1 and 2.

At a high level, the results helped to confirm the minimum time that should be allocated to the pilot within the DAA timeline: 15 sec for pilot response time (i.e., to detect, determine and execute a maneuver), and 10 sec for pilots' interaction with ATC to obtain a clearance approval (Figure 59). These numbers, along with aircraft maneuver time of approximately 30 sec⁵, inform both the alerting thresholds and the minimum surveillance ranges necessary to detect aircraft at those thresholds. In order to detect intruders at roughly 90 sec to CPA, which includes ATC interaction time, the minimum required surveillance range is approximately 8 nm (RTCA, 2015d). Surveillance range capabilities greater than 8 nm will provide additional time for pilot and aircraft response times (including interaction with ATC), as well as a buffer for surveillance uncertainty and communication latencies between the GCS and air vehicle. However, this will be optional for manufacturers to implement. The specific inputs to the alerting, maneuver guidance and display elements sections of the DAA MOPS are discussed below.

⁵ At the time the Experiment 3 was executed, minimum aircraft performance, and more specifically, the amount of time to be allotted to aircraft maneuvering in the DAA timeline, had not yet been defined for the DAA MOPS. After Experiment 3, the allocation for aircraft maneuvering was defined as 30 s. Details for this decision can be found in "Minimum Operational Performance Standards (MOPS) for Unmanned Aircraft Systems (UAS) Detect and Avoid (DAA) Systems Appendix D: UAS Maneuver Performance Requirements DRAFT" (RTCA, 2015d).

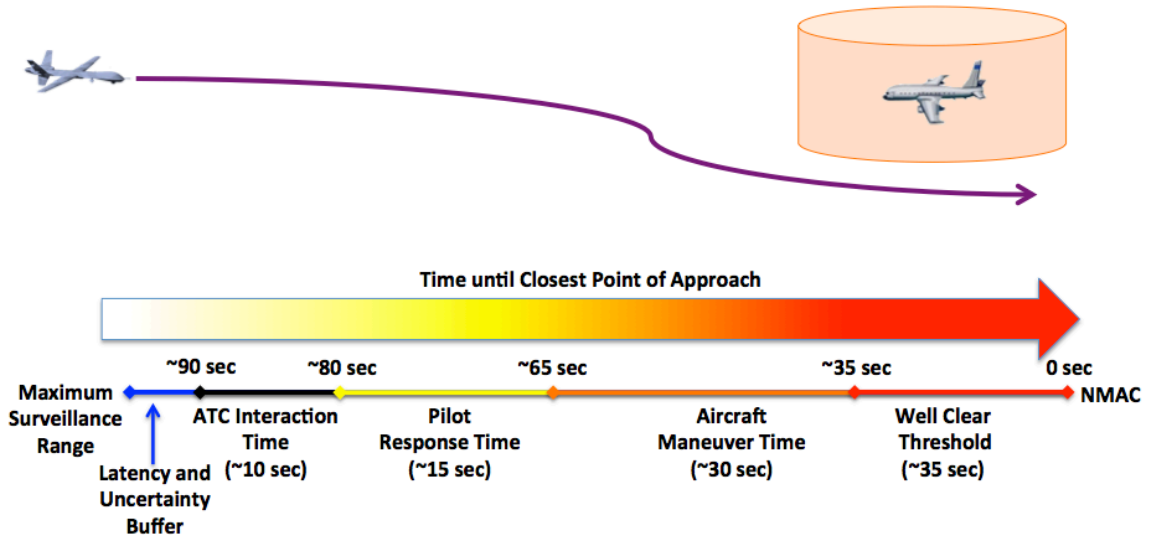


Figure 59. DAA timeline with expected pilot response time, ATC interaction time, and aircraft maneuver time.

Alerting

The results of the Experiment 3 helped to define both the alert levels of the DAA MOPS alert structure as well as the associated thresholds for each alert for the current (at time of writing) draft of the DAA MOPS [see section 2.2.4.2.4 “Alerting” of the “DRAFT Detect and Avoid (DAA) Minimum Operational Performance Standards (MOPS) for Verification and Validation” (RTCA, 2015a) for the full list of alerting requirements]. The three required alert levels are Preventive, Corrective and Warning. (Based on subjective feedback, the Proximal alert level was no included as a required alert level, however a different alert level is currently being proposed to indicate traffic that are not an active Preventive, Corrective or Warning alert, but that are causing yellow or red suggestive maneuver guidance bands.) Each of these alert levels is associated with both a must and must not alert threshold that defines the trade space for manufacturers of

when alerts should and should not be generated, rather than specifying a particular implementation. For each alert, the τ_{mod}^* (modified tau), HMD^* (horizontal miss distance), h^* (vertical distance), and DMOD (distance modifier) parameters define the alert threshold for future predicted ownship and intruder states for the Within Time. The HMD_p and d_{h_p} parameters define the must not threshold for the More Than Time. Thus, for the Preventive and Corrective alerts, their must alert time is within 55 seconds (to loss of well clear) with a modified tau of 35 sec, and a horizontal miss distance and distance modifier of 0.66 nm. They differ only in the vertical distance threshold; the Corrective alert vertical threshold is equivalent to the well clear threshold of 450 ft while the Preventive alert vertical threshold is 700 ft, providing the vertical buffer as specified by the FAA whitepaper on well clear (Walker, 2014; Figure 36).

Table 26. *DAA Draft MOPS alerting thresholds. Reprinted with permission from: RTCA (2015a), DRAFT Minimum Operational Performance Standards (MOPS) for Unmanned Aircraft Systems (UAS) Detect and Avoid (DAA) Systems.*

		Preventive Alert	Corrective Alert	Warning Alert
Alert Level		Caution	Caution	Warning
Must Alert Threshold	Within Time (Seconds)	55	55	25
	τ_{mod}^* (Seconds)	35	35	35
	DMOD and HMD* (NM)	0.66	0.66	0.66
	h* (Feet)	700	450	450
Must Not Alert Threshold	More-Than Time (Seconds)	75	75	35
	HMD _p (NM)	> 2.0	> 1.5	> 1.0
	d _{h,p} (Feet)	> 800	> 450	> 450

There are two important differences to note about the draft DAA MOPS thresholds compared to those utilized in Experiment 3. First, the horizontal miss distance and distance modifier used in Experiment 3 was 0.75 nm or roughly 4500 ft, whereas for the well clear definition and the draft MOPS alert structure in Table 26 they are 0.66nm or roughly 4000 ft. The reason for this 500 ft discrepancy is that Experiment 3 used a slightly larger horizontal threshold to account for the uncertainty associated with surveillance and aircraft performance expected in real world operations. In order to achieve the safety target or a of 0.045 loss of well clear probability, some amount of

buffering will be required in the alerting structure to allow for this uncertainty. However, this number has not yet been specified in the draft DAA MOPS. The second difference between the must alert thresholds in Table 26 and the thresholds used in Experiment 3 is the time threshold. Experiment 3 used a time threshold of 75 sec to loss of well clear compared to a must alert requirement of 55 sec. This is because the must alert threshold specifies the latest time at which an alert can be triggered. The must alert threshold of 55 sec was the minimum time calculated from the DAA timeline whereby aircraft performance is allocated 30 sec and pilot response time is allocated 25 sec (10 sec for ATC coordination and 15 sec for GCS interaction; Figure 59). The pilot response time number with ATC coordination was derived from the average total response time for pilots responding to Corrective alerts in Experiment 3 (Table 20). The Corrective and Preventive alert thresholds for Experiment 3 (75 sec) were the furthest time out it was judged to be operationally acceptable for pilots to request a maneuver to remain well clear based on previous studies (Mueller, Isaacson, & Stevens, unpublished), which established the current must not alert thresholds for the draft DAA MOPS alert structure. However, the Warning alert threshold from Experiment 3 is identical to the must alert threshold specified in the draft DAA MOPS, which allocates roughly 10 sec to pilot response time (based on pilots' total response times to Warning alerts, Table 20) and 10 sec to aircraft maneuvering⁶. The must not alert indicates that a Warning alert could

⁶ Although 10 s for aircraft maneuvering (given a 25 s to loss of well clear threshold for a Warning alert) is likely not sufficient to prevent all imminent losses of well clear, concerns about ATC interoperability and having pilots maneuver up to a minute prior to CPA without prior ATC coordination forced the committee

happen as early as 35 sec to loss of well clear, but not earlier. The must not alert threshold is critical in terms of ATC interoperability given the concerns of having pilots maneuvering without notifying ATC beforehand. The horizontal and vertical must not alert thresholds establish the maximum buffer around the well clear boundary that can be used in alerting and are set to minimize unnecessary alerts and resulting maneuvers. Thus, the must alert thresholds are established according to the minimum performance expectations of the pilot and aircraft in order to achieve the loss of well clear safety target, while the must not alert thresholds are established according to operational considerations. Alert symbology was selected based on the Aircraft Surveillance Applications MOPS (RTCA, 2011) and are specified in section 2.2.5.6.2.3.2.4 of the draft DAA MOPS (RTCA, 2015a). The auditory alerts utilized in Experiment 3 (Table 16), based on subject matter expert input, were accepted as minimum requirements in the draft DAA MOPS (section 2.2.5.9.1.1) with the exception of the aural alert of the Corrective alert. “Traffic, Separate” was changed to “Traffic, Coordinate” based on terminology concerns from the ATC community with using “separate” and “separation”.

Maneuver Guidance

Experiment 3 provided strong objective data to support requiring suggestive guidance in the form of banding as a minimum requirement for DAA systems. Section 2.2.4.3.1 (RTCA, 2015a) specifies the functional requirements for the DAA maneuver

to accept a shorter time threshold than would be necessary to prevent all losses of well clear. By comparison, a TCAS RA, for which pilots are allowed to maneuver with prior ATC coordination, occurs roughly 25 - 35 s to CPA.

guidance, e.g., “DAA guidance **shall** include ranges of horizontal maneuvers.” Section 2.2.5.6.2.3.2.6 (RTCA, 2015a) specifies the HMI requirements for the maneuver guidance, e.g., “An amber/yellow caution band **shall** be displayed with its associated predicted DAA preventive/corrective alert in any direction, and a red band **shall** be displayed with its associated predicted DAA warning alert in any direction.” As this requirement shows, it is necessary for the maneuver guidance to differentiate between trajectories that are predicted to result in a DAA warning alert with red bands and those predicted to cause a corrective DAA alert with yellow bands. This was one of the major differences in the HMI for the No-Fly Bands and Omni Bands configurations, and was chosen given the better performance of Omni Bands over No-Fly Bands in maintaining well clear. While the use of green bands was not established as a minimum requirement, and instead recommended to be saved for future system guidance (i.e., collision avoidance guidance), there was a requirement included for the display to differentiate between trajectories that are and are not being probed. The reason for this requirement is so that the lack of banding information (if not all trajectories are probed or are not returning conflict information) is not confused with trajectories predicted to not cause a conflict, (i.e., if no colored bands are used for showing trajectories predicted to be safe).

Display Elements

The results of Experiment 3 prompted a return to the minimum display requirements outlined in the “Minimum Operational Performance Standards (MOPS) for Aircraft Surveillance Applications (ASA) System” (RTCA, 2011) as a baseline for the DAA MOPS. It was determined that the suggestive maneuver guidance was the primary

source of decision support, and contributed to acceptable performance, for pilots remaining well clear of other aircraft, and that only basic traffic elements were necessary. For traceability purposes, each requirement within the display elements section that came from the ASA MOPS references the line in the original document where the requirement can be found, e.g., “A data tag for airborne traffic **shall** (3032) include: traffic altitude.” In this requirement, which is located on line 3677 of the DAA MOPS (RTCA, 2015a), the line from the original MOPS is also referenced in brackets (i.e., 3032). Where relevant, additions and subtractions were made to accommodate specific DAA characteristics.

Other MOPS Considerations

Experiment 3 also helped to formalize pilot operational responsibilities with respect to interacting with ATC when maneuvering to maintain well clear, which had previously been largely ambiguous. Before Experiment 3, the operational assumption was that pilots were required to obtain a clearance to prior to maneuvering to maintain well clear unless there was a safety of flight concern. The challenge for pilots was in subjectively determining when a safety of flight concern occurred. Not surprisingly, there was a lot of variability in how pilots judged a safety of flight concern – some pilots were observed to maneuver as soon as an alert occurred (at 110 sec to CPA) without contacting ATC whereas other pilots were observed to penetrate the well clear threshold while attempting to contact ATC. With the introduction of the DAA Warning alert, the safety of flight concern was defined objectively for pilots as 25 sec to loss of well clear. The establishment of clear operation responsibilities (whether or not to contact ATC) coupled

with HMI requirements that leverage established human factors regulations to elicit appropriate pilot responses, resulted in more predictable pilot behaviors and overall improved performance.

Chapter 8: Discussion

As old technologies evolve into new ones, there emerges new opportunities to introduce increasingly sophisticated forms of automation. With these new opportunities, however, also comes the question of how work will be coordinated and accomplished between humans and machines. As new generations of machine capabilities replace old ones, roles and responsibilities within the system change, which means new requirements for supporting human-machine cooperation. In safety critical systems, this question is as important to regulators as it is to the designers of the system, as regulators make predictions about the future performance of the new technology as part of their safety assurance processes.

Traditional function allocation methods, such as the use of MABA-MABA lists and levels of automation, have fallen short of their promises to provide easy answers to who or what will accomplish certain tasks or functions when new technology is introduced. This is often because these paradigms typically ignore how human-machine cooperation will be supported. Indeed, traditional function allocation approaches ignore questions about issues that go beyond specific functions, such as: what are the roles and responsibilities of various human and machine agents, how do agents coordinate across their respective roles and across different levels to accomplish their goals, and how are activities synchronized (Murphy & Shields, 2012). These questions concern the

underlying human-automation architecture of the system. Unfortunately, function allocation activities in the design of new systems typically make implicit assumptions about the underlying human-automation architecture, rather than making explicit the coordination and synchronization needed by agents in order to fulfill their relative roles and responsibilities and accomplish the goals of the system. When these needs are ignored by system designers, failures in human-machine cooperation are likely to occur, leading to undesirable system outcomes. Thus, the human-automation architecture and support for human-machine interaction and cooperation, often in the form of HMIs, need to be sufficiently evaluated in the safety assurance process for safety critical systems, where undesirable outcomes could potentially be catastrophic.

A methodology is needed that enables regulators to ensure that the human-machine cooperation needs of new technology are accounted for and supported by designers. Examining the changes in the underlying human-automation architecture of a new technology will elucidate how existing roles in the domain will be impacted by its introduction into the larger system. This is because no new technology can be introduced into a system without changing the roles and interactions between roles that are needed to get work accomplished. Thus, a methodology that makes explicit the coordination and synchronization needs of the roles required to effectively operate a new technology will enable regulators to better anticipate the reverberations and side effects that will occur when that technology becomes operational.

The experiments discussed in the last four chapters detail a process that was implemented to help determine the minimum HMI requirements for future DAA systems

for UAS. Not only did these studies meet the short term UAS community needs of informing the DAA MOPS, but they also provided key insights into how the functions of remaining well clear will likely be accomplished by UAS pilots in the future operational environment. How UAS pilots would meet the regulatory requirement to remain well clear was not well understood by the UAS community or the federal advisory committee tasked with developing MOPS when this process started. While well clear is currently accomplished by pilots of manned aircraft, there is little data or insight (other than anecdotal) on how it is accomplished, let alone how it might be accomplished by ground-based pilots with the assistance of a DAA system. A key revelation of the reported studies, is that how, and “how well”, the critical functions of a system get accomplished is highly dependent on the DAA HMI and underlying human-automation architecture.

For UAS pilots, maintaining well clear has to be accomplished through the DAA and control and navigation systems located both in the GCS and onboard the aircraft. The pilot on the ground can neither directly detect potential threats nor directly input maneuvers to the aircraft. Instead they have to rely on receiving information and sending aircraft commands through the machine components on the ground via a communication link. Because of the reliance on electronic detection of potential threats, the regulatory responsibility of maintaining well clear was changed from a subjective out-the-window assessment to a new separation standard based on a mathematical equation. With this change came a new ability to objectively assess whether well clear has been “lost” (i.e., the well clear separation threshold has been penetrated) and a new safety standard that the pilot and DAA system, together, have to meet.

The reported studies revealed that because of the complex nature of the well clear threshold, which utilizes a time-based, or tau, component in addition to a spatial component, UAS pilots were not very successful at maintaining well clear without some automation assistance. More specifically, pilots struggled with one of the three well clear functions in particular: determining an appropriate resolution maneuver. The results of the studies showed that pilots were slower and more likely to make ineffective maneuvers that lead to losses of well clear when they only had a display of traffic information elements. On the other hand, when pilots had use of an algorithm that presented either a single or a range of maneuvers that were predicted to resolve the potential conflict, they were quicker and significantly less likely to penetrate the well clear threshold. These findings lead directly to the inclusion of the requirement for maneuver guidance in the DAA MOPS (RTCA, 2015a). In addition, the studies pointed strongly to the need of HMI requirements that support pilot coordination with ATC. Although coordination is not necessary to accomplish the well clear functions, it is a critical requirement to ensure interoperability with current ATC roles and responsibilities, and to reduce the disruption to the aviation system when UAS are deployed.

The alerting and maneuver guidance requirements that were instantiated in the draft DAA MOPS point to the underlying human-automation architecture of the DAA system, whereby the DAA system and the pilot share and coordinate the well clear function of determining a resolution maneuver (Figure 60). In this architecture, the automation provides a range of maneuvers that are predicted to be clear of conflict to the pilot, who then decides on a single maneuver and uploads that to the aircraft using the

GCS control and navigation automation. In addition, the DAA system does not simply pass surveillance information of nearby aircraft through to the DAA display located in the ground-based control station, but it also utilizes automation to provide information to the pilot about the predicted threat level of, calculated as the predicted time to loss of well clear with, potential conflicts through alerting. The requirements derived from the three reported studies therefore not only assist the pilot in determining “how” to maneuver to avoid a loss of well clear, but it also assists the pilot in determining “when” to maneuver to avoid a loss of well clear. In this human-automation architecture, the DAA automation also supports and mediates the human coordination between pilots and ATC, which is an operational pilot function, not a well clear function. However, serves to also support the well clear functions by giving pilots information on when they need to prioritize maintaining well clear over ATC coordination. This point highlights how different goals change over the temporal timeline of a DAA event; further out in the timeline, ATC coordination takes priority, but as the event develops and the threat becomes more imminent, maintaining well clear takes over priority. This pacing and synchronization of the pilot’s roles and responsibilities relative to other users of the system (i.e., other pilots/aircraft and ATC) over time is supported by the DAA system.

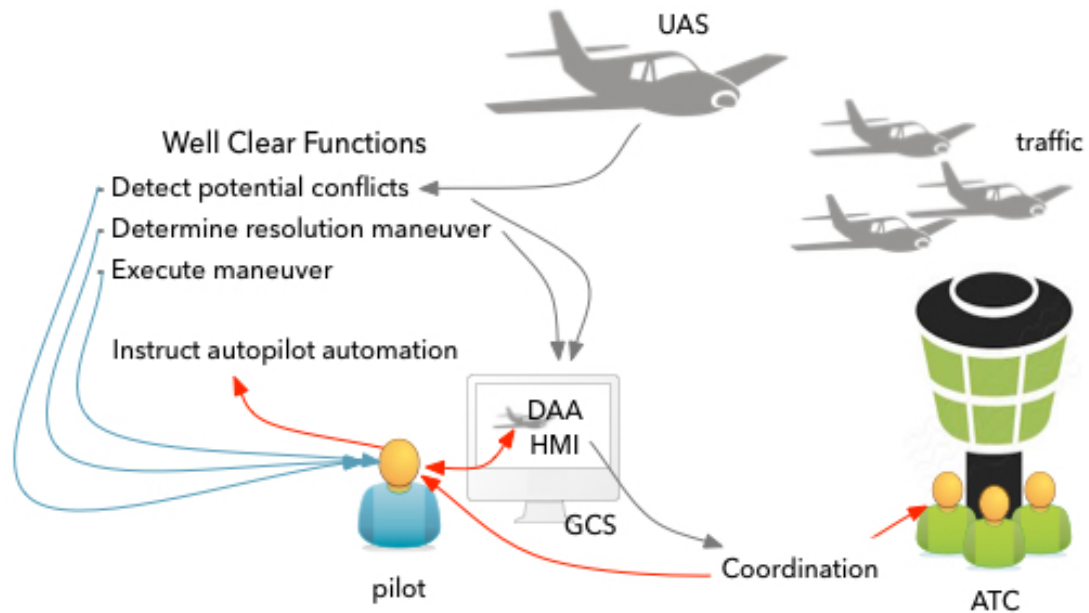


Figure 60. How the well clear functions are accomplished according to the draft DAA MOPS (RTCA, 2015a).

In the first chapter of this dissertation, two key conditions were identified that had to be met in order for a methodology to successfully aid in the transition of new technologies to highly regulated domains. These conditions were derived from CSE principles based on previous research and observations of people at work in complex systems. The first condition was that the methodology needed to be able to uncover the shifts in human and machine roles and responsibilities so that the reverberations resulting from a change in technology within the existing system could be predicted. The second condition was that those predictions would be based on empirical results abstracted from observations within a relevant context. The next section will evaluate the methodology implemented in the reported research program with respect to these two conditions and from an overall CSE perspective.

Reflecting on the Methodology for Determining HMI Requirements for DAA Systems
from a CSE Perspective

To begin the evaluation of the DAA HMI research program, a high level overview of the methodology that was implemented is provided in Figure 61. The research program began with the generation of candidate HMIs for DAA systems. This step was followed by the empirical evaluation of the prototype HMIs and then the analysis of the results of those evaluations. At the end of the third step, results both served as inputs to the refinement of existing HMI prototypes or the development of new ones and became inputs to the DAA requirements (step 4). The process was repeated as lessons learned and new HMI concepts drove subsequent evaluations. Throughout the process, the underlying human-automation architecture was assumed both in the candidate HMIs that are developed in the first step, and in the derived requirements in the last step.

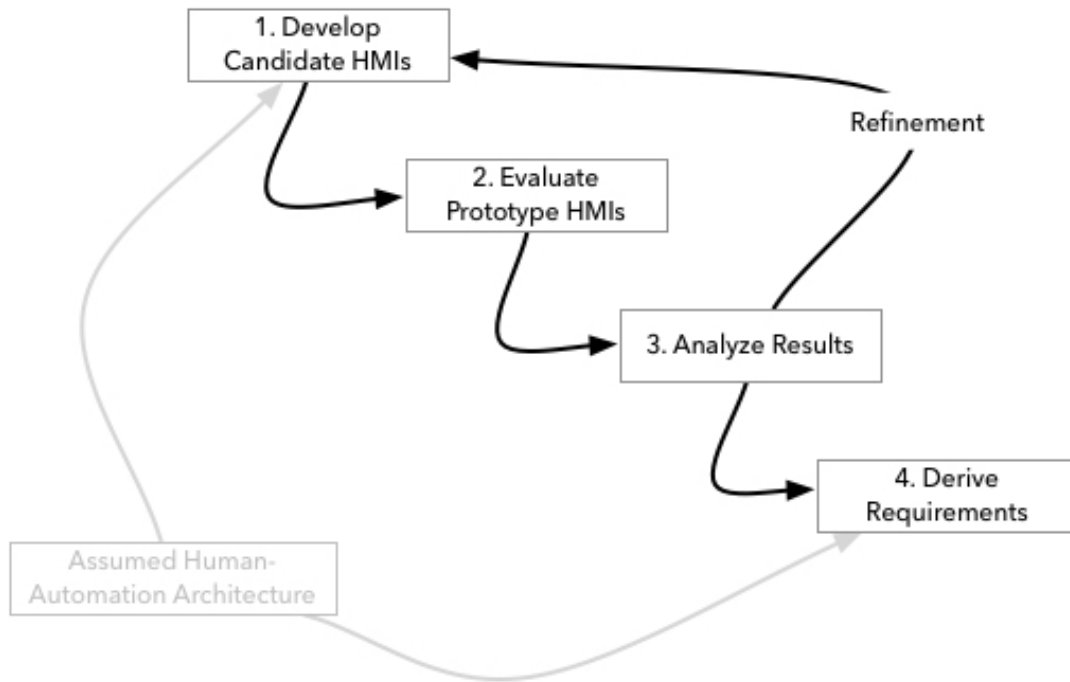


Figure 61. Methodology implemented for the DAA HMI requirements development.

Capturing Shifts in Roles and Responsibilities

A key limitation of the methodology presented in Figure 61 is the absence of a step that captures the shift in the roles and responsibilities from the previous generation of technology to the new one. This can only be done by analyzing the underlying human-automation architecture of the current system and comparing to it candidate prototypes. Figure 62 compares the human-automation architectures for how the well clear functions are accomplished for both manned and unmanned aircraft, assuming the architecture determined by the DAA HMI research program for the MOPS. This comparison is derived from Figure 1 and Figure 60 with the ATC coordination links removed, so it therefore captures only the accomplishment of the three well clear functions.

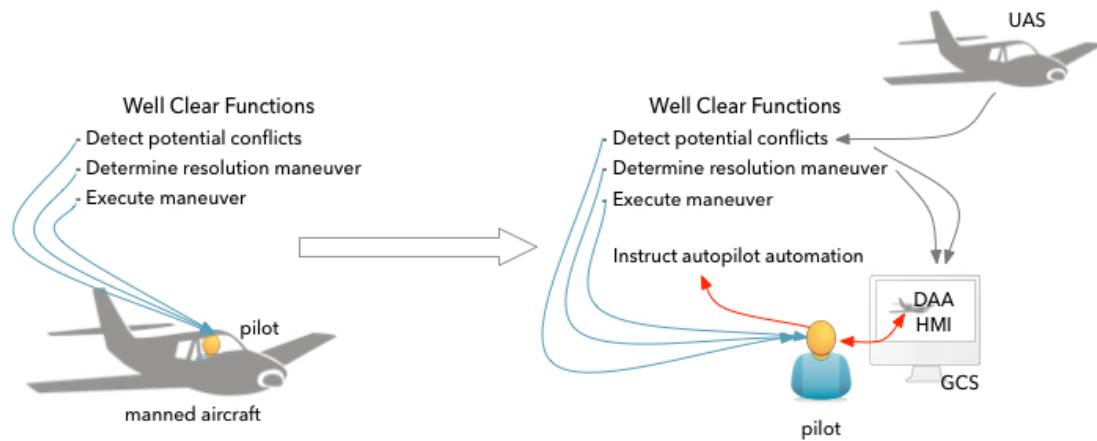


Figure 62. How the well clear functions are accomplished when coordinated between pilot and DAA system for UAS (left) compared to how they are accomplished by manned aircraft (right).

Figure 62 underscores how the introduction of technology substantially changes the roles within a system. While the UAS pilot on the left maintains responsibility for all of the same three well clear functions as the manned pilot on the right (shown with the blue lines), there are new coordination links (red lines) between the pilot and automation capabilities via the GCS. This is because the DAA system takes on a new role with responsibilities for two of the three well clear functions (grey lines). The machine's role is mediated through the HMI which the pilot relies on to coordinate and synchronize his/her activities with the machine's, as well as other agents or roles in the larger aviation system (e.g., ATC and other pilots). In addition, the UAS pilot also remotely controls the unmanned aircraft through the aircraft's autopilot system located within the GCS via a command and control link.

By comparing the old and new human-automation architectures side-by-side, it becomes much easier to identify the new roles and coordination between roles that need to be supported by the new technology's HMI. For UAS DAA systems, the introduction of the machine role into the well clear functions, and especially the new coordination links between the machine and human roles, are a critical area of needed support.

Making Predictions Based on Empirical Observations in a Relevant Context

The DAA HMI research program was successful in generating empirical data within a relevant operational context that allows the responsible regulators to make predictions about the future performance of UAS DAA systems. In particular, the research program provided predictions about the expected performance of UAS pilots remaining well clear of other traffic in the airspace in terms of response times and proportions of encounters that result in a violation of the well clear threshold. These predictions not only gave justification to the several of the minimum requirements instantiated in the DAA MOPS as documented in the previous chapters, but they also provided objective data as input into other analyses (e.g., fast time and Monte Carlo simulations, as well as engineering analyses) that will result in minimum requirements for other components of the DAA system not directly related to the HMI (e.g., surveillance equipment performance).

One reason for the ability of the research program to successfully do this, was the substantial analysis of the future envisioned world that was provided within the Operational Services and Environment Description Document (RTCA, 2015b). More important, however, was the ability to bring significant resources to the research program

in order to build up a simulation environment that supported the development of prototype DAA HMIs and supporting DAA systems (i.e., JADEM and Stratway+) that were capable of realistically emulating surveillance equipment and alerting and guidance algorithms, as well as being able to network the airspace and ATC simulation with the UAS GCS.

Unfortunately, the DAA HMI research program employed only a single instantiation of all of the possible combinations of DAA HMIs, systems (i.e., alerting, guidance, and surveillance fusion and tracking), UAS GCSs, and operational environments. Thus, the predictions made about the performance of the future DAA system based on this research are only good to the extent that the actual operational context matches the simulated environment. Given that the boundaries of the simulated operational context are relatively narrow relative to actual operational environments, the risk is high that a DAA system resulting solely from this research program would have narrow operational boundaries and be brittle when deployed in conditions it was not tested in. For example, a recent study UAS DAA study also conducted in support of the SC-228 DAA MOPS uncovered potential brittleness when the DAA system did not make accurate predictions about the autopilot functionality in the UAS GCS (Comstock, Consiglio, Ghatas, & Vincent, 2016). This unanticipated mismatch in the predicted and actual autopilot performance caused the maneuver guidance algorithms and DAA HMI to provide incorrect guidance to the pilot, e.g., it indicated that particular trajectories were conflict-free when in fact they weren't based on the way the autopilot would attempt to capture the heading when commanded by the pilot. When this incorrect guidance was

followed by the UAS pilot, it resulted in violations of the well clear threshold. Thus it is critical that the DAA system instantiated in the final MOPS be tested in various contexts and environments, by various methods (i.e., human-in-the-loop and fast time simulations, flight test, engineering analyses, etc.) and under multiple possible configurations of GCSs, HMIs, and DAA systems so that the regulators can get a more robust assessment of the performance, potential risks, operational boundaries, and impacts on other roles in the NAS. A critical part of this testing process is the development of scenarios, under all methods, that are designed to test the boundaries of the system. As will be discussed further in a later section, the scenarios utilized in the DAA HMI research program were limited in the extent to which they tested both the bounds of the DAA system performance (i.e., alerting and maneuver guidance algorithms) as well as the human-machine coordination supported by the HMI.

Summary

Overall, the DAA HMI research program partially met the conditions outlined in the first chapter for what would be required of a methodology that could aid in the transition of new technologies to highly regulated and complex domains. While it was able to provide objective results and observations to support predictions about how future DAA systems will perform and impact the future NAS into which it will be deployed, those predictions are based on a narrow characterization of the envisioned world. In addition, the methodology failed to explicitly identify the transition in roles and responsibilities for accomplishing the well clear functions that is expected to occur with the introduction of UAS equipped with DAA systems. The result of these two limitations

is an increased risk that the new human and machine roles, and the coordination between them, will not be properly supported by the DAA HMI and that the testing of the DAA prototypes (including the HMIs and underlying processing systems) did not sufficiently identify where the system boundaries are and when breakdowns might be predicted to occur. One proposed way to overcome these limitations is to modify the methodology presented in Figure 61 by adding a step of generating potential human-automation architectures before candidate HMIs are developed. The contribution of this step to the overall methodology is discussed in the next section.

Improving the Methodology for Developing HMI Requirements for New Technologies

Figure 63 presents a slight modification to the methodology that was utilized in the DAA HMI research program. The generation of potential human-automation architectures is a critical first step as it serves as a precursor to other activities that feed the overall process. Outlining the art of the possibility in human-automation architectures prior to evaluation allows for a clearer definition and design of the prototypes to be tested. In addition, by analyzing the different coordination and synchronization needs for potential human-automation architectures and the change in architecture from the previous generation technology, HMIs can be designed that are specifically designed to tailored to support them. In this way, the human-automation architecture drives the design of the HMI. When the HMI is designed independent of the human-automation architecture, there is a risk that the underlying assumptions about the architecture, and the related human-machine cooperation requirements, will not or cannot be met in the overall technology design. If this happens, a mismatch between the human-automation

architecture needs, in terms of coordination and synchronization, will not be appropriately supported by the HMI, resulting in human-automation cooperation breakdowns.

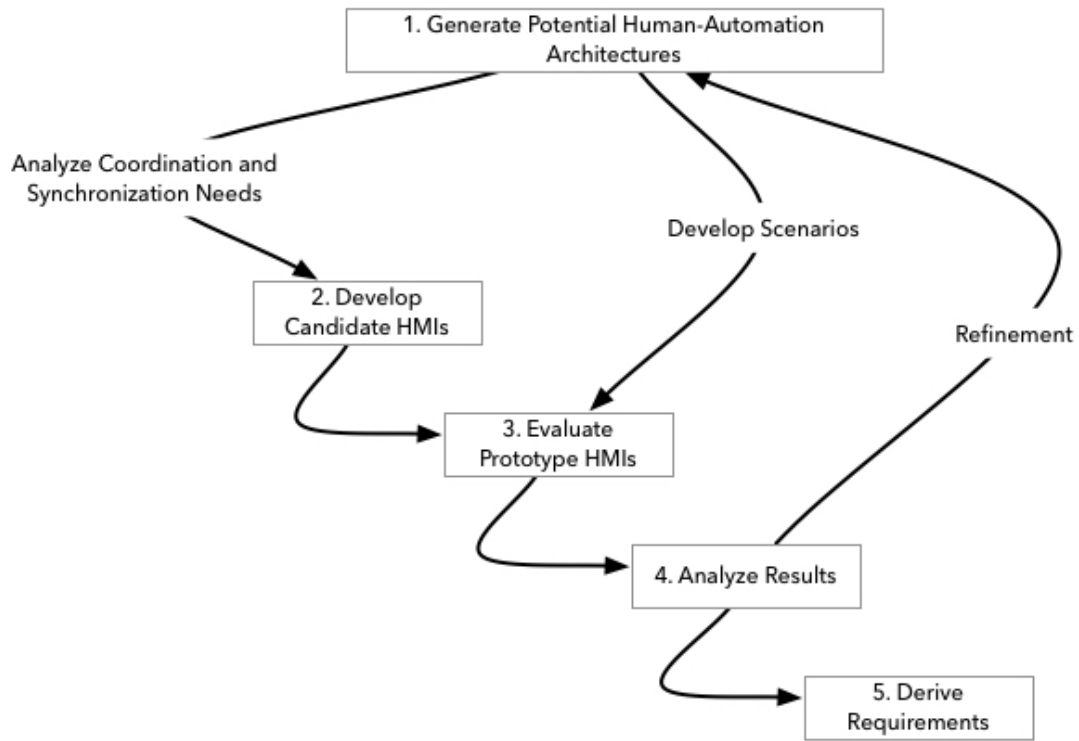


Figure 63. Modified methodology for developing HMI requirements for new technologies.

In the DAA HMI research program design decisions were being made at the HMI level, which were making implicit assumptions about the underlying allocation of functions between the automation and pilot, rather than allowing the design decisions to be driven by the specific human-automation architecture. The main research question

underlying the HMIs that were evaluated across the three experiments essentially boiled down to whether, and what type, of maneuver guidance should constitute the minimum requirement for UAS DAA systems. Instead of this being a question about the appropriate minimum HMI requirement, the fundamental question is really about what the appropriate human-automation architecture for DAA systems is. This is because while technical requirements will typically be framed at the level of the HMI, the HMI is really just a proxy for the human-machine coordination support that is needed by the human-automation architecture. Thus, to properly derive HMI requirements, one should start at the architecture, or possible architectures, of the new technology.

The informative, suggestive, and directive maneuver guidance display configurations evaluated in the reported experiments varied in how the “determine resolution maneuver” well clear function was accomplished by the automation and pilot. Because of the original assumptions by the SC-228 advisory committee that there would be a pilot-in-the-loop with the DAA system that would be responsible for executing any resolution maneuvers, the execute maneuver function was always allocated to the pilot (who essentially directs the autopilot system to execute the maneuver) while the detect potential conflicts function was necessarily shared between the pilot and automation. How the determine resolution maneuver function would be accomplished by the pilots and/or machines, however, was not pre-determined but rather implied based on the HMI requirements. This baseline architecture is shown in Figure 64. The key question the current research program was explicitly trying to answer, was whether an informative display, which does not provide any maneuver guidance, would be a sufficient minimum

HMI requirement for the DAA system. What this question was really asking, however, is how the determine resolution function should be accomplished by humans and machines in order to achieve the desired level of performance of the DAA system.

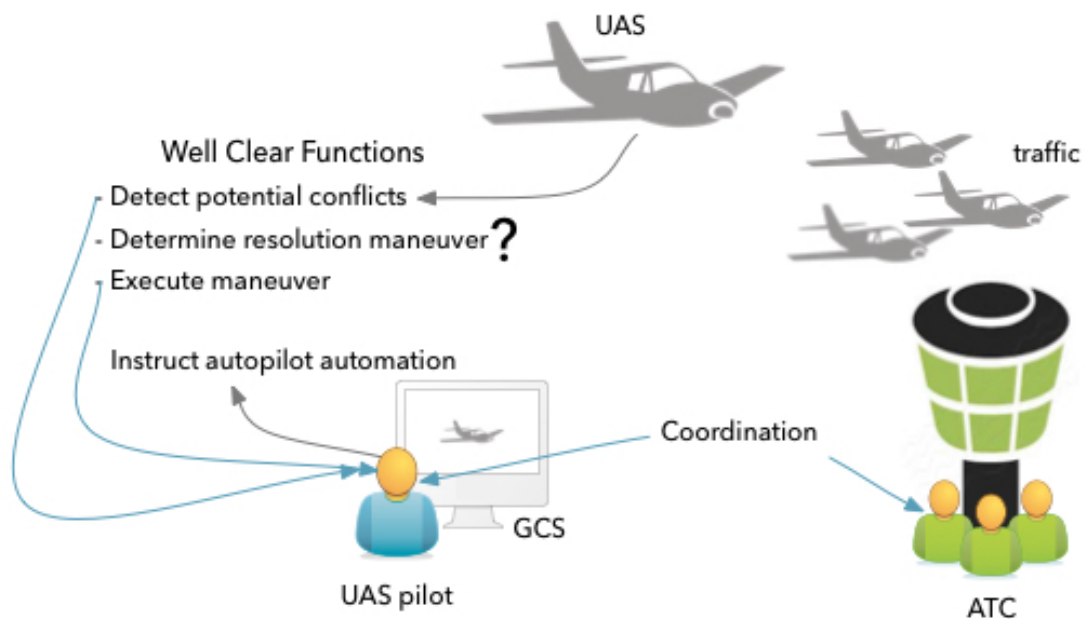


Figure 64. Baseline human-automation architecture assumed by SC-228 DAA MOPS.

An *informative* display solution would provide essential hazard information to the pilot along with other information to develop and execute an avoidance maneuver; no maneuver guidance is provided. This solution implies a human-automation architecture where the determine resolution maneuver function is accomplished solely by the pilot based on the presentation of potential conflicts on the DAA display. Alternatively, a *directive* HMI solution would provide a specific recommended solution to the pilot who would then execute the maneuver manually. In this human-automation architecture, the

function of determining a maneuver is accomplished solely by the machine unless the pilot rejects the machine's solution and is forced to determine an alternate maneuver manually. Finally, a *suggestive* HMI would provide a range of potential resolution maneuvers to the pilot who would then decide which maneuver was most appropriate given the operational context, and execute it. This last HMI solution implies an architecture where the human and machine jointly accomplish the function of determining a safe resolution maneuver. Table 27 shows how the determine resolution maneuver function is accomplished for each HMI configuration.

Table 27. How the determine resolution maneuver function is accomplished for each of the three evaluated HMI configurations.

HMI Configuration	Machine	Human
Informative	Show potential conflicts	Determine maneuver
Suggestive	Provide set of potential resolution maneuvers	Choose from set of maneuvers
Directive	Determine a single resolution maneuver	Accept single maneuver or determine alternate maneuver

By not generating the potential human-automation architectures prior to developing the candidate DAA HMIs as depicted in Figure 63, three key opportunities were missed: 1) generation of human-automation architectures based on alternative human-machine cooperation approaches; 2) development of candidate HMIs based on the human-machine cooperation needs of the underlying architecture; and 3) development of scenarios designed to stress the boundaries of the prototype system, especially those that

depend on human-machine coordination. How these three missed opportunities can contribute to a more comprehensive methodology for developing HMI requirements is discussed in turn.

Alternative Approaches to Human-Automation Architectures

The human-automation architectures exhibited by the three HMI configurations described in Table 27 follow the same general approach to human-machine cooperation: the machine does something, then the human does something. This serial approach to human-machine cooperation is typical across many systems and in many domains. The typical manifestation of this approach is that the machines does some of the work required by the system and then the human has to evaluate or critique the work done by the machine, and then act based on that evaluation of the machine's work and an assessment of the current situation. When this serial approach to human-automation interaction is taken, the same surprising (or not-so-surprising) effects continue to occur: human practitioners find themselves ineffectively trying to redirect brittle machines and/or unable to effectively assess or evaluate the work of the machine (e.g., Roth et al., 1987, Smith, McCoy, & Layton, 1997). These effects can lead to complete reliance on either the machine's or their own decision making, resulting in poorer decisions and less desirable outcomes than may have occurred with effective human-machine cooperation.

An alternative approach to designing human-automation architectures is to have the humans and machine work in parallel, rather than in serial. In this approach, both the human and the automation are working simultaneously to solve the problem. This type of human-automation architecture may have the automation monitoring, evaluating and/or

critiquing the human's progress in accomplishing the work, rather than vice versa.

Studies have shown that employing automation to critique the human's work can improve the joint decision making of the human-machine system (e.g., Guerlain, Smith, Obradovich, Smith & Svirbley, 1995; Guerlain, Smith, Obradovich, Rudmann, Strohm, Smith, Svirbely, & Sachs, 1999).

For UAS DAA systems, where time, or the running out of time, is a critical driver of the timeline to maintain well clear, an interactive critiquing architecture may not be appropriate. Instead, the automation could monitor how the pilot is responding to a potential conflict and act in extremis, i.e., when a resolution maneuver has not been executed in time (either due to a delay in the pilot's response or as a result of a late system detection of the conflict). In this architecture, the automation is monitoring both how fast the world is changing, i.e., how fast the potential threat is increasing, relative to the pilot's response. Essentially, the automation is monitoring for when the pilot has run out of time to avoid a loss of well clear or collision by executing a maneuver on their own. This is the same type of human-automation architecture that underlies envelope protection and automatic safety assist (e.g., automation braking) systems.

The decision to design for parallel work versus serial work is fundamentally a decision about where to apply design effort. In the latter case, the design work is typically concentrated in designing the machine capabilities relative to the system goals, and leaving the rest to the human. In the former case, design effort must also be put into building the supporting automation. By laying out all of the potential human-automation architectures at the beginning of the design or requirements development process,

different approaches, such as whether to implement a serial or parallel architecture, can be evaluated to determine which will be assessed as candidates for the future system.

Deriving Candidate HMI's Based on the Human-Automation Architecture

Once a set of potential human-automation architectures has been decided upon, they can be used to drive the HMI designs for the resulting prototype by analyzing the coordination and synchronization needs of each architecture. For each of the HMI configurations evaluated in the research program, the research questions essentially asked, “what information does the pilot need to execute the well clear functions?” and different HMI candidates were developed based on the informative vs suggestive vs directive display conceptualization. A better question would have been, “given potential human-automation architectures, how will the well clear function be coordinated between the pilot and machines?”. The last question looks to the human-automation architecture to identify the coordination and synchronization that is needed to effectively execute the main functions or goals of the technology. For the DAA system, the main function to be accomplished is maintaining well clear and the required HMI depends directly on how well clear is accomplished between the humans and machines. For coordination needs are different across the informative, suggestive and directive human-automation architectures. Specifically, what information does the human need from the machine in order to effectively finish their joint task of remaining well clear of another aircraft, given the capabilities that have been given to the machine. To do this, the pilot and machine need to coordinate their activities, i.e., the accomplishment of the three well clear functions, in a timely manner. That means that the machine must detect and send conflict

information to the pilot early and quickly enough that the pilot has time to respond and execute a maneuver. If the machine is going to determine a single or set of maneuvers, that also must be provided to the pilot in a timely and clear manner so that the pilot can evaluate the machine's solution(s). Finally, how will the HMI support the pilot's ability to coordinate with ATC.

While the research program reported in the previous chapters decided on a final HMI design (and underlying human-automation architecture) that was shown empirically to effectively support the pilot's ability to maintain well clear, the process relied on a lot of trial and error through iterations of testing. Display features (information element, maneuver guidance, and alerting) were added, subtracted, and modified based on the evaluation of pilot performance and an effort to improve human-automation and human-human cooperation. For example, the successive iterations of the alerting finally resulted in an alert structure that supported pilots' coordination both with other roles in the system (i.e., ATC) as well as the DAA system. By laying out these coordination needs within the human-automation architecture prior to developing candidate HMIs, the process of identifying the minimum DAA requirements may have been much more efficient. Given that the type of simulation testing that was utilized for the DAA HMI MOPS can be very resource intensive and therefore not available to all new technology endeavors, making the methodology more efficient can help to ensure that it is more likely to be employed.

Developing Scenarios from Human-Automation Architectures

Another benefit of laying out the human-automation architectures for potential prototypes to be evaluated is that they can be used to drive the scenarios used in the

evaluation phase. The coordination and synchronization links between humans and other humans and machines can be stressed by particular scenario designs. In this approach, scenarios are designed top down, based on a desire to explore the boundaries of the system and especially where coordination between humans and machines are likely to break down. For DAA systems, coordination is more likely to break down based on the increasing urgency inherent in the temporal nature of the well clear events when either the human or machine cannot work fast enough. Scenarios should then be designed to determine when a given prototype system and pilot “run out of time”. However, coordination can also break down when the pilot and DAA system simply cannot cooperate across any of the three well clear functions. In addition, the alerting and guidance algorithms that drive the HMI will also have boundaries that can be explored through scenario design. Scenarios need to explore various ways and reason for why breakdowns in the system occur. Sophisticated scenario design during simulation, flight testing and other analysis methods can help to uncover where the system is brittle and likely to fail quickly beyond its operating capacity, giving designers and regulators an opportunity to address potential issues and make the system more robust prior to deployment in the real operational setting.

The scenarios utilized in the DAA HMI research effort were largely derived bottom up, that is, based on the constraints on running a scenario given the simulation capabilities to produce scenarios with face validity to the UAS and operational community. While the goal of the research was to generate encounters that varied in their predicted distance at CPA, this is only one test of the timeliness of a pilot and DAA

system accomplishing the well clear function. Further, due to the dynamic nature of the full mission scenarios that were utilized, the alert levels and distance at CPA could not be guaranteed, nor could they be replicated consistently across trials. However, this limitation points more to the tradeoffs in research approaches more than to the specific scenario design. At the heart of the issue is that the scenarios were not designed based on fundamental questions about a specific DAA human-automation architecture and the coordination and synchronization demands that arise from it, or about testing the bounds of the prototyped system performance. While scenarios will always be subject to the constraints of the testing environment, starting scenario design from a more principled perspective intended to test the human-automation architecture, will result in a more robust evaluation of the prototype systems.

Conclusions

The research program presented in this dissertation for determining the DAA HMI minimum requirements represent substantial progress in how HMI standards are developed in safety critical systems. This is especially true in the aviation domain where standards have typically been derived from subject matter expertise or a single empirical investigation. Features of the methodology implemented by the DAA HMI research program have supported the needs of designers and regulators alike. The systematic evaluation of candidate DAA systems provided empirical human performance data on the effect of their associated HMIs on pilots' ability to accomplish the DAA task of remaining well clear. Analysis of this empirical data supported the comparison and selection of the minimum information elements, alerting and guidance requirements that

are predicted to support acceptable performance in the future operational world. In order to do this, the research effort was embedded in an environment that captured key relevant features of the future envisioned world, namely the temporal nature of DAA events and the interactions between the pilot and ATC. In summary, the research program helped to fulfill a pressing near term need by the UAS community, and especially, the FAA regulators to develop minimum HMI performance standards for a UAS DAA system. While the research program could not provide all of the necessary HMI requirements for the SC-228 MOPS, it contributed to the determination of arguably the most critical ones – alerting and maneuver guidance.

However, the research program is also just a starting point from which to build a comprehensive methodology, derived from CSE principles, that might be used to guide the design and regulation of new technologies. By basing the methodology upon the generation and analyses of potential human-automation architectures, prototype HMIs for future systems can be designed specifically for supporting the coordination and synchronization needs of the underlying architecture. Systematic evaluation of these prototypes, utilizing scenarios created specifically to test the bounds of the new system, then help to inform the final technical requirements which ensure effective human-automation cooperation. Thus, the DAA HMI research effort has provided the opportunity to not only learn from the direct results of the studies about how to support the ability of UAS pilots to maintain well clear through a DAA system, but also to reflect and learn from the process for developing HMI requirements for any new technology. The lessons learned here can not only be applied within the broader aviation domain,

including new technologies for manned aircraft, but also across a broader spectrum of domains where human-machine cooperation is necessary to accomplish work in complex systems.

References

- Bell, S., Drury, J., Estes, S., & Reynolds, C. (2012). GDTI: a ground station display of traffic information for use in sense and avoid operations. In *Digital Avionics Systems Conference (DASC), 2012 IEEE/AIAA 31st*, Williamsburg, VA, October 14-18.
- Billings, C.E. (1997). *Aviation Automation: The Search for a Human-Centered Approach*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Birmingham, H. P., & Taylor, F. V. (1954). A design philosophy for man-machine control systems. *Proceedings of the IRE*, 42(12), 1748-1758.
- Calhoun, G. L., Miller, C.A., Hughes, T. C., & Draper, M. H. (2014). UAS sense and avoid system interface design and evaluation. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Chicago, IL, October 27-31.
- Chamberlain, J. P., Consiglio, M. C., Comstock, J. R., Ghatas, R. W., & Munoz, C. (2015). *NASA controller acceptability study 1 (CAS-1) experiment description and initial observations*. NASA/TM-2015-218763, 2015.
- Chen, J. Y., Barnes, M. J., & Harper-Sciarini, M. (2011). Supervisory control of multiple robots: Human-performance issues and user-interface design. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 41(4), 435-454.
- Cook, S. P., Brooks, D., Cole, R., Hackenberg, M. D., & Raska, M. V. (2015). Defining Well Clear for Unmanned Aircraft Systems. In *Proceedings of AIAA Infotech@Aerospace*, Kissimmee, FL, January 5-9.
- Comstock, J. R., Consiglio, M. C., Ghatas, R. W., Vincent, M. J. (2016). Unmanned aircraft systems (UAS) collision avoidance, self separation, and alerting times (CASSAT) overall results.
- Dahm, W. J. A. (2010). *Report on technology horizons, a vision for air force science and technology during 2010–2030*, AF/STTR1001PR, United States Air Force.
- Dekker, S. W., & Woods, D. D. (2002). MABA-MABA or abracadabra? Progress on human–automation co-ordination. *Cognition, Technology & Work*, 4(4), 240-244.

Department of Transportation (2013). *Advisory Circular: Air Carrier Operational Approval and Use of TCAS II*. AC No. 120-55C.

Department of Transportation (2010). *Advisory Circular: Flightcrew Alerting*. AC No. 25.1322-1.

Draper, M., Calhoun, G., Ruff, H., Mullins, B., Lefebvre, A., Ayala, A., & Wright, N. (2008). Transition display aid for changing camera views in UAV operations. In *Proceedings of the first conference on Humans Operating Unmanned Systems (HUMOUS'08)*, Brest, France.

Draper, M. H., Miller, C. A., Benton, J., Calhoun, G. L., Ruff, H., Hamell, J., & Barry, T. (2013). Multi-unmanned aerial vehicle systems control via flexible levels of interaction: an adaptable operator-automation interface concept demonstration. In *Proceedings of the Infotech@ Aerospace 2013 Conference*, Boston, MA, August 19 – 22.

Draper, D. H., Pack, J. S., Darrah, S. J., Moulton, S. N., & Calhoun, G. L. (2014). Human-machine interface development for common airborne sense and avoid program. *Proceedings of the Human Factors and Ergonomics Society 58th Annual Meeting*, Chicago, IL, October 27-31.

Federal Aviation Act of 1958, Pub. L. No. 85-726 (1958).

Federal Aviation Administration, (2013a). *Integration of civil unmanned aircraft systems (UAS) in the national airspace system (NAS) roadmap, first edition*. Washington, D.C, 2013.

Federal Aviation Administration (2013). *Operational Use of Flight Path Management Systems: Final Report of the Performance-based Aviation Rulemaking Committee/Commercial Aviation Safety Team Flight Deck Automation Working Group*.

Feitshans, G. L., Rowe, A. J., Davis, J. E., Holland, M., & Berger, L. (2008). Vigilant spirit control station (VSCS)—‘The face of COUNTER’. In *Proceedings of AIAA Guidance, Navigation and Control Conf. Exhibition*, Honolulu, HI.

Fern, L., Flaherty, S. R., Shively, R. J., & Turpin, T. (2011). Airspace deconfliction for UAS operations. *Proceedings of the 16th International Symposium on Aviation Psychology*, Dayton, OH, May 3-5.

Fern, L., & Shively, R. J. (2011). Designing airspace displays to support rapid immersion for UAS handoffs. *Proceedings of the Human Factors and Ergonomics Society 55th Annual Meeting*, Las Vegas, NV, September 19-23.

Fitts, P. M. (1951). *Human engineering for an effective air-navigation and traffic-control system*, Washington D. C.: National Research Council.

Friedman-Berg, F., Rein, J. & Racine, R. (2014). Minimum visual information requirements for detect and avoid in unmanned aircraft systems. *Proceedings of the Human Factors and Ergonomics Society 58th Annual Meeting*, Chicago, IL.

General Operating and Flight Rules, 14 C.F.R. § 91, Sec. 91.113 (2004).

Guerlain, S. A., Smith, P. J., Obradovich, J. H., Rudmann, S., Strohm, P., Smith, J. W., & Svirebely, J. (1995). Dealing with brittleness in the design of expert systems for immunohematology. *Immunohematology/American Red Cross*, 12(3), 101-107.

Guerlain, S. A., Smith, P. J., Obradovich, J. H., Rudmann, S., Strohm, P., Smith, J. W., Svirebely, J., & Sachs, L. (1999). Interactive critiquing as a form of decision support: An empirical evaluation. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 41(1), 72-89.

Hancock, P. A., Jagacinski, R. J., Parasuraman, R., Wickens, C. D., Wilson, G. F., & Kaber, D. B. (2013). Human-automation interaction research past, present, and future. *Ergonomics in Design: The Quarterly of Human Factors Applications*, 21(2), 9-14.

Hart, S., & Staveland, L. (1988). Development of NASA TLX (task load index): Results of empirical and theoretical research, In P. Hancock and N. Meshkati (Eds.), *Human Mental Workload* (pp. 139-183). Elsevier, Amsterdam, 1988.

Hoekstra, J. M., van Gent, R. N., & Ruigrok, R. C. (2002). Designing for safety: the 'free flight' air traffic management concept. *Reliability Engineering & System Safety*, 75(2), 215-232.

Hollnagel, E. (1999). From function allocation to function congruence. *Coping with computers in the cockpit(A 00-40958 11-54)*, Aldershot, United Kingdom and Brookfield, VT, Ashgate Publishing, 1999, 29-53.

Johnson, W. W., Battiste, V., & Bochow, S. H. (1999). *A cockpit display designed to enable limited flight deck separation responsibility* (No. 1999-01-5567). SAE Technical Paper.

Jordan, N. (1963). Allocation of functions between man and machines in automated systems. *Journal of applied psychology*, 47(3), 161.

Kemeney, J.G. (1979). *Report of the President's Commission on the Accident at Three Mile Island*. The Commission, Washington D.C..

Mueller, E., Isaacson, D., & Stevens, D. (unpublished). *Air traffic controller acceptability of unmanned aircraft system detect-and-avoid thresholds*.

Muñoz, C., Narkawicz, A., Hagen, G., Upchurch, J., Dutle, A., Consiglio, M., & Chamberlain, J. (2015, September). DAIDALUS: detect and avoid alerting logic for unmanned systems. In *Digital Avionics Systems Conference (DASC), 2015 IEEE/AIAA 34th*

Murphy, J., & Kim, S. K. (2013). Live Virtual Constructive Distributed Test Environment Characterization Report.

Murphy, R., & Shields, J. (2012). Defense science board-task force report—the role of autonomy in DOD systems. *Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, Washington, DC, Tech. Rep CD, 1172*.

National Transportation Safety Board (1992). *Safety Recommendation*. A-92-39, Washington, D.C..

Norman, G. (2010). Likert scales, levels of measurement and the ‘laws’ of statistics. *Advances in Health Science Education, 15*(5), 625-632.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on, 30*(3), 286-297.

Perrow, C. (1984). *Normal Accidents*. New York: Basic Books.

Prevot, T. (2002). Exploring the many perspectives of distributed air traffic management: The Multi Aircraft Control System MACS. In *Proceedings of HCI-Aero*, Menlo Park, CA.

Rorie, R. C., & Fern, L. (2014). UAS measured response: the effect of GCS control mode interfaces on pilot ability to comply with ATC clearances. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Chicago, IL, October 27-31.

RTCA (2011). *Minimum operational performance standards (MOPS) for aircraft surveillance applications (ASA) system*. RTCA Inc., Washington, DC.

RTCA (2013a). *Draft DAA white paper – revision for FRAC*, paper No. 236-13/SC228-008. RTCA Inc., Washington, DC.

RTCA (2013b). *Terms of reference RTCA special committee 228 minimum performance standards for unmanned aircraft systems*, paper No. 109-13/PMC-1089. RTCA Inc., Washington, DC.

RTCA. (2015a). *DRAFT detect and avoid (DAA) minimum operational performance standards (MOPS) for verification and validation*. Paper No. 261-15/PMC-1400, RTCA Inc., Washington, D.C..

RTCA. (2015b). *Minimum operational performance standards (MOPS) for unmanned aircraft systems (UAS) detect and avoid (DAA) systems: Appendix A detect and avoid operational services and environment description (OSED) - DRAFT*. RTCA Inc., Washington D.C..

RTCA (2015c). *DRAFT minimum operational performance standards (MOPS) for unmanned aircraft systems (UAS) detect and avoid (DAA) systems: Appendix C – development of well clear*. RTCA Inc., Washington, D.C..

RTCA. (2015d). *Minimum operational performance standards (MOPS) for unmanned aircraft systems (UAS) detect and avoid (DAA) systems Appendix D: UAS maneuver performance requirements DRAFT*. RTCA Inc., Washington, D.C..

Santiago, C., Abramson, M., Refai, M., Mueller, E., Johnson, M., & Snow, J. (unpublished). *Java architecture for detect and avoid (DAA) modeling and extensibility (JADEM)*.

Santiago, C., & Mueller, E. R. (2015). Pilot evaluation of a UAS detect-and-avoid system's effectiveness in remaining well clear. In *Proceedings of the Eleventh USA/Europe Air Traffic Management Research and Development Seminar (ATM2015)*, June 23-16, Lisbon, Portugal.

Sarter, N. B. & Woods, D. D. (1994). Pilot interaction with cockpit automation II: An experimental study of pilots' model and awareness of the Flight Management System (FMS). *International Journal of Aviation Psychology*, 4, 1-28.

Sarter, N. B., Woods, D. D., & Billings, C. E. (1997). Automation surprises. *Handbook of human factors and ergonomics*, 2, 1926-1943.

Sheridan, T. B. (1987). Supervisory control. In G. Salvendy (Ed.), *Handbook of Human Factors/Ergonomics*. NY: Wiley.

Sheridan, T. B. (1997). Task analysis, task allocation and supervisory control. *Handbook of human-computer interaction*, 87-105.

Shively, R. J., Vu, K. P. L., & Buker, T. J. (2013). Unmanned Aircraft System Response to Air Traffic Control Clearances Measured Response. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, San Diego, CA, September 30-October 4.

Suarez, B., Kirk, K., & Theunissen, E. (2012). Development, Integration and Testing of a Stand-Alone CDTI with Conflict Probing Support. In *Proceedings of the AIAA Infotech@Aerospace Conference, Garden Grove, CA*,

Theunissen, E., Suarez, B., & de Haag, M. U. (2013). From spatial conflict probes to spatial/temporal conflict probes: Why and how. In *Digital Avionics Systems Conference (DASC)*,

Vu, K. P. L., Morales, G., Chiappe, D., Strybel, T. Z., Battiste, V., Shively, J., & Buker, T. J. (2013). Influence of UAS pilot communication and execution delay on controller's acceptability ratings of UAS-ATC interactions. In *Digital Avionics Systems Conference (DASC), 2013 IEEE/AIAA 32nd*, East Syracuse, NY, October 5-10.

Walker, D. (2014). *FAA position on building consensus around the SARP Well-Clear definition*. RTCA Inc., SC-228 WG1 Whitepaper, Washington, D.C..

Wing, D. J., Adamns, R. J., Duley, J. A., Legan, B. M., Barmore, B. E., & Moses, D. (2001). *Airborne use of traffic information in a distributed air-ground traffic management concept: experiment design and preliminary results*, NASA/TM-2001-21154, Hampton, VA.

Woods, D. D., & Hollnagel, E. (2006). *Joint cognitive systems: Patterns in cognitive systems engineering*. CRC Press.