

**Project Report
ATC-434**

A Model of Unmanned Aircraft Pilot Detect and Avoid Maneuver Decisions

**R.E. Guendel
M.P. Kuffner
D.E. Maki**

24 January 2017

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



This material is based on work supported by the Department of the Air Force under
Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001.

Approved for public release: distribution unlimited; Case Number 88ABW-2017-1669, 10 Apr 2017.

This report is the result of studies performed at Lincoln Laboratory, a federally funded research and development center operated by Massachusetts Institute of Technology. This material is based on work supported by the Department of the Air Force under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Department of the Air Force.

© 2017 MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Notwithstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

Massachusetts Institute of Technology
Lincoln Laboratory

A Model of Unmanned Aircraft Pilot Detect and Avoid Maneuver Decisions

*R.E. Guendel
M.P. Kuffner
D.E. Maki
Group 42*

Project Report ATC-434

24 January 2017

Approved for public release: distribution unlimited; Case Number 88ABW-2017-1669, 10 Apr 2017.

Lexington

Massachusetts

This page intentionally left blank.

ABSTRACT

For unmanned aircraft to share airspace with manned aircraft, extensive testing is first required to ensure that such vehicles can fly safely with manned traffic. Safe operation includes not only avoiding collisions with other traffic but also complying with the Federal Aviation Regulations to remain well clear of other traffic. One method for investigating the safety of unmanned aircraft operations is fast-time Monte Carlo simulation of encounters between unmanned and manned aircraft. As part of that simulation, one must model how the pilots of unmanned aircraft react to the encounters. To that end, an empirical, rule-based stochastic model of responses of unmanned aircraft pilots has been constructed based on data collected from a succession of human-in-the-loop experiments. This report details the main elements of that model and demonstrates its use in a safety analysis.

This page intentionally left blank.

ACKNOWLEDGMENTS

The authors wish to thank personnel at the Air Force Simulation and Analysis Facility for their role in building and conducting the human-in-the-loop experiments, in particular Sara Darrah and Jeremy Warren. We further thank all of those who volunteered to participate in the experiments.

We also extend our thanks to personnel at NASA Langley and Ames Research Centers for their assistance with the suggestive guidance algorithms, for generously sharing data collected in their own experiments, and for their feedback on our experiments.

This page intentionally left blank.

TABLE OF CONTENTS

	Page
Abstract	iii
Acknowledgments	v
List of Figures	ix
List of Tables	xi
1. INTRODUCTION	1
1.1 Operational Context for Detect and Avoid	2
1.2 Background	4
1.3 Modeling Approach	7
2. DATA COLLECTION	9
2.1 Experiment 1: Non-maneuvering Intruders	9
2.2 Experiment 2: Maneuvering Intruders	11
2.3 Experiment 3: Maneuver Updates	14
2.4 Key Experimental Findings	16
3. MODELING	21
3.1 Maneuver Timing	22
3.2 Maneuver Direction	24
3.3 Maneuver Magnitude	31
3.4 Deterministic Mode	34
3.5 Assumptions and Limitations	34
4. VALIDATION	37
5. SAFETY ANALYSIS EXAMPLE	43
6. CONCLUSION	45
Appendix A: Pilot Feedback and Human Factors Conclusions	47
A.1 Experiment 1: Non-maneuvering Intruders	47
A.2 Experiment 2: Maneuvering Intruders	52
A.3 Experiment 3: Maneuver Updates	54

TABLE OF CONTENTS
(Continued)

	Page
Appendix B: Model Parameter Summary	57
Glossary	61
References	63

LIST OF FIGURES

Figure No.		Page
1	Experiment 1 Pilot Workstation Displays	10
2	Experiment 2 Pilot Workstation Displays	13
3	Experiment 3 Pilot Workstation Display	16
4	All Pilot Responses for Encounter #1 in Experiment 2	17
5	Avoidance Maneuver Direction Choices Observed Over All Encounters	17
6	Observed Avoidance Maneuver Direction Choices by Maneuvering Protocol	18
7	Observed Avoidance Maneuver Direction Choices by Ownship Type	18
8	Pilot Model Architecture	22
9	Observed and Modeled Initial Delay	23
10	Observed and Modeled Update Delay for Corrective or Warning Alert	25
11	Example of Suggestive Guidance and Associated Minimum Suggestions	27
12	Pairwise Elimination Procedure for Choice of Maneuver Direction	28
13	Observed and Modeled Turn Magnitudes for Minimum-Guidance Compliant Maneuvers	32
14	Observed and Modeled Turn Magnitudes for Maneuvers not Following Minimum Guidance	33
15	Observed and Modeled Climb/Descend Magnitudes for all Vertical Maneuvers	33
16	One Response Observed in Experiment 2 and a Simulated Response	37
17	Comparison of Observed and Simulated Losses of Well Clear	38
18	Experiment 2 Encounter #9 Trajectories	39
19	Experiment 2 Encounter #12 Trajectories	40
20	Experiment 2 Encounter #16 Trajectories	40
A.1	Factors Contributing to Choice of Horizontal Maneuver in Experiment 1	47
A.2	Factors Contributing to Choice of Vertical Maneuver in Experiment 1	49
A.3	Influence of Maneuver Protocol on Maneuver Strategy	52
A.4	Factors Contributing to Choice of Horizontal Maneuver in Experiment 2	53

LIST OF FIGURES
(Continued)

Figure No.		Page
A.5	Factors Contributing to Choice of Vertical Maneuver in Experiment 2	53

LIST OF TABLES

Table No.		Page
1	Summary of HITL Experiments	9
2	Experiment 1 Alert Thresholds	12
3	Experiment 2 and 3 Alert Thresholds	13
4	Gauss-Markov Error Parameters for Experiment 3	15
5	Experiment 3 Encounters Segments	15
6	Pilot Model Maneuver Timing	24
7	Horizontal Pairwise Elimination Rules	30
8	Vertical Pairwise Elimination Rules	30
9	Maneuver Plane Rules	30
10	Maneuver Plane Selection Rule Weights	31
11	Pilot Model Maneuver Magnitude Distributions	34
B.1	Pilot Model Parameter Values	57

This page intentionally left blank.

1. INTRODUCTION

In 2012 Congress mandated that the Federal Aviation Administration establish rules for the integration of unmanned aircraft into the U.S. National Airspace System (NAS)—allowing them to operate together with existing manned air traffic [1]. Doing so requires the development of policies, procedures, and equipment to ensure that operations within the NAS can be conducted safely. For safety purposes, existing aviation regulations require that aviators “see and avoid” nearby traffic [2]. Without a pilot in the cockpit, unmanned aircraft rely on electronic surveillance systems to locate other air traffic and relay this information to the ground control station where the pilot can observe and act on it. These systems have been called Detect and Avoid (DAA) systems¹ and are the subject of a (forthcoming) Minimum Operational Performance Standard (MOPS) document [3] developed by RTCA Special Committee 228 (SC-228)².

The current operational assumptions of the MOPS have a remote pilot responsible for maneuvering the unmanned aircraft away from other air traffic. The question that must be asked, then, is whether a DAA system, as defined by the MOPS, is adequate to allow that remote pilot to operate his or her unmanned aircraft in the NAS safely. Safe operation in the NAS includes avoiding a collision with another aircraft and also maintaining what is considered a safe distance from other aircraft—*well clear* in regulatory terms.

The most straightforward way to test this is to build a system, integrate it into an aircraft, and then run flight tests. This approach, though, would be impractical and potentially dangerous. Such an approach would only test one implementation of a DAA system, not the requirements in toto; moreover, it would only test the system for the specific aircraft it was installed on; finally, it is simply not practical to flight test enough encounters to draw statistically sound conclusions about safety. Rather than flight testing, then, a standard tool for conducting safety analysis is modeling and simulation [4]. Specifically, fast-time Monte Carlo simulation facilitates time- and cost-efficient testing of numerous equipment configurations and encounter types. Whereas a flight test may entail on the order of one hundred encounters, these simulations typically consider millions of encounters, thereby allowing any weaknesses of or particular challenges for the system to emerge. The challenge, though, is since a pilot is responsible for choosing and executing avoidance maneuvers, a model of that pilot’s choices³ is needed that can realistically emulate how he or she uses the DAA system and responds when encountering other air traffic.

The model of the pilot is particularly important for analysis of the SC-228 DAA system concept because of the comparatively large role of the pilot. Safety analyses in the past at Lincoln Laboratory have typically assessed collision avoidance systems that are either directive or automated involving limited or no pilot interaction and decision-making. Directive systems identify a specific maneuver that the pilot should execute. The pilot’s contribution to this can be reasonably

¹ The term Sense and Avoid (SAA) is also widely used and is synonymous with DAA.

² RTCA, formerly the Radio Technical Commission for Aeronautics, is a Federal Advisory Committee sponsored by the the Federal Aviation Administration (FAA) and comprised of government and industry representatives. Through its Special Committees, RTCA develops guidance for aviation systems to ensure safety and reliability.

³ It is important to emphasize that the model referred to here is not a model of the pilot’s mental processes but one of the outcome of those processes; throughout this report, the terms “pilot model” and “model of the pilot” are used to refer to the latter unless otherwise specified.

modeled as a delay in executing that maneuver—see, for example, the analysis of using TCAS⁴ for collision avoidance on a Global Hawk unmanned aircraft [5]. Automated systems do not involve a pilot at all in the execution of the maneuver—avoidance maneuvers are determined and are executed automatically, often with the maneuver being continuously updated as the encounter situation evolves.

The pilot’s role in the SC-228 DAA system concept is much more considerable: the pilot must monitor the display, assess the situation, consider the suggestive guidance, and choose when and how to maneuver. Usually there are many reasonable maneuvers the pilot could choose—including the choice to not maneuver. The pilot can make multiple maneuvers depending on how the encounter evolves, but will probably not make as many as an automated system. A model is needed that encompasses the wide latitude given to the pilot to determine when and how to maneuver.

Filling that need is the focus of this work. Since the specific purpose of this model is to determine the adequacy of the DAA requirements in the MOPS, this report begins by reviewing the context and key human-machine interface requirements of the MOPS. That is followed by a review of prior approaches to modeling pilots and research on pilot preferences when maneuvering to avoid other traffic. This section closes with a summary of the approach taken for a pilot model that will suit the needs of DAA safety analysis. Section 2 describes the data collection experiments and summarizes some of the main findings, and then the design of the model based on that data is reviewed in Section 3. The report ends with a discussion of validation activities to show that the model is indeed realistic and an example of the usage of the model in a safety analysis.

1.1 OPERATIONAL CONTEXT FOR DETECT AND AVOID

It is important to understand the context within which this model must work, some of which has already been mentioned above. DAA system requirements are identified in a soon-to-be-published Phase 1 DAA Minimum Operational Performance Standard (MOPS). Some of the key aspects of the system are as follows:

- Unmanned (remotely piloted) aircraft heavier than 55 pounds
- Speeds exceeding 40 knots
- Instrument Flight Rules (IFR) flight plan and operations
- Transit operations in Class D, E, or G airspace en route to Class A or Special Use Airspace (SUA)
- On-board ADS-B⁵, radar, and active (transponder-based) surveillance
- Control station display incorporating DAA alerts and suggestive guidance

⁴ Traffic Alert and Collision Avoidance System, a collision avoidance system presently in wide use.

⁵ Automatic Dependent Surveillance – Broadcast, a surveillance system using satellite navigation information that each aircraft broadcasts to other aircraft in the vicinity.

- Collision avoidance system (e.g., TCAS) optional

The pilot model described here is applicable specifically to unmanned aircraft operating within this context. The operating context is expected to change in a future Phase 2 DAA MOPS, and pilot modeling will need to be readdressed.

As the unmanned aircraft nears another aircraft, the first priority is to stay well clear. For manned operations, well clear has not been strictly defined, but for unmanned operations staying well clear means the following condition does not occur:

$$0 \leq \tau_{mod} \leq \tau_{mod}^* \wedge HMD \leq HMD^* \wedge -h^* \leq d_h \leq h^* \quad (1)$$

where $\tau_{mod}^* = 35$ seconds, $HMD^* = 4000$ feet, and $h^* = 450$ feet [6]. The variable τ_{mod} is a measure of the time to closest point of approach adopted from the TCAS logic. HMD is the predicted horizontal separation at the point of closest approach, and h is the current vertical separation.

That region is difficult to visualize, so the DAA system incorporates *alerts* to help the pilot identify other aircraft on the traffic display that may violate the well-clear condition if both aircraft remain on their current course. Alerts are indicated both visually on the traffic display and aurally. The MOPS defines three levels of severity of alerts. A preventive alert is the lowest alert level, and it occurs primarily when the aircraft are separated vertically by 450–700 feet; it is intended to alert the pilot not to maneuver vertically to avoid causing a loss of well clear. The next alert level is the corrective alert. This alert indicates that a loss of well clear is predicted and maneuver is deemed necessary; however, there is still sufficient time to coordinate an avoidance maneuver with Air Traffic Control (ATC) in advance. Finally, a warning alert indicates that a well-clear violation is impending, an immediate avoidance maneuver is necessary, and coordination with ATC before maneuvering is not required.

Suggestive guidance works in conjunction with the alerts to assist the pilot in choosing a maneuver that will resolve that alert. When a nearby aircraft causes an alert, the suggestive guidance identifies maneuvers that will prevent loss of well clear by indicating where it is safe to fly horizontally and vertically. The suggestive guidance algorithm takes into account the dynamic capabilities of the ownship aircraft. It is intended to give the pilot a range of choices rather than the single specific maneuver of a directive system. In the event the well-clear violation cannot be avoided, the system displays *recovery guidance* that indicates maneuvers that will best resolve the well-clear violation. Suggestive guidance is also displayed for aircraft that are not currently alerting but will if the ownship maneuvers towards them. The caveat for the pilot using the suggestive guidance is that the future trajectory of the other aircraft is uncertain but current implementations assume non-accelerating flight by the intruder. Depending on what the other aircraft does, the guidance may change rapidly. Furthermore, even if the other aircraft holds steady, the guidance does not remain static; a maneuver designated as appropriate to maintain well clear at one moment is not necessarily still appropriate a few seconds later because the longer the pilot waits to maneuver, the larger the maneuver needs to be to remain well clear.

1.2 BACKGROUND

Prior to commencing this work, the technical literature was reviewed to determine whether a model useful for DAA safety analysis purposes already existed in part or in whole. Numerous types of pilot models were found and are briefly recapped here; none were directly suitable but some ideas did influence the model as will be noted in subsequent sections.

Broadly, there are two categories of pilot models in the literature. The first category is the representation of the pilot as a physical (or neuro-physical) system. These models tend to focus on *how well* a pilot performs a certain task. A prominent example of this is the McRuer and Krendel linear controller model and subsequent refinements of it [7]. These models focus on the physical interaction of the pilot with the aircraft's flight controls and are intended to inform control system design (e.g., pilot inputs needed to maintain a level course). Various other aspects of the pilot as a physical system have also been modeled—e.g., visual detection of targets [8], response time to a stimulus [9], workload [10], etc. None of these models are appropriate for this work because they do not address the decision-making aspects of piloting tasks.

The second major category of models, more relevant to this work, represents the pilot as a cognitive (or psycho-physical) system. Rather than trying to measure how well pilots perform a specific task, these models tend to focus on *what* a pilot does. Various approaches have been proposed, which are broken down here, roughly in order of applicability to this project, into the following subcategories:

- Intent prediction models
- Aerial combat models
- Delay models
- Regression models
- Bayesian network models
- Heuristic models

Intent prediction models have had a primary goal of understanding the air traffic likely to be faced, particularly in a terminal-area environment. That is, if you observe another aircraft, you might be able to infer something about its likely path. As such, this approach is better suited as a model of the intruder aircraft pilot than the ownship's. Lowe and How's work focused on identifying a Markov model of trajectory changes based on navigational intent [11]. Ogaard and Marsh used data mining to create a probabilistic model of manned aircraft behavior [12]. These models were included in this survey on the chance they could shed some light on the ownship pilot's decision-making when faced with an intruder with uncertain intentions; however, they are computationally intensive and hence better suited to refine the trajectory predictions for suggestive or directive guidance algorithms (to incorporate intruder accelerations, for instance) or for decision-making by autonomous systems.

Artificial intelligence (AI) techniques have been used to create models for simulating aerial combat for training pilots. Ernest et al. used genetic fuzzy trees [13] to create an AI called ALPHA to fly red-force aircraft in an aerial combat simulator. But their work does not show that the decisions of their AI are similar to human decisions; perhaps part of the challenge for an experienced fighter pilot of going head-to-head with ALPHA is that it does not behave like a human. A requirement of the model of DAA maneuver decisions, though, is that it reflects actual pilot behavior, and so this machine learning-approach was deemed unsuitable. Another model of air combat maneuvering decisions was developed by Virtanen et al., theirs based on multistage influence diagrams [14]. This approach can be adapted to human-derived preferences, but it is ultimately a deterministic approach that does not address the variability of responses given the same information about the conflict traffic.

A very simple model of a pilot that has been used in some studies is a response delay [15,16]. This sort of model is only useful when there is some other source providing the desired maneuver such as TCAS [17] or the next-generation collision avoidance system under development, ACAS X [18]. These directive systems give the pilot a specific maneuver that must be executed that is explicitly coordinated with the other aircraft where possible. The assumption of most delay models is that the pilot always follows the directed guidance; however, a more complex model presented by Lividas et al. incorporates the choice of the acceleration applied and permits TCAS advisories to be ignored [19]. These models fall short because they do not address the choice of maneuver.

Extending on the delay approaches above, Taguchi et al. propose using logistic regression to model “maneuver” decisions in a different domain: driving [20]. Their model uses data from a human-in-the-loop experiment (HITL) to simulate a human driver’s decision to make a turn across crossing traffic. The delay is thus dependent on real-time observed features rather than an independent distribution. Their approach clearly has parallels to pilot decisions about aircraft maneuvering. In an early phase of this work, logistic regression with two encounter features—predicted time to closest approach and horizontal range at closest approach—to model horizontal maneuver timing was explored and the resulting distribution was found to be unrealistic.

Lee and Wolpert explore modeling pilot decisions in the context of responses to resolution advisories from TCAS [21]. TCAS assumes that pilots will follow the advisories within five seconds of issuance, while studies cited by Lee and Wolpert suggest that actual responses vary more widely. Their model, built on a Bayes network, allows the pilot the choice of whether to maneuver or not given an RA and observational data. This is a stochastic approach that addresses the variability of responses but has not addressed how to learn the appropriate parameters to represent actual pilot preferences, nor is the framework yet available to choose from multiple maneuvers rather than the single maneuver directed by TCAS.

Heuristic approaches use a set of rules for decision-making; TCAS, in fact, uses a heuristic algorithm to choose a maneuver. Maki et al. introduced a model of pilots’ horizontal maneuver decisions using a “dynamic protection zone” for timing and a heuristic based on right-of-way rules for turn direction [22]. A subsequent live test showed a reasonably good fit between the model’s predictions and actual pilot decisions. However, this work has several shortcomings. Among those, variations in maneuver timing are not addressed. In addition, the right-of-way heuristic only ad-

addresses maneuvers in the horizontal plane—it neither offers a choice between climbing and descending nor a choice between a horizontal maneuver and a vertical maneuver. Finally, the model does not adequately address what happens if the first avoidance maneuver is insufficient to achieve the desired separation since continued reliance on right-of-way rules may lead to unrealistic reversals.

In addition to these models of pilot behavior, there have also been numerous studies characterizing pilot preferences when in conflict situations with other aircraft [23–26]. These studies typically entailed some form of human experiment. All (with the exception of the NASA experiments discussed further below) used pilots of manned aircraft, and the emergent preferences differed from study to study—perhaps because many of the studies used homogenous subject populations such as students from a single flight school. Summarizing the reported findings, though, the following general traffic avoidance preferences (or lack thereof) can be inferred:

- Single-axis vertical or horizontal maneuvers (e.g., climb) are preferred over multi-axis maneuvers (e.g., climb and turn together),
- Vertical maneuvers are preferred over horizontal maneuvers,
- No preference between right and left turns regardless of right-of-way rules, and
- Little preference for airspeed maneuvers.

Whether these preferences are also applicable to pilots of unmanned vehicles is an open question. Manned aircraft pilots generally have a three-dimensional head-on view of the surrounding airspace, while unmanned aircraft pilots must rely entirely on displays that typically give a two-dimensional overhead view. To begin to answer this question, studies of the preferences of air traffic controllers [27, 28], who view the airspace and air traffic from a similar perspective as unmanned aircraft pilots, were also explored. Those studies suggested that controllers have an even stronger preference for vertical maneuvers over horizontal maneuvers, postulating that horizontal maneuvers are more disruptive to the overall air picture. This highlights that while unmanned aircraft pilots and air traffic controllers may share a similar viewing perspective of the airspace, their overall priorities are different. Pilots want to steer their aircraft safely around threats while controllers are trying to organize all of the traffic in a given airspace.

One body of work that does address the maneuvers choices of unmanned aircraft pilots is a series of experiments conducted by NASA Ames Research Center [29, 30]. Those experiments were conducted in concert with the drafting of the DAA MOPS and were intended to identify requirements for the display. Consequently, those experiments explored various display concepts including information-only displays that showed only the relative location of nearby aircraft and suggestive displays that showed nearby aircraft and maneuvers that could be selected to stay well clear of it. Based on pilot response times with the various display concepts, NASA recommended to SC-228 that a suggestive display be required for DAA. This project began making use of the data NASA collected; however, purpose-built data collection experiments were ultimately preferred due to concerns that the flight plans the subject pilots were to fly confounded interpretation of any maneuvers observed—e.g., if flying a left-handed circuit, a preponderance of left turns may be attributable to the flight plan rather than an avoidance maneuvering preference.

1.3 MODELING APPROACH

The goal of this work is to model a pilot’s interpretations of a display showing nearby air traffic and his or her decisions of how to maneuver to avoid that traffic and remain well clear. To achieve the realism desired, the model is based on observed responses. The literature surveyed earlier gives a general idea of pilot preferences, but does not address the unique operational context, as described above, of detect and avoid by unmanned aircraft. Consequently, multiple data collection experiments have been conducted consistent with that operational context to form the empirical foundation of the model. Those experiments are described in more detail in the next section.

It is reasonable to expect, and this report will later show, that pilot responses to nearby air traffic are highly variable. When in an encounter with another aircraft, there are many possible choices of how to respond; the timing, direction, and size of the maneuver will vary from pilot to pilot and, possibly, from occasion to occasion for the same pilot. This variability may emerge from various factors including training, experience, and the pilot’s degree of risk acceptance. The model addresses this variability with a stochastic formulation whereby decisions at various levels are based on random draws from probability distributions. This differs from several of the approaches described above that are inherently deterministic—e.g., in the neural network approach laid out by Ernest et al., the AI identifies, through training, a function mapping a specific situation to the single best maneuver option. That work also highlights an important distinction: the model should represent the *actual* characteristics of pilot behavior not the desired or possible characteristics (i.e., an ideal pilot). Nothing guarantees that the AI will identify a maneuver that a human pilot would choose.

If the goal of this work is a model of representative and variable maneuver decisions, how can the decision itself be addressed? Unlike a directive system like TCAS, which provides a specific maneuver to the pilot who can then be modeled as a simple delay, the DAA operational context requires the individual pilot to decide when and how to maneuver. The maneuver decision could be approached through the field of decision theory. That field has a large body of literature, only some highlights of which will be address here.

Classical decision theory rests on the notion of rational choice, asserting that a rational decision maker should choose the option that is best for itself, knowing all possible choices and possessing universal knowledge of the relative efficacy of each [31]. This effectively treats a decision as a global optimization problem. Decision-making under this paradigm is termed substantive rationality. Many researchers argue, however, that these assumptions are too broad to accurately represent human decision-making [31–33]. Pomerol [34] outlines the major objections:

- All possible choices rarely come to mind.
- Knowledge of the consequences of each choice is usually fragmentary.
- It is rarely possible to determine a complete ordering of the consequences of all of the choices.

These objections are reinforced by psychological research demonstrating that human decision-makers are subject to numerous cognitive biases that influence their choices [35]. To address

these concerns, researchers have proposed models utilizing so-called bounded rationality [36], e.g., satisficing [31], prospect theory [37], fast and frugal heuristics [38], etc.

With any of these methods, detailed modeling of how a pilot interprets and reacts to what he or she sees on a display is not trivial. The approach taken for this model uses the notion of bounded rationality but recognizes that the ends do not require modeling of the actual decision-making process and the cognitive mechanisms employed in that process, only the outcome of this process. This simplifies the challenges of the model considerably. Insights of cognitive scientists and psychologists can be used to influence and/or validate the model, but it is not necessary to accurately determine and represent all the complexity of a pilot's thought process. The model will be referred to, then, as a decision model rather than a decision-making model. To model those decisions, behavioral patterns identified by the data collection experiments are used.

2. DATA COLLECTION

The model was developed iteratively, with its structure and parameters informed by data collected in a series of three HITL experiments. Table 1 outlines the major features of each experiment, with detail and a summary of the results following. Data provided by NASA Ames Research Center was also used to build the model, particularly for modeling coordination with Air Traffic Control.

The trajectory of each aircraft, the alert status, the suggestive guidance being displayed, and maneuvers commanded by the pilot were recorded in all of the experiments. The first two experiments also recorded the timing of simulated ATC coordination prior to maneuvering. Subjects were encouraged to think aloud during the experiment and provide feedback afterward. A catalog of that feedback and a summary of human factors conclusions drawn from the experiments are included in Appendix A.

TABLE 1

Summary of HITL Experiments

	Experiment 1	Experiment 2	Experiment 3
Guidance Algorithm	OmniBands	DAIDALUS	DAIDALUS
Level of Fidelity	Medium	Medium	Low
Number of Subjects	26	27	50
Number of Encounters	23	18	50
Encounter Duration	4–6 min	3–4 min	50–70 s
Surveillance	Perfect	Perfect	Perfect/Imperfect

2.1 EXPERIMENT 1: NON-MANEUVERING INTRUDERS

The first experiment was conducted in July 2015 in cooperation with the Air Force Simulation and Analysis Facility (SIMAF). Twenty-six UAS pilots from various Air Force and NASA locations participated as subjects. The unmanned aircraft flying time for this group ranged from 50 to 9000 hours with a median of 800 hours.

The experiment was conducted using a portable aircraft encounter simulator developed by SIMAF. The pilot workstation comprised two monitors, shown in Figure 1, a keyboard and a mouse. The left-hand monitor showed a moving map and the Primary Flight Display (PFD) with the aircraft’s attitude, heading, altitude, and speed. The pilot’s ATC communication controls and ability to declare all clear (i.e., end the simulation) were also on the PFD. The right-hand monitor,

the DAA display, showed nearby traffic and suggestive guidance. The traffic's location on the display and the guidance computation used perfect surveillance data, i.e., no errors in the traffic's position and velocity and no limits on detection.



Figure 1. Experiment 1 pilot workstation displays. The multifunction display on the left includes a moving map and the PFD. The DAA display is on the right.

One challenge of this project was that the requirements in the MOPS were in flux while the experiments were on-going. Alerting and guidance for Experiment 1 was consistent with the then-current draft of the MOPS. The alerting thresholds used are summarized in Table 2. The OmniBands algorithm from NASA Ames Research Center [30] provided the suggestive guidance for the display. Suggestive guidance did not include recovery guidance because it had not yet been introduced as a requirement, which meant that guidance saturated, turning all red, if a loss of well clear could no longer be avoided.

Maneuvers were commanded by entering the desired heading and/or altitude in text boxes on the PFD and clicking an 'execute maneuver' button. Pilots could choose a horizontal maneuver, a vertical maneuver, or both simultaneously (a multi-axis or "combo" maneuver) and were permitted to make as many successive maneuvers as desired. Horizontal maneuvers were executed at $3^\circ/\text{s}$ and vertical maneuvers at 500 fpm with a generic ownship dynamic model, consistent with the SC-228 assumptions.

Pilots were trained to coordinate with ATC prior to maneuvering if possible via simulated ATC interaction. When ready to maneuver, the pilot clicked a button on the PFD, and after a delay drawn randomly from a gamma distribution with mean 11 seconds (fit to ATC coordination times from NASA's PT5 experiment [29]), a message was displayed indicating that the pilot may proceed. There was no voice communication and all maneuver requests were approved. Flight controls were *not* disabled during this delay, so the pilot was able to command maneuvers before receiving authorization if he or she chose to do so.

The experiment was conducted one pilot at a time. Each test commenced with a training briefing to explain the controls, display, operational procedures (e.g., coordinating maneuvers with ATC), and the objective of avoiding loss of well clear while minimizing course deviations. After the training briefing, the pilot was guided through several training encounters for hands-on familiarization with the displays and controls before beginning the experiment.

In the experiment, each pilot flew a generic UAS ownship (with dynamics consistent with the minimum aircraft performance defined in the DAA MOPS) in a total of twenty-three single-intruder encounters. Without intervention from the pilot, the ownship followed a predetermined flight path (due north). The intruder aircraft and all background traffic followed predetermined trajectories. The intruder followed a straight-line trajectory in all but three of the encounters; the remaining three encounters included simple maneuvers by the intruder. All encounters occurred in Class E airspace below 10,000 feet.

One within-subjects variable was included in the experiment: maneuver protocol. The encounters were divided into two sets to explore how responses differed depending how the pilots were instructed to respond when nearing an intruder. In one set, the pilot was asked to wait for a DAA alert before commanding any maneuvers, and in the second set, the pilot was permitted to maneuver at will. The order of the sets, the constituent encounters, and the order of encounters within the sets were counterbalanced across subjects and were not known to them in advance.

Overall, this experiment is considered to be at a medium level of fidelity for several reasons. First, the experiment was encounter-centric rather than mission-centric: the simulation stops and resets after each encounter rather than presenting the pilots with a continuous series of encounters along a planned flight path as in the higher-fidelity NASA HITLs. In addition, because it was believed that the encounter-centric approach (with relatively short encounters) made it unnecessary, the pilots had no secondary tasks to perform. Finally, the intruder aircraft followed scripted trajectories and did not respond to ownship maneuvers; in NASA's experiments, on the contrary, simulation confederates actively flew the intruder aircraft. These choices in the experiment design were made deliberately under the assertion that they do not fundamentally alter the behavior under study; because there is less control over the exact encounter conditions in the higher fidelity simulation, it is harder to compare responses between pilots and to distinguish avoidance maneuvers from navigation maneuvers.

At completion of the experiment, pilots responded to a brief post-test regarding maneuver decision-making. The total test time, including training, each block of encounters with a break in between, and the survey, was approximately 2.5 hours. A demographics survey was completed separately to record each subject's flying experience. Kuffner et al. [39] gives more details about the experimental setup and outcome.

2.2 EXPERIMENT 2: MANEUVERING INTRUDERS

The second experiment was conducted in cooperation with SIMAF with the main goal of collecting responses to maneuvering intruders using varying ownship dynamic models with an updated interface consistent with the latest MOPS requirements. This included switching to the

TABLE 2

Experiment 1 Alert Thresholds

Alert Level	Lookahead	τ_{mod}^*	HMD*	h*	Aural Alert
Proximate Traffic Advisory	85 s	35 s	1.5 nm	1200 ft	None
Preventive Alert	75 s	35 s	1.0 nm	700 ft	Tone
Corrective Alert	75 s	35 s	0.75 nm	450 ft	Tone
Warning Alert	25 s	35 s	0.75 nm	450 ft	Tone

DAIDALUS algorithm from NASA Langley Research Center for DAA alerting and suggestive guidance [40]. Twenty-seven UAS pilot participants completed the experiment in May 2016. The median unmanned aircraft experience of this group was 500 flight hours. Eight of these pilots previously participated in Experiment 1.

The pilot workstation again comprised two monitors, shown in Figure 2. The left-hand monitor was updated from the first experiment to include a system status panel and a chat window. Those new components were used for the secondary tasks introduced in this experiment. During an encounter, the pilot periodically received messages via the chat window to perform one of several system checks (e.g., check and report on fuel level). The pilot’s accuracy on secondary tasks was not assessed. The DAA display, incorporating a new DAA status panel, appeared on the right-hand monitor. The suggestive guidance had a similar look to the previous experiment but some notable differences. Now powered by the DAIDALUS algorithm, the display no longer differentiated between preventive, corrective, and warning guidance, instead showing any heading or altitude predicted to lead to a DAA alert in red. The new display also incorporated recovery guidance, indicating the best headings to regain well clear after it has been lost in dashed green.

Since the first experiment, alerting thresholds in the draft MOPS had been revised. These revisions included the elimination of the proximate traffic advisory. The thresholds used in Experiment 2, shown in Table 3, were adjusted accordingly.

Several other changes were introduced in this experiment. First, the pilots were instructed to always wait for an alert before maneuvering. As already noted, secondary tasks were incorporated to force the pilots to monitor both displays. Finally, an aural notification of ATC’s maneuver approval was added due to concern that pilots may not immediately see the text-only approval used in Experiment 1. Again the pilots were instructed that their objective was to avoid a loss of well clear while minimizing deviations from course.

The experiment consisted of eighteen encounters with a single intruder. As a within-subjects variable, all pilots flew two types of ownship aircraft in the experiment. Half of the encounters were flown with the generic UAS ownship from the first experiment, and the other half were flown with a UAS ownship similar to either a MQ-9 Reaper (referred to as the “MALE” configuration) or a RQ-4

Global Hawk (referred to as the “HALE” configuration), depending on the pilot’s real-world aircraft qualifications. Both the MALE and HALE configurations offered greater dynamic capabilities over the generic ownship. The intruder did not maneuver in seven of the encounters, and the remaining eleven encounters included varying degrees of intruder maneuvering. The order of the ownship models and the order of the encounters was counterbalanced over all pilots. Completion of the experiment, including training and a post-test survey, took about 2 hours.



Figure 2. Experiment 2 pilot workstation displays. The multifunction display on the left includes a moving map, the PFD, system status indicators, and the chat window. The DAA display is on the right.

TABLE 3

Experiment 2 and 3 Alert Thresholds

Alert Level	Lookahead	τ_{mod}^*	HMD*	h^*	Aural Alert
Preventive Alert	60 s	35 s	0.9 nm	700 ft	“Traffic, Monitor”
Corrective Alert	60 s	35 s	0.9 nm	450 ft	“Traffic, Avoid”
Warning Alert	35 s	35 s	0.9 nm	450 ft	“Traffic, Maneuver Now”

2.3 EXPERIMENT 3: MANEUVER UPDATES

The final experiment was conducted at Lincoln Laboratory with an interface consisting of a single display. The experiment was focused specifically on the timing of maneuver updates. It is difficult to collect sufficient data on update maneuvers because the need for updates is contingent on the initial maneuver choice—many times an update maneuver is not necessary. This experiment aimed to overcome a data gap by having a large pool of subjects complete many short encounters. The experiment also examined the effect of surveillance error on pilot responses—the previous experiments used only perfect surveillance. Fifty pilots with time in civil, commercial and/or military *manned* aircraft participated in the experiment in May–June 2016. The median flying experience for this group was 350 hours.

The display for this experiment incorporated the key components of the previous experiments, the PFD and the DAA display, onto a single screen. The encounters were brief and there was no specific route to fly, so the map was eliminated. With the brevity of the encounters, the secondary tasks of Experiment 2 and the associated display components were also eliminated. DAIDALUS was used for alerting and guidance with the same alerting thresholds as in Experiment 2 (see Table 3). Figure 3 shows an example of the display.

Each participant completed a total of 50 encounters with an average duration of 60 seconds. Those encounters were broken down into four segments as summarized in Table 5. The first segment of ten encounters included a mix of maneuvering and non-maneuvering intruders. The display used perfect surveillance (no errors or detection limits) for alerting and guidance. Pilots chose initial and update maneuvers as necessary. The second segment of ten encounters was similar only with imperfect surveillance approximating the performance of an ADS-B system. For this, the position and velocity errors were modeled as independent Gauss-Markov processes:

$$\delta x_k = \delta x_{k-1} e^{-\Delta t/\tau} + \mathcal{N}(0, \sigma^2(1 - e^{-\Delta t/\tau})) \quad (2)$$

The parameters used with this model are summarized in Table 4. As surveillance error increases, the alerting and suggestive guidance can become difficult to use: the alert state can toggle back and forth and suggested maneuvers can change quickly. The error level for this experiment was intentionally kept low to avoid causing mistrust of the system.

For the third and fourth segments, a procedural change was introduced. Rather than choosing their own initial maneuver, the pilots took control only after a prescribed initial maneuver was started. This procedure was intended to focus more closely on the timing of update maneuvers. Both segments included fifteen encounters, all with a maneuvering intruder; alerting and guidance used perfect surveillance in the third segment and imperfect surveillance in the fourth. The low-end generic UAS ownership was used in all fifty encounters. The pilots were instructed to wait for an alert before maneuvering in all encounters. There was no coordination with ATC in this experiment. It took about 90 minutes for each participant to complete the experiment.

There were some risks to the usefulness of this experiment, foremost the use of manned aircraft pilots. The resulting data was only considered valid and useful if similar features were observed to the previous two experiments using unmanned aircraft pilots.

TABLE 4**Gauss-Markov Error Parameters for Experiment 3**

State Dimension	σ	τ	Δt
Position East	300 ft	10 s	1 s
Position North	300 ft	10 s	1 s
Altitude	0 ft	0 s	1 s
Velocity East	12 ft/s	10 s	1 s
Velocity North	12 ft/s	10 s	1 s
Altitude Rate	1 ft/s	10 s	1 s

TABLE 5**Experiment 3 Encounters Segments**

Segment	No. of Encounters	Intruder Behavior	Surveillance	Ownship Initial Maneuver
1	10	Maneuvering & Non-maneuvering	Perfect	Pilot Choice
2	10	Maneuvering & Non-maneuvering	ADS-B	Pilot Choice
3	15	Maneuvering	Perfect	Scripted
4	15	Maneuvering	ADS-B	Scripted

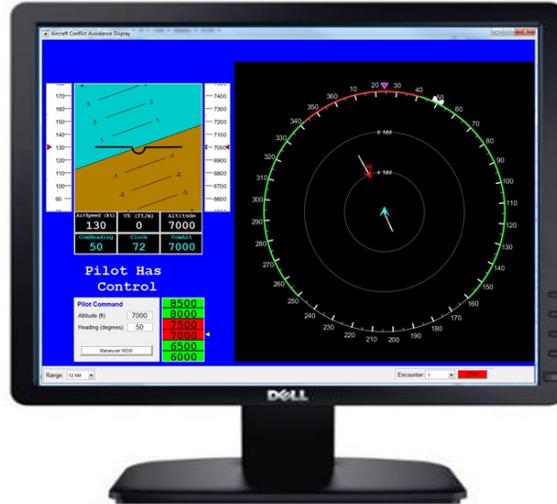


Figure 3. Experiment 3 pilot workstation with PFD and traffic on a single display.

2.4 KEY EXPERIMENTAL FINDINGS

The three experiments yielded 800, 703, and 3299 recorded avoidance maneuvers, respectively. As expected, the responses were highly variable. Figure 4 shows all of the responses to encounter 1 in Experiment 2—all of which were successful in remaining well clear. Variations in maneuver timing, direction, and plane are all evident.

Figure 5 shows the overall preferences for maneuver direction. Of note from this first view of the data is the small number of multi-axis maneuvers that were selected, a result consistent with the earlier studies mentioned above. Filtering the data yields additional insights. Figure 6 reveals that pilot behavior was starkly different depending on whether they chose their initial maneuver before an alert or after an alert. Figure 7 shows that maneuver preferences changed depending on ownship capabilities. Recall that ownship type was a within-subjects variable—each pilot flew both the generic ownship and the MALE or HALE ownship.

One question was whether the head-on out-the-window view of a pilot in a cockpit might lead to different preferences than the overhead view of a pilot operating with a traffic display. Experiment 3 suggests that perspective does matter. As noted above, other researchers have observed a general preference for vertical maneuvers over horizontal among manned aircraft pilots. Experiment 3 involved mostly manned aircraft pilots operating without a head-on view, just the overhead view of a DAA display. As Figure 5 shows, there was a strong preference for horizontal maneuvers. In part, that preference may be that the suggestive guidance showed that vertical maneuvering would not be effective. However, after filtering out those instances, out of the remaining 1231 instances with both horizontal and vertical suggested maneuvers, 58.8% of the maneuvers were horizontal, 34.7% were vertical, and 6.5% were multi-axis (a significant preference for horizontal maneuvers over vertical maneuvers with $p < 0.01$). So the same type of pilot (manned aircraft) has different preferences depending on the perspective. It's possible that the outcome of a horizontal

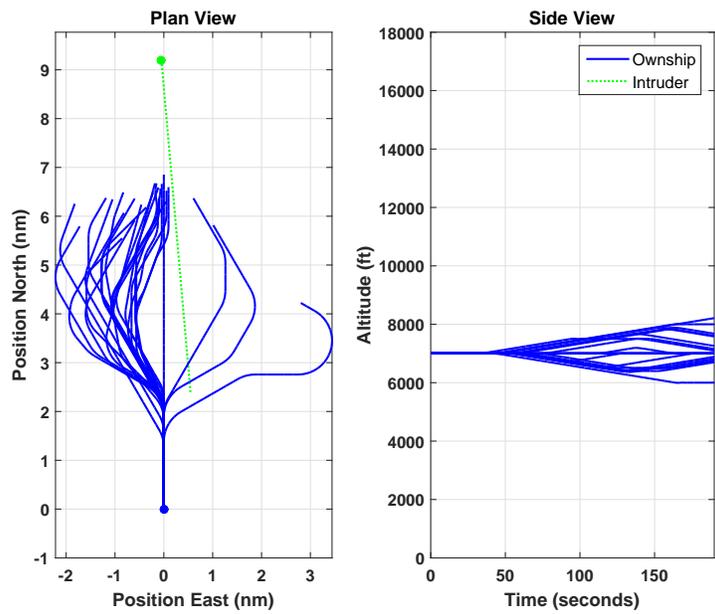


Figure 4. All pilot responses for encounter 1 in Experiment 2.

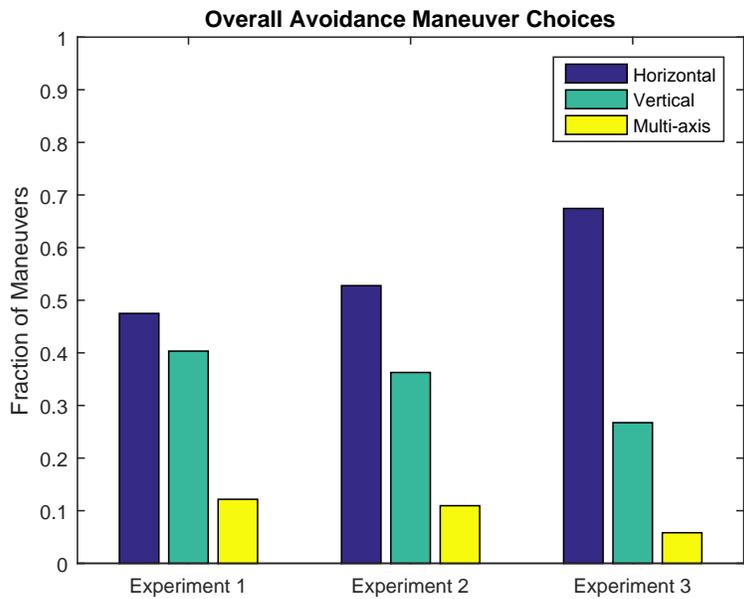


Figure 5. Avoidance maneuver direction choices observed over all encounters.

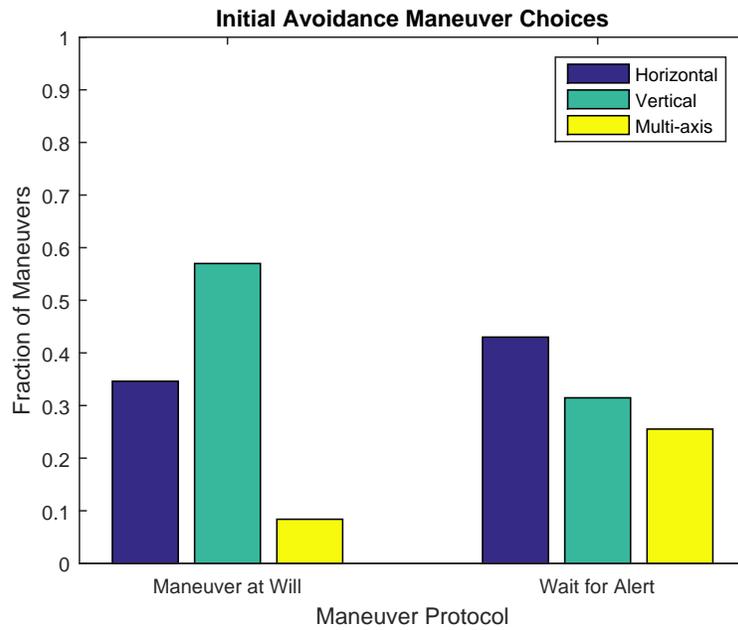


Figure 6. Observed avoidance maneuver direction choices by maneuvering protocol (Experiment 1).

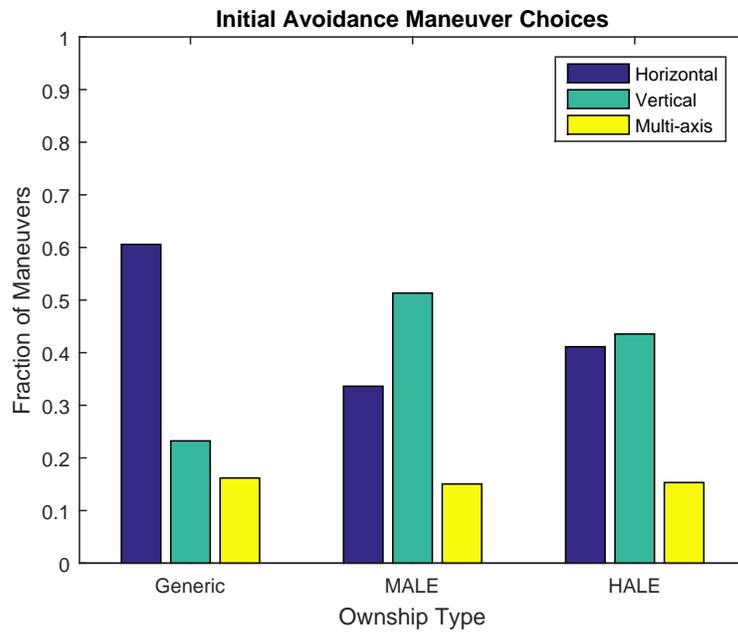


Figure 7. Observed avoidance maneuver direction choices by ownship type (Experiment 2).

maneuver can be more easily visualized on a map-like display (or a vertical maneuver may be difficult to visualize) while the opposite may be true from a cockpit. That the manned pilots in Experiment 3 demonstrated similar preferences for maneuver plane to the unmanned pilots in the other experiments lends support to use of that data for model-building.

Interestingly, the same conclusion cannot be so easily drawn for the unmanned aircraft pilots in Experiment 2. In that case, the preference for horizontal maneuvers was evident in the trials conducted with the generic ownship aircraft, but the preference was weakened or reversed when the pilots flew one of the ownship models similar to their real-world experience. This suggests that the hypothesis above is subject to various confounding factors. Real-world training may account for some of the differences as may awareness of the capabilities of the aircraft being flown. A further confounding factor was suggested by Experiment 1: when pilots were permitted to maneuver before any DAA alert had occurred, they preferred to maneuver vertically ($p < 0.01$). It is worth noting that in this case there was no suggestive guidance on the display (because the intruder was still too distant) for the pilot to use in decision-making.

The data from all of the experiments indicates that the previously reported preference for single-axis maneuvers holds here, too. Depending on the experiment and how the data were filtered, multi-axis maneuvers (horizontal and vertical maneuvers commanded simultaneously) accounted for 5–20% of all maneuvers. This result is understandable since under time pressure the subjects may prefer to quickly execute a single-axis maneuver than to spend the extra time deciding on and entering a desired heading and altitude. The unmanned aircraft pilots in Experiments 1 and 2, perhaps more familiar with flying their aircraft via keyboard, were somewhat more likely to use multi-axis maneuvers than the manned aircraft pilots in Experiment 3.

Exploration of response time characteristics revealed that pilots were slower to respond when the intruder was climbing or descending than when it was level in altitude. This was first identified in Experiment 1, in which the initial response times (time between first DAA alert and the beginning of ATC coordination) with a climbing or descending intruder had mean of 8.62 seconds while with a level intruder the mean was 4.35 seconds. A two-sample t-test indicates this is a significant difference with $p < 0.01$. This result was repeated, albeit with closer means, in Experiment 2, where the mean initial response time with a climbing or descending intruder was 4.5 seconds as compared to 3.63 seconds with a level intruder (significant with $p < 0.04$). The difference in response time is not surprising. Changes in the intruder heading are more salient on the display than changes in altitude. The pilot is left to mentally determine when and where an intruder's vertical trajectory will intersect the ownship's.

Another response time characteristic was revealed by Experiment 2: responses are quicker when recovery guidance is displayed (i.e., a loss of well clear has occurred or can't be avoided). The average time to maneuver after a DAA warning alert was 14.61 seconds (including initial and update maneuvers), while the average time to maneuver after the appearance of recovery guidance was 6.41 seconds. The number of instances of the latter was small ($N=34$), but the difference is significant with $p < 0.01$. Note, however, that these are not always independent events since a warning alert usually precedes the appearance of recovery guidance. This result is a reassuring one since the purpose of the recovery guidance is to facilitate regaining well clear as quickly as possible.

A final observation from Experiment 3 is that the pilot responses were not significantly different between the perfect and imperfect surveillance cases. This result *does not* suggest that surveillance noise does not matter. The surveillance errors in the experiment were small, and some participants mentioned that they hadn't even noticed the change between perfect and imperfect surveillance. This suggests that the pilots were easily able to integrate small perturbations on the display. Further research is needed to determine what level of jitter leads to different response characteristics.

3. MODELING

In the operational context of the DAA Phase 1 MOPS, the real-world process of avoiding traffic that is to be modeled is as follows:

1. Surveillance systems onboard the unmanned aircraft detect nearby aircraft.
2. That information is transmitted to the ground station where it is processed and the traffic is depicted on a DAA display for the pilot's situational awareness.
3. When the pilot notices or is alerted to the presence of an intruder aircraft that is on a path that may lead to a loss of well clear, he or she chooses an avoidance maneuver, using the suggestive guidance shown on the DAA display as a decision aid if desired.
4. The pilot then contacts ATC to coordinate the desired maneuver.
5. After given authorization for that maneuver, the pilot maneuvers the aircraft.
6. Steps 3 through 5 are repeated until the conflict is resolved.
7. When clear of conflict, the pilot coordinates with ATC to maneuver back to the original course.

There is one exception to this process: when the DAA system issues a warning alert, the pilot should take immediate action without contacting ATC first. (In this event, ATC should be advised after the conflict is resolved.)

For modeling this process, there are two primary decisions the pilot must make: when to maneuver and how to maneuver. The latter can be decomposed into two subdecisions, which direction to maneuver and how much to maneuver. The model was designed around each of these decisions and so is structured around three primary subfunctions: maneuver timing, maneuver direction, and maneuver magnitude. The latter two elements are referred to collectively as maneuver selection. Each of these elements are described in detail in the following sections.

Parameter values for each of the elements were derived from the data collected in Experiment 2 unless otherwise noted. Experiment 1 data was used for an initial version of the model, but with subsequent changes to the alerting requirements and the suggestive guidance algorithm, those data were not used for the final model parameterization. Experiment 3, with a much larger set of subjects and encounters at lower fidelity, was intended mainly to confirm any trends observed in Experiment 2 that had limited supporting data. In the end, sufficient update maneuvers were observed in Experiment 2 that the Experiment 3 data did not have to be relied upon directly. Also note that for modeling simplicity not every feature observed in the data was captured in the model.

The resulting overall model architecture is depicted in Figure 8. The model runs continuously on a 1 Hz cycle. Inputs to the model are the current DAA alert level and the suggestive guidance. The alert level is both a trigger to activate the model and a decision parameter. The model outputs appropriately timed maneuver commands.

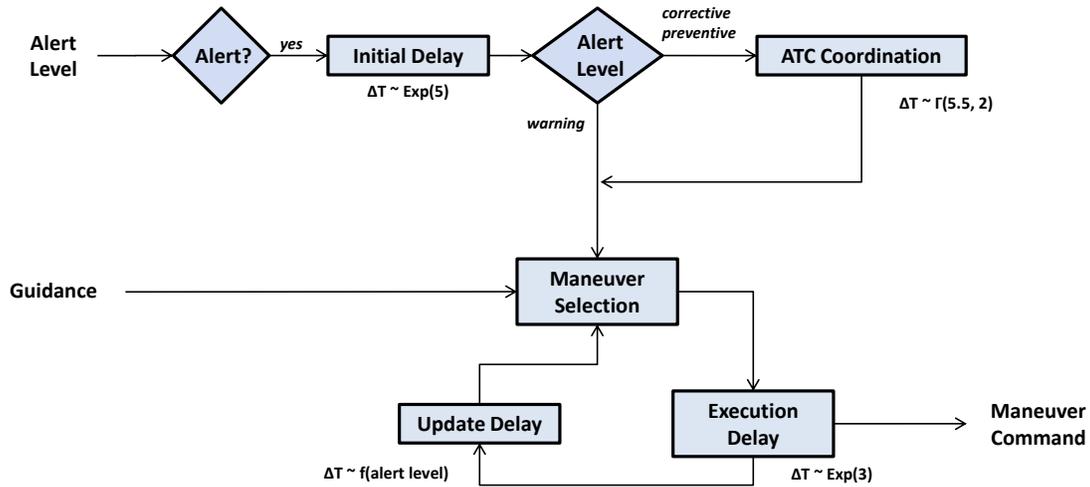


Figure 8. Pilot model architecture.

3.1 MANEUVER TIMING

The first subfunction of the model is timing, which is viewed as a delay between some triggering event and a maneuver. The assumed triggering event is the occurrence of a DAA alert. The first experiment demonstrated that sometimes pilots opt to maneuver before an alert if permitted, but no distinguishing encounter features prompting that early response could be identified. Instead, the model takes the conservative (and MOPS-consistent) course of assuming avoidance maneuvering does not commence until an after alert is issued. The alert must still be active for an initial maneuver to be selected, but after the initial maneuver, additional maneuvers may be selected in the absence of an ongoing alert.

The total delay is decomposed into several successive delay processes. First is an *initial delay* representing the time it takes for the pilot to notice the alert and formulate a plan. Next there is a *coordination delay* representing the time it takes to contact ATC and gain approval for the desired maneuver from an air traffic controller. The maneuver selection processes (described in the next two sections) are executed at this point and are followed by the *execution delay*, which represents the time it takes for the pilot to enter the desired maneuver at the workstation and transmit it to the aircraft, at which point the aircraft begins executing the commanded maneuver according to its control system logic. There is one caveat to this succession of delays: if a DAA warning alert occurs, the model skips the coordination delay or interrupts it if it has already begun. Finally, since a single maneuver may not be sufficient to resolve the conflict, especially if the intruder is maneuvering, subsequent maneuvers may be chosen after an *update delay*. The experimental data were used to

parameterize probability distributions for each of these delay processes. The distributions for each of the delays are summarized in Table 6, and the rationale for those model choices are detailed in the next several paragraphs.

The initial delay is modeled with an exponential probability distribution with mean 5.0 seconds, the approximate average over Experiments 1 and 2. Initial delays were shorter in Experiment 2, with mean 3.9 seconds. There was no difference in alert annunciation or coordination between the two experiments that would account for the faster response time; in fact, the addition of secondary tasks in Experiment 2 were expected to have slowed responses. The more conservative 5-second mean was chosen because of two concerns. First, in the experimental setting with a clear goal of maneuvering to avoid other air traffic, the pilots may have been primed for a quick response. Second, since ATC coordination in the experiments was simulated, with no actual interaction with a controller, the pilots may have learned to initiate the coordination clock while still deciding on a maneuver. The Experiment 2 data and the modeled initial delay are shown in Figure 9.

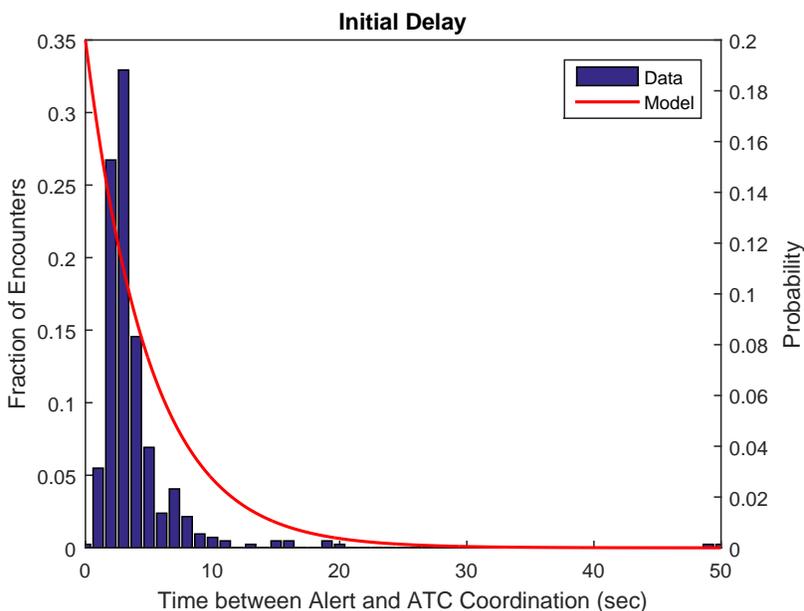


Figure 9. Observed and modeled initial delay (Experiment 2 data).

The coordination delay model is based on data from NASA Ames Research Center’s Part Task 5 experiment, which used live radio communication between each subject-pilot and a confederate controller [29]. The audio recordings provided by NASA were parsed to find the elapsed time between the beginning of the pilot’s request and the end of the controller’s response. Those data were best fit with a gamma-distribution with mean 11.0 seconds. The model skips or interrupts the coordination delay if the alert is at or increases to the warning level.

The execution delay is modeled with an exponential distribution with mean 3 seconds. After observing unreasonably long delays between maneuver approval and maneuver execution in the first experiment, an aural indication for the pilot to precede with his/her desired maneuver at the end of the simulated ATC coordination was added for Experiment 2. This is more consistent with actual communication with a controller.

The update delay, too, is modeled with an exponential distribution, but the parameters of the distribution are dependent on the current alert level. The initial hypothesis was that the time between maneuvers would fall as the severity of the alert rose, but in fact no significant difference between update timing for corrective and warning alerts was observed. Updates did occur more slowly with a preventive alert or no alert at all. It was previously noted that the presence of recovery guidance led to the fastest response; this observation was incorporated into the model by considering the presence of recovery guidance to be an additional (the highest) alert level. As is evident in Figure 10, the distributions of the actual update times closer in form to a gamma distribution; however, for modeling convenience the exponential distribution was used for its memoryless property, which allows the delay to change as the encounter changes. Together with a minimum delay at the current alert level, the update delay distribution can be sampled at each simulation time step to determine whether to continue holding or to choose an update maneuver. Table 6 summarizes the modeled update delays for each alert condition, and Figure 10 compares the observed time between maneuvers with a corrective or warning alert with the model.

TABLE 6

Pilot Model Maneuver Timing

Delay Type	Model
Initial	$\Delta t_{init} \sim Exp(5)$
Coordination	$\Delta t_{coord} \sim \Gamma(5.5, 2)$
Execution	$\Delta t_{exec} \sim Exp(3)$
Update	$\Delta t_{upd} \sim \begin{cases} Exp(12) + 12 & \text{if no alert,} \\ Exp(6) + 9 & \text{if preventive alert,} \\ Exp(3) + 6 & \text{if corrective or warning alert,} \\ Exp(3) & \text{if recovery guidance} \end{cases}$

3.2 MANEUVER DIRECTION

The foundation of the model's maneuver direction choice is observed heuristics. Heuristics in Kahneman's sense [35] are cognitive shortcuts that are used, often subconsciously, to expedite decision-making. The term is used here in a slightly different sense: a heuristic is a modeling

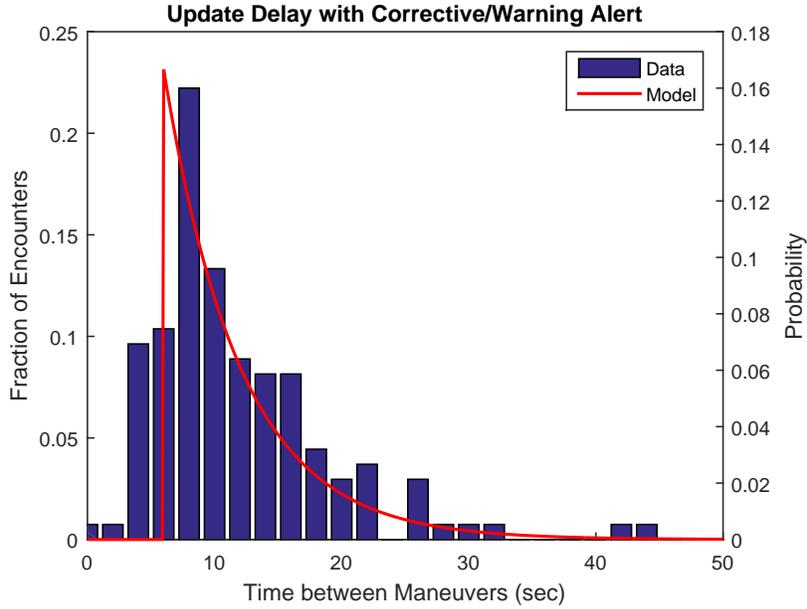


Figure 10. Observed and modeled update delay for corrective or warning alert (Experiment 2 data), shown here with the 3-second average execution delay included.

shortcut whereby an input is connected with an output without understanding the real world happenings in between. Specifically, the model uses binary choice heuristics—e.g., the choice of a left or a right turn under given conditions. The approach here differs from traditional heuristic approaches, however, in an important way. Traditionally heuristics define an outcome to certain conditions deterministically. For example, if encountering a stop sign at a road intersection, then stop. This is too idealized to truly represent the real world, in which drivers may not stop at a stop sign for any number of reasons. Acknowledging the unlikelihood of finding a heuristic that perfectly predicts pilot decisions in all their variety, the model instead establishes fuzzy-logic-like heuristics, or rules, that are not always followed. Each rule expresses a choice of action that is paired with a weight value (ranging from 0 to 1) identifying the likelihood it will be followed.

Symbolically, for the set of actions $\mathcal{A} = \{A, B\}$, conditions (or features) $F = \{f_1, f_2, \dots, f_n\}$, and rule weight $s \in [0, 1]$, the rule R states:

$$R(\mathcal{A}; F; s) = \begin{cases} A & \text{if } x \leq s \\ B & \text{if } x > s \end{cases} \quad (3)$$

for $x \sim F(x)$, where F is a probability distribution. A uniform distribution, $x \sim U(0, 1)$, is used for all rules in this model. This construct represents tendencies or preferences for particular responses, so an observation such as 80% of pilots choose a right turn under certain conditions and 20% choose to turn left can be represented with the rule to turn right with weight 0.8.

A rule with weight 0.5 is essentially a coin-flip. If the weight is less than 0.5, meaning the rule is more likely to be not followed than to be followed, the rule can be rewritten in a positive form as follows:

$$R(\mathcal{A}; F; s) = R'(\mathcal{A}, F; 1 - s) \quad (4)$$

where R' is the complement of R :

$$R'(\mathcal{A}; F; s) = \begin{cases} A & \text{if } x > s \\ B & \text{if } x \leq s \end{cases} \quad (5)$$

This quite simply says that, for example, a 20% preference to turn left is equivalent to an 80% preference to turn right. Finally, if $s = 1$, the traditional form of the heuristic is restored, where the conditions deterministically define the outcome.

The model assumes an overall decision space of up, down, right, or left, so $\mathcal{A} = \{\text{up, down, left, right}\}$ —it does not presently permit multi-axis maneuvers. This is consistent with past research and with this study's experiments, which showed pilots chose multi-axis maneuvers less than 20% of the time. The goal of the maneuver direction subfunction is to choose one of these maneuver directions. That decision is broken down into a succession of three binary choices: left or right ($\mathcal{A}_1 = \{\text{left, right}\}$), up or down ($\mathcal{A}_2 = \{\text{up, down}\}$), and then horizontal or vertical ($\mathcal{A}_3 = \{\text{horizontal, vertical}\}$). The conditions for each will be addressed separately.

Beginning with the first choice, Maki et al. previously introduced a heuristic, mentioned above, based on right-of-way rules for the choice between turning left or turning right that was initially considered for this model. However, when applied to the experimental data, that heuristic was only about 60% accurate in predicting turn direction. A more accurate heuristic, with about 80% predictive accuracy, was formulated based on the suggestive guidance. This heuristic, referred to as the minimum-guidance heuristic, predicts that the pilot will turn in the direction in which the guidance suggests a smaller turn and/or lower alert level. Thus, if the ownship is heading due north and the guidance suggests either a 45° turn to the right or a 10° turn to the left, the pilot will likely turn left. The same heuristic was employed for the choice to descend or climb as it has even better predictive accuracy of about 90%.

To use this minimum-guidance heuristic, the model ingests the suggestive guidance data that gives expected alert levels over a range of headings ($\pm 135^\circ$ as implemented) and altitude (+1500/−1000 ft). It pares this continuous data down to four *minimum suggestions*—the smallest maneuver in each direction that will achieve the lowest alert level. For the example suggestive guidance shown in Figure 11, the minimum suggestions are as follows:

- Left turn: 28° (heading 332) to no alert
- Right turn: 13° (heading 13) to no alert
- Climb: 0 feet (altitude 15,000 feet) to corrective alert

- Descend: 1000 feet (altitude 14,000 feet) to no alert

In this example, the pilot could choose to turn in either direction or descend to avoid a loss of well clear. The minimum-guidance heuristic suggests that if a turn is selected, it will probably be to the right. Because the intruder is descending from above and the ownship's vertical rate is limited (500 feet per minute), any climb will result in a loss of well clear; on the other hand, descending 1000 feet will avoid the loss of well clear.

One important note about this methodology: its present implementation is only valid when there is a single contiguous horizontal threat region spanning the current course. This precludes use of the model in encounters with more than one intruder. In those cases, the suggestive guidance will often include multiple threat regions.

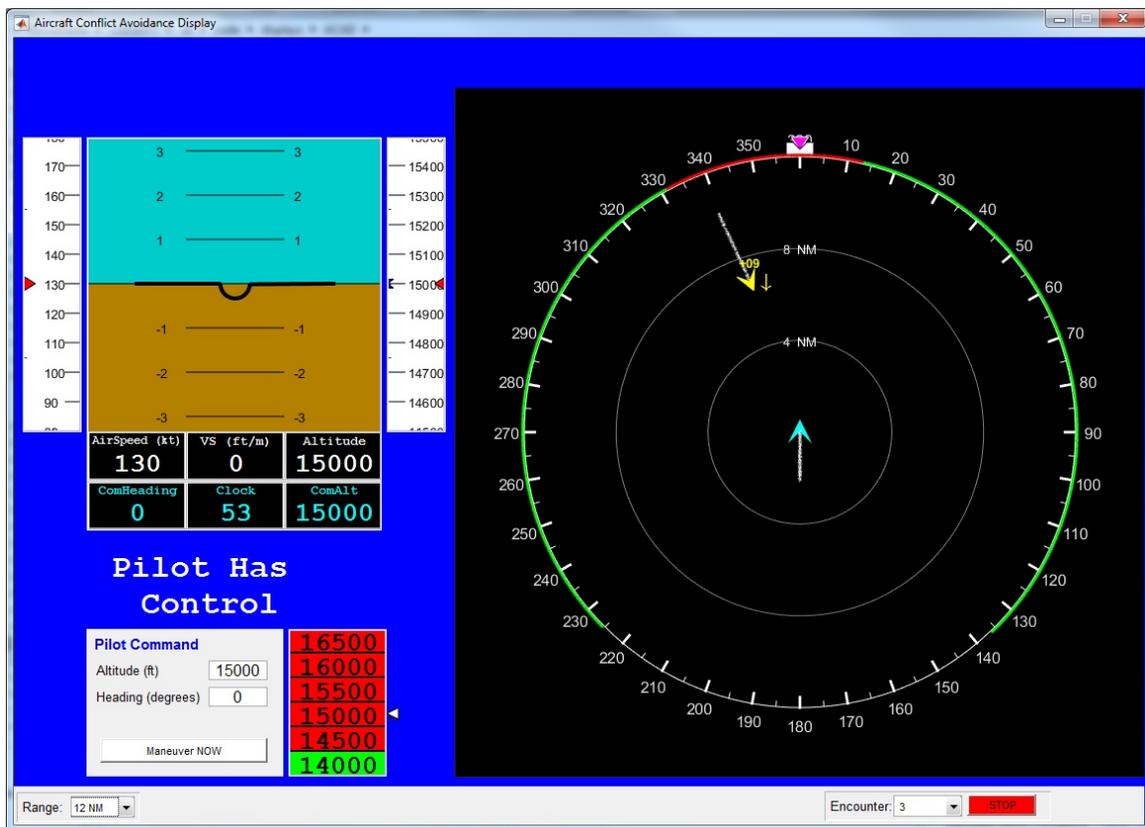


Figure 11. Example of suggestive guidance and associated minimum suggestions.

With the four minimum suggestions in hand, the model then uses a binary decision-making approach referred to as *pairwise elimination* to choose a maneuver direction. That is, it first chooses its preferred horizontal maneuver and then its preferred vertical maneuver, reducing the decision space from a set of four options to two. It then chooses whether to maneuver vertically or horizontally. Figure 12 illustrates this process.

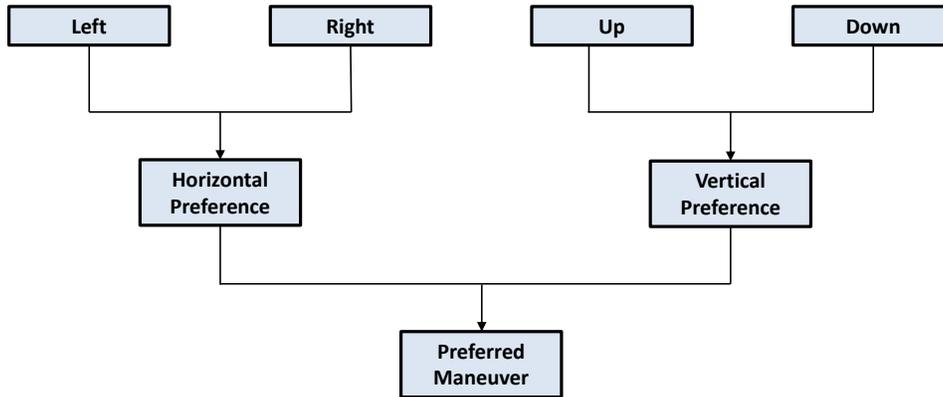


Figure 12. Pairwise elimination procedure for choice of maneuver direction.

Each step of the pairwise elimination uses rules that are defined below. The basic procedure is this:

1. If the two minimum suggestions have different predicted alert levels, choose the maneuver with the lower alert; if not, continue to the next step.
2. If the suggested maneuver magnitudes are not equal, draw $X \sim U(0, 1)$ and choose the smaller maneuver if $X \leq P$ or the larger maneuver if $X > P$, where P is a rule weight; otherwise, continue to the next step.
3. If the suggested maneuver magnitudes are the same, choose the maneuver corresponding to a random draw against preference Q , where Q is a rule weight.

That procedure can be codified as follows. Each action $\mathcal{A}_i \in \mathcal{A}$ has an associated minimum suggestion, as described above, that can be denoted as a tuple $\mathcal{S}_i = \{\Delta_i, a_i\}$, where Δ_i is the magnitude of the minimum maneuver suggestion and a_i is the corresponding predicted alert level. Binary decision-making involves a choice between two actions $\mathcal{A} = \{\mathcal{A}_i, \mathcal{A}_j\}$ and the two corresponding minimum suggestions $\mathcal{S} = \{\mathcal{S}_i, \mathcal{S}_j\}$. The pair of maneuver magnitudes is similarly denoted as $\Delta = \{\Delta_i, \Delta_j\}$, and the pair of alert levels as $a = \{a_i, a_j\}$.

Now the following choice relations can be established:

$$\mathcal{S}_i = \mathcal{S}_j \text{ if } (\Delta_i = \Delta_j \wedge a_i = a_j) \quad (6)$$

$$\mathcal{S}_i \neq \mathcal{S}_j \text{ if } (\Delta_i \neq \Delta_j \vee a_i \neq a_j) \quad (7)$$

$$\mathcal{S}_i < \mathcal{S}_j \text{ if } (\Delta_i < \Delta_j \wedge a_i = a_j) \vee a_i < a_j \quad (8)$$

If $\mathcal{S}_i \neq \mathcal{S}_j$, define the function

$$\operatorname{argmin} \mathcal{S} = \begin{cases} i & \text{if } \mathcal{S}_i < \mathcal{S}_j \\ j & \text{if } \mathcal{S}_i > \mathcal{S}_j \end{cases} \quad (9)$$

Using this notation, rules (conditions and associated weights) were derived from the experimental data. For horizontal pairwise elimination, yielding a single turn direction, the rules are summarized in Table 7.

The experimental data demonstrates that pilots generally prefer to turn in the direction of the smaller suggested turn—but they sometimes choose the direction of the larger suggested turn. The model assumes that the degree of preference is the same regardless of encounter conditions and the associated suggestions. Tests with an initial version of the model revealed that this assumption could lead to unreasonable responses. Choosing to turn in the direction of the larger suggestion seems reasonable when the suggested turn magnitudes are similar but less so with highly asymmetric suggestions. For example, if the guidance were to suggest a left turn of 70° and a right turn of 10° , it seems unlikely that a pilot would choose to turn left. Indeed, this conclusion is borne out in the data. To remedy this unreasonable behavior of the model, the conditions in rules 1.2a and 1.2b were added so that the direction of the minimum suggested turn is always selected if the difference in magnitudes is greater than 40° .

The rules for vertical pairwise elimination are summarized in Table 8. One deviation from the observed behavior has been incorporated. In the experiment, the minimum altitude of any encounter was 3000 feet. Anticipating the need to apply the model over a wider range of altitudes than strictly valid, including altitudes down to ground level, logic was incorporated that prevents the choice of a descent altitude below 500 feet, as dictated by rules 2.1a, 2.1b, 2.2a, and 2.2b.

The final pairwise elimination step is to choose between the preferred turn and the preferred vertical maneuver. This decision is based on the alert levels for the maneuvers in each plane and the observed overall preference for horizontal and vertical maneuvers. Table 9 summarizes the rules.

Anecdotal evidence suggested that the preference for horizontal or vertical maneuvers could vary by pilot population. While professional training may have a role shaping this preference, it was postulated that preferences are also dependent on the capabilities of the ownship aircraft. Experiment 2 explored this dependency by having the subjects use various ownship models. In the resulting data, preferences were identified for horizontal maneuvers over vertical maneuvers for four categories of ownship aircraft based on achievable turn rate ($\dot{\psi}_{max}$) and achievable vertical rate (\dot{h}_{max}). Table 10 summarizes those categories and the associated rule weights for turning over vertical maneuvers. These results largely conform to expectations, e.g., when the ownship is more agile horizontally than vertically, pilots prefer horizontal avoidance maneuvers. With an agile ownship in both maneuver planes, pilots preferred vertical maneuvers, a finding consistent with preferences reported in previous studies as described in Section 1.2.

For simplicity, the same rule weights were used in rules 1.2a and 2.2a even though the data indicated that pilots are more inclined to follow the minimum guidance in their vertical maneuvers

TABLE 7

Horizontal Pairwise Elimination Rules for $\mathcal{A}_1 = \text{Turn Left}$ and $\mathcal{A}_2 = \text{Turn Right}$

Rule	Under the conditions	Choose \mathcal{A}_i , where $i =$	With strength
1.1	$a_1 \neq a_2$	argmin \mathcal{S}	1.0
1.2a	$\mathcal{S}_1 \neq \mathcal{S}_2 \wedge \max \Delta - \min \Delta \leq 40^\circ$	argmin \mathcal{S}	0.8
1.2b	$\mathcal{S}_1 \neq \mathcal{S}_2 \wedge \max \Delta - \min \Delta > 40^\circ$	argmin \mathcal{S}	1.0
1.3	$\mathcal{S}_1 = \mathcal{S}_2$	1	0.6

TABLE 8

Vertical Pairwise Elimination Rules for $\mathcal{A}_1 = \text{Climb}$ and $\mathcal{A}_2 = \text{Descend}$ and Ownship
Altitude h in Feet

Rule	Under the conditions	Choose \mathcal{A}_i , where $i =$	With strength
2.1a	$a_1 \neq a_2 \wedge h - \Delta_2 \geq 500$	argmin \mathcal{S}	1.0
2.1b	$a_1 \neq a_2 \wedge h - \Delta_2 < 500$	1	1.0
2.2a	$\mathcal{S}_1 \neq \mathcal{S}_2 \wedge h - \Delta_2 \geq 500$	argmin \mathcal{S}	0.8
2.2b	$\mathcal{S}_1 \neq \mathcal{S}_2 \wedge h - \Delta_2 < 500$	1	1.0
2.3	$\mathcal{S}_1 = \mathcal{S}_2$	1	0.5

TABLE 9

Maneuver Plane Rules for $\mathcal{A}_1 \in \{\text{turn left, turn right}\}$ and $\mathcal{A}_2 \in \{\text{climb, descend}\}$

Rule	Under the conditions	Choose \mathcal{A}_i , where $i =$	With strength
3.1	$a_1 \neq a_2$	argmin a	1.0
3.2	$a_1 = a_2$	1	See Table 10

TABLE 10

Maneuver Plane Selection Rule Weights

	$\dot{\psi}_{\max} \leq 2^\circ/\text{s}$	$\dot{\psi}_{\max} > 2^\circ/\text{s}$
$\dot{h}_{\max} \leq 1000 \text{ fpm}$	0.5	0.65
$\dot{h}_{\max} > 1000 \text{ fpm}$	0.5	0.4

than for horizontal maneuvering. Furthermore, the model uses the same rules and weights for initial maneuver decisions and updates despite a much stronger preference for horizontal maneuvers in the latter case. Both simplifications could be readdressed in subsequent updates of the model.

It is important to reiterate that this decision-making process is not proposed to be that used by a pilot during decision-making; rather, the weighted rules here describe the outcome of decision-making. As noted previously, this approach does not explicitly model the decision-making process but rather the decisions themselves.

3.3 MANEUVER MAGNITUDE

The final subfunction of the model is selection of magnitude of the maneuver—the specific heading or altitude desired. Examination of the data collected in the HITL experiments revealed that the maneuver magnitudes selected were well-modeled by a gamma distribution relative to the minimum suggested maneuver in the direction selected. That is, over all of the recorded right turns, the magnitude of the turn was gamma-distributed with respect to the displayed suggested minimum right turn. Moreover, not only were the left turns also gamma-distributed with respect to the minimum suggested left turn, that distribution was approximately equal to the right turn distribution. The same trend was observed for vertical maneuvers and is consistent with a hypothesis that pilots would tend to use the maneuver guidance plus some discretionary margin when deciding upon a maneuver magnitude.

Because gamma distributions have domain $x \geq 0$, the distributions were shifted to allow selection of maneuvers smaller than the minimum suggested, a behavior that appeared in the experimental data. The suggestive guidance constantly changes as the encounter evolves, and those cases of the pilot commanding a maneuver smaller than the minimum suggested are attributed to the pilot’s attention being focused away from the suggestive guidance to entry and execution of the desired maneuver. The reference guidance for the pilot is some time in the past. It is difficult or impossible to identify the guidance state at the moment of decision-making, but since this model is not directly concerned with the actual cognitive processes, referencing the guidance state at the moment of the maneuver command is acceptable.

Further examination of the maneuver magnitudes in the experimental data revealed one other phenomenon that was incorporated into the model. The maneuver magnitude distributions differed

for horizontal maneuvers with direction consistent with the minimum-guidance heuristic (referred to subsequently as compliant) and those opposite (noncompliant). Compliant maneuvers were approximately twice as large as noncompliant maneuvers, suggesting that when choosing the direction requiring the larger maneuver the pilots tended to buffer their maneuvers less, perhaps in an attempt, conscious or unconscious, to comply with their assigned objective of minimizing deviations from course. This dependency on turn direction is reflected in the model with separate distributions. The experimental data and the models fit to those data for compliant and noncompliant turns are shown in Figures 13 and 14, respectively. These distributions give average turn magnitudes of 30.1° (compliant) and 15.1° (noncompliant) relative to the guidance.

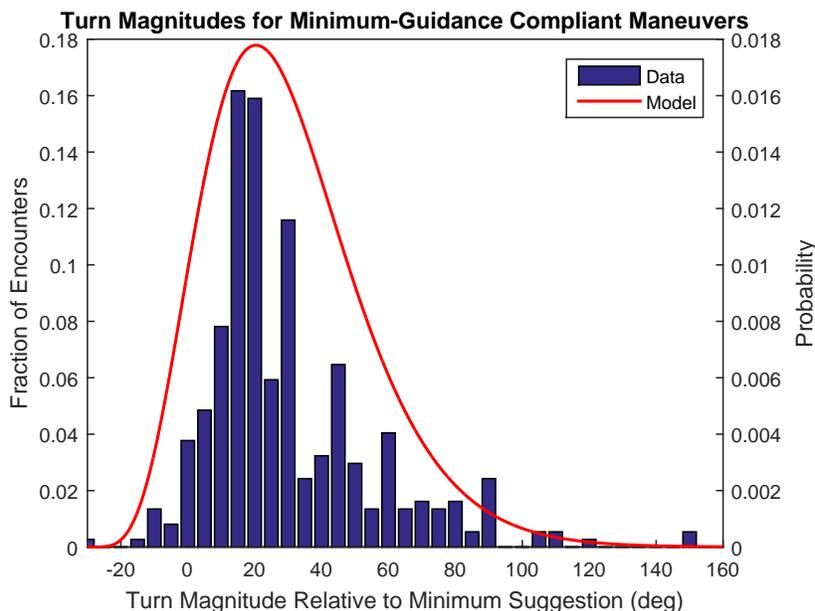


Figure 13. Observed and modeled turn magnitudes for minimum-guidance compliant maneuvers (Experiment 2 data).

This behavior was not found in the vertical dimension, so the model uses a single distribution for maneuver magnitudes in that dimension regardless of whether the heuristic was followed or not. Figure 15 shows the data and model for vertical maneuver magnitude, where the average altitude change relative to the guidance is 502 feet. Table 11 summarizes the gamma distribution parameters for each of the maneuver magnitude models.

The output of the model is either a desired heading or desired altitude. This is attained by drawing a heading change from one of these distributions and adding it to the current heading or altitude. The resulting desired heading is discretized to 5° and desired altitude is discretized to 100 feet—the smallest discretization observed in the data.

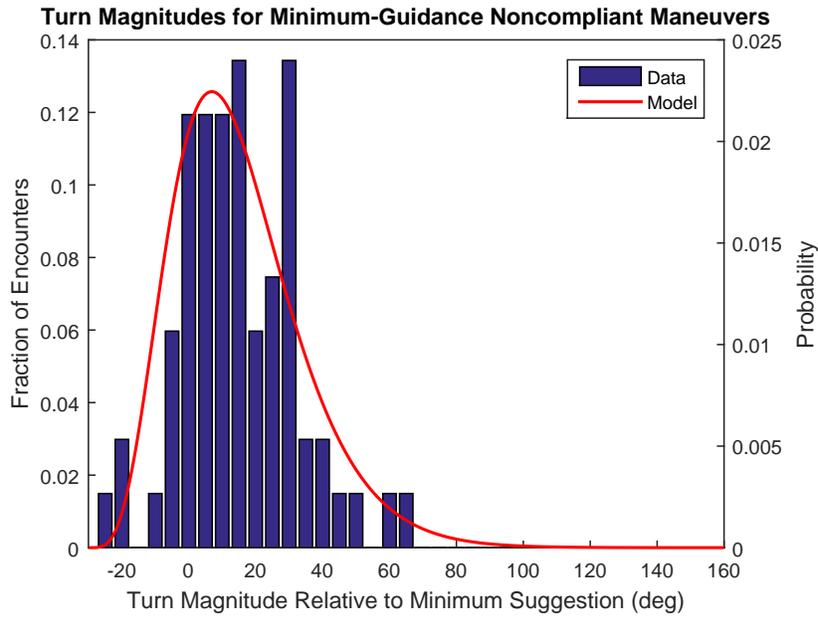


Figure 14. Observed and modeled turn magnitudes for maneuvers not following minimum guidance (Experiment 2 data).

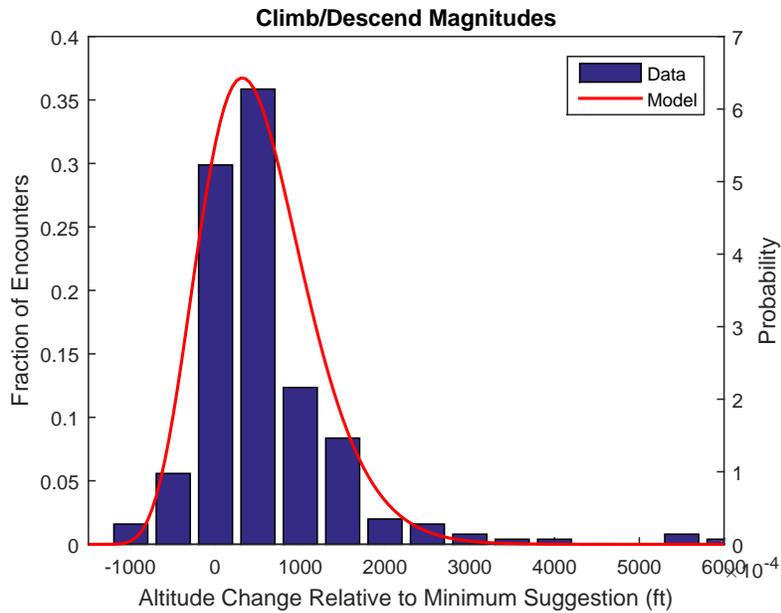


Figure 15. Observed and modeled climb/descend magnitudes for all vertical maneuvers (Experiment 2 data).

TABLE 11**Pilot Model Maneuver Magnitude Distributions**

Maneuver Type	Model
Horizontal (Compliant)	$\Delta\psi \sim \Gamma(6.21, 9.67) - 30^\circ$
Horizontal (Noncompliant)	$\Delta\psi \sim \Gamma(5.47, 8.25) - 30^\circ$
Vertical	$\Delta h \sim \Gamma(9.37, 207.98) - 1500 \text{ ft}$

A few undesirable behaviors (not consistent with the experimental data) emerged after the initial implementation of these maneuver magnitude distributions. Several overrides were added to tame those behaviors. First, an altitude floor of 500 feet was imposed to prevent the selection of clearly unrealistic desired altitudes. In addition, because the distributions are independent of the magnitude of the suggested guidance and the distributions include negative changes relative to that guidance, it is possible for the maneuver magnitude drawn to be inconsistent with the maneuver direction picked per the rules above. For example, a -20° maneuver magnitude paired with a suggested turn 15° right would actually yield a left turn. To prevent this, the model bounds the desired maneuver by the current heading or altitude, so for the previous example, the resulting maneuver would be a 0° “turn”—maintain current heading, in other words. Finally, because altitude changes often take a long time to complete, the update maneuver timing distributions above can lead to multiple changes to the desired altitude before the originally desired altitude would have been reached—a behavior not observed in the experiments. Consequently, vertical maneuver updates were limited to instances where the suggestive guidance indicates the original maneuver is no longer sufficient. Turns are usually completed much more quickly, and no modification to the basic structure was needed.

3.4 DETERMINISTIC MODE

The model also includes a deterministic mode that can be selected for analyses in which variability in the pilot response may obfuscate the effects of other subsystems under test. In this mode, each probability distribution is replaced with its mean—e.g., a 5-second initial delay in every instantiation of the model. All rule weights are set to unity. When run in this mode, the model will always output the same maneuver commands for given alerting, suggestive guidance, and ownship maneuverability input. If any of those inputs change, the maneuver output may also change.

3.5 ASSUMPTIONS AND LIMITATIONS

It is important to note that the pilot model described here models pilot behavior only under certain circumstances: when facing a potential loss of well clear with a single intruder aircraft and flying with a DAA system in the context of the Phase 1 DAA MOPS. The key assumptions and

limitations of the model that must be considered when using it in a safety analysis are summarized here.

- The unmanned aircraft is flown by keyboard: that is, to perform a maneuver, the pilot enters a desired altitude and/or a desired heading at his or her workstation and an on-aircraft control system actually executes the maneuver rather than directly controlling the aircraft with a stick and rudder. Operational systems provide examples of both modes of control: the ground control station for the Predator includes the standard manual flight controls while the Global Hawk's ground control station does not. In [25], Thomas and Wickens noted that the mode of control influenced manned aircraft pilots' maneuver decisions. The same may be true for pilots of unmanned aircraft. Further work would be necessary to identify any effect on the model.
- The model is based on suggestive guidance specifically from the DAIDALUS algorithm. However, the similarity of the results from Experiment 1, which used the OmniBands algorithm for suggestive guidance, and Experiment 2, using DAIDALUS, suggests that the model may be broadly applicable to any band-style suggestive guidance algorithm. It should be used more cautiously with other suggestive guidance approaches such as shaded keep-out regions similar to terrain and weather displays.
- Avoidance maneuvering commences only after a DAA alert. As previously noted, during the experiments, pilots often recognized an impending conflict and maneuvered before an alert if permitted. In this regard, the model is more conservative than the observed responses.
- All aspects of the maneuver decisions are correlated to the encounter geometry only through the suggestive guidance. Further, all rule weights are fixed and uncorrelated to the encounter geometry.
- Update maneuvers are uncorrelated from any previous maneuvers. In fact, the experiments showed that an initial vertical maneuver was frequently followed by a horizontal maneuver, which is not currently reflected in the model.
- Each maneuver decision is for a single axis maneuver only. Multi-axis maneuvers were selected by pilots 5–20% of the time in the experiments, a small but not negligible number. The model approximates multi-axis maneuvers by independently choosing the maneuver plane in each successive maneuver decision.
- Suggestive guidance data has been smoothed. Noisy surveillance sources can cause considerable jitter in the suggestive guidance. The pilot model uses the guidance only at the current time step, which may yield unrealistic maneuvers in the presence of jitter. Further work is necessary to model interpretation of noisy data.
- Maneuver decisions are based on the instantaneous own and intruder states (or state estimates) at the moment of maneuver selection and linear projections of future states. The DAA MOPS requires that the unmanned aircraft be flying an IFR flight plan, which means

its route is preplanned. Any intended maneuvers of the ownship to follow that route are not considered by the model during avoidance maneuvering.

- Maneuvers are always coordinated with ATC unless a warning alert occurs. NASA's experiments suggest that this is a conservative assumption, as in those experiments pilots coordinated in advance only about 70% of the time [41].
- A hard floor of 500 feet altitude is imposed that the model will not descend below. In cases when the ownship is already below 500 feet altitude, the model will not permit the ownship to descend to avoid traffic. The experiments did not explore encounters at very low altitudes (the lowest was 3000 feet), and this behavior is merely asserted.
- Return-to-course decisions are not modeled, nor does the model provide any indication that the conflict has been cleared. Returning to course prematurely was a cause of some of the losses of well clear in NASA's experiments [29]; in this regard, the model may be considered slightly optimistic.
- Single intruder encounters. The logic can only interpret suggestive guidance with a single contiguous band of no-fly headings and altitudes, which frequently is not the case in multi-intruder encounters. Encounters with multiple intruders are not very likely in the context of the Phase 1 MOPS, which excludes terminal area operations.

The model's parameter values, as summarized in Appendix B, are strictly valid only under these assumptions. However, one of the benefits of this modeling approach is that the model parameters themselves are meaningful, something not always true for more complex approaches like neural networks or hidden Markov models. With careful manipulation of the parameters, the model can be used to explore the effects of changing pilot behavior. For example, one could study the ramifications of training unmanned aircraft pilots to maneuver primarily in the vertical plane by adjusting the associated rule weight.

4. VALIDATION

This model of pilot traffic avoidance maneuver choices has been constructed by breaking those choices down into subdecisions and extracting data from human-in-the-loop experiments supporting those subdecisions. The question that remains is whether the behavior that emerges when those pieces are put back together is anything like the behavior that was originally observed. More broadly, the model must behave in a way that is representative of pilot behavior in the real world to be considered valid.

Validating a model like this one is challenging for two main reasons. First, there is no “right” answer when maneuvering to avoid nearby traffic. There are an infinite number of paths in time and space that can be followed to remain well clear of another aircraft. Comparison of any one path generated with the model and any one path observed in an experiment, as in Figure 16, can lend credence to the model but does not say anything about the model’s overall validity. Instead, the distributions of paths selected must be compared. But generating an adequately refined *distribution* of observed paths is impractical—that’s why the model is needed in the first place.

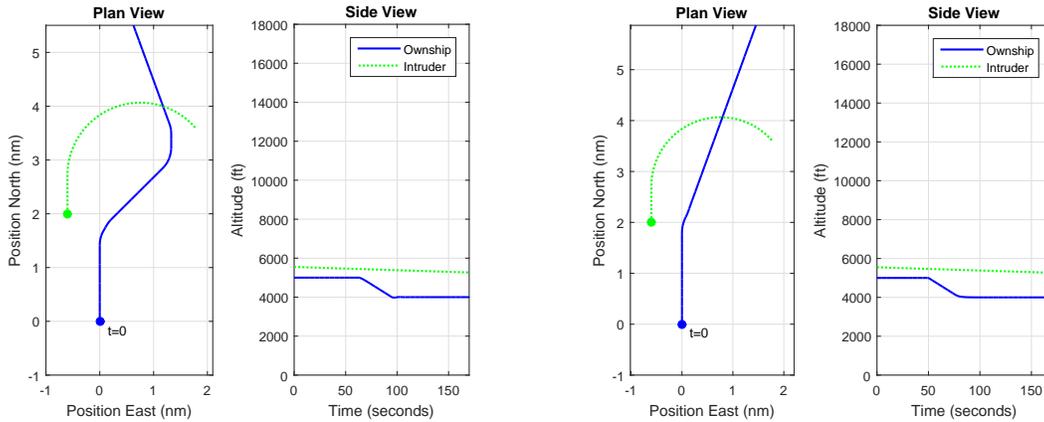


Figure 16. One response observed in Experiment 2 (left) and a simulated response (right).

One approach used to validate the model given this challenge is to compare its performance with observed performance at a higher level. Rather than comparing trajectories, encounter outcomes are compared: frequency of loss of well clear and frequency of near midair collision. Validation typically uses a different dataset than was used to build the model to avoid overfitting. Since as much data as could practically be collected was required to build the model, by necessity the same dataset had to be used for validation; however, a measure of independence was achieved by validating with features in the data not used in building the model: loss of well clear and near midair collision outcomes. Using a Simulink implementation of the model and the MIT Lincoln Laboratory Collision Avoidance and System Safety Assessment Tool (CASSATT) simulation software [42], 1000 trials were executed with the model for each of the 18 encounters from Experiment 2. Normalizing this dataset to the observed results, there were 34.24 losses of well clear (LoWC)

and 0.14 near midair collisions (NMAC) with the model⁶. This compares favorably with the 38 losses of well clear and 0 NMACs observed in Experiment 2.

This assessment can be taken a step further by looking at the results for each encounter. Losses of well clear with the model are compared with the 95% confidence bounds of the observed results in Figure 17. The model results are within the confidence bounds for most of the encounters. In two encounters (9 and 12), the model outperforms the pilots; in one encounter (16), the pilots were better than the model. The nominal trajectories for each of those encounters are shown in Figures 18–20.

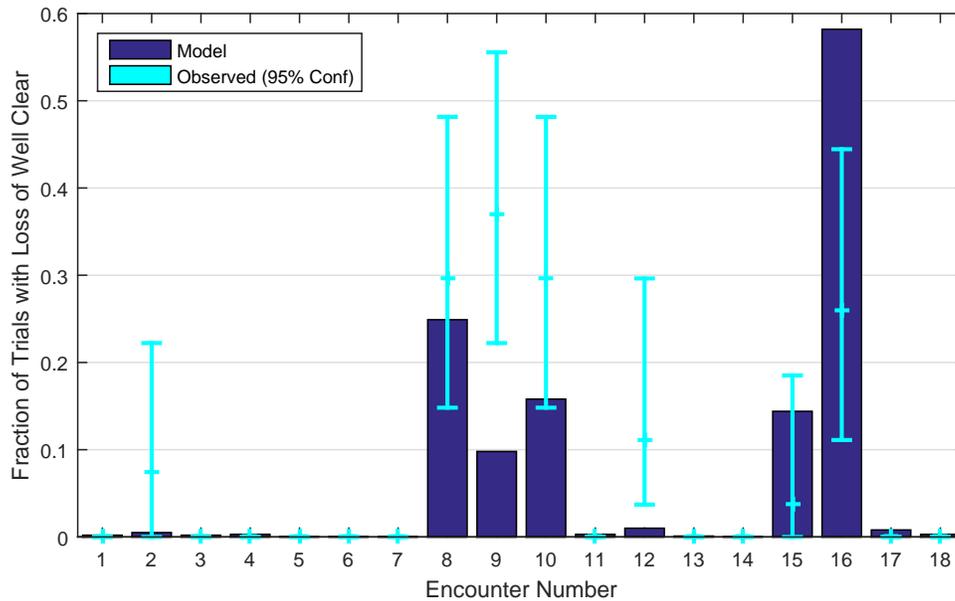


Figure 17. Comparison of observed and simulated losses of well clear.

The circling intruder in encounter 16 causes two periods of alerts, one early in the encounter and the second late and nearly coincident with loss of well clear, that reveals a model shortcoming. If the model maneuvers within the first alert period, approximately 10 seconds long, it is fairly successful: of the 1000 trials this occurs in 340, and a subsequent loss of occurs in only 17% of those—within the observed confidence bounds. On the other hand, if an initial maneuver does not occur within that first alert period, the pilot model is disabled until the second alert period, and in that case, 80% of the trials result in a loss of well clear. Pilots, on the other hand, were permitted to execute a maneuver chosen during the first alert period even after the alert ended. This model shortcoming could potentially be addressed in a future version, and in any case, it makes the model conservative.

⁶ The criteria for loss of well clear were defined earlier. A near midair collision is said to have occurred when the two aircraft approach within 500 feet horizontally and 100 feet vertically of one another [17].

On the other hand, both encounter 9 and 12 have only a single alert period with the intruder turning towards the ownship. In these cases, the model outperforms the pilots—substantially in encounter 9, the encounter most challenging for the pilots. In both of these cases, following the minimum initial guidance prompts a disadvantageous right-hand turn. What appears to mark the difference between the pilot responses and the model responses, though, is that the pilots are slower to command an initial maneuver. In both cases, the initial corrective alert quickly progresses to a warning alert. In these cases, the pilot model always skips ATC coordination. It is possible that some pilots waited for ATC approval before maneuvering or found these encounters mentally challenging. It is worth noting that the intruder in encounter 12 is descending as well as turning, and it was previously shown that pilots were slower to react to climbing or descending intruders.

Though this assessment has revealed some instances in which the model does not perfectly mimic actual pilot behavior, the model does generally behave in a manner consistent with pilot behavior. Further refinement of the model may improve specific cases of model mismatch.

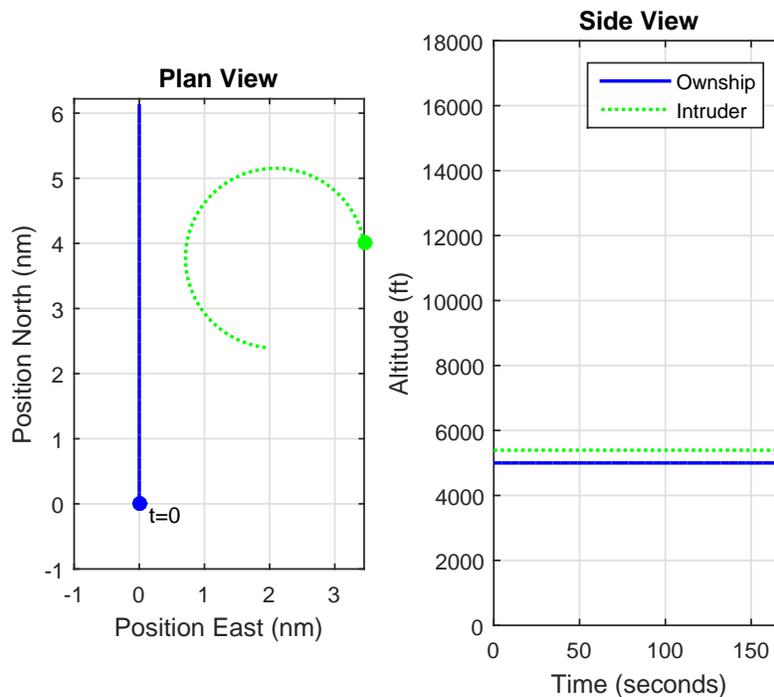


Figure 18. Experiment 2 encounter #9 trajectories.

Another validation activity that would be valuable but has not been completed is a qualitative validation approach based on the Turing test [43], wherein a human views a response and is asked to judge whether that response comes from another human or a machine. A new set of UAS pilot subjects would view a set of encounters, some of which are as-flown by a human pilot (reused from the HITL data collection experiments, perhaps) and others artificially generated using the pilot model. The subjects would be asked to determine which were from the HITL and which from the

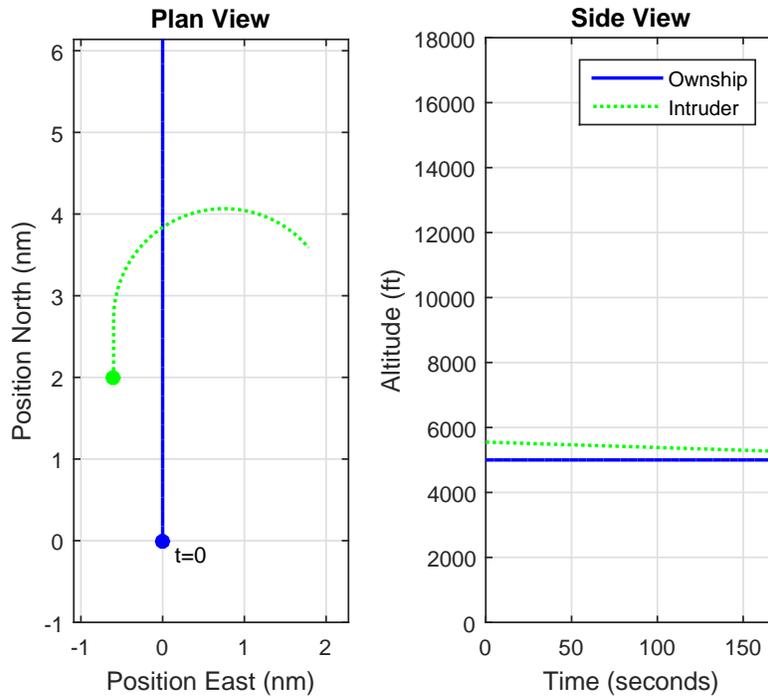


Figure 19. Experiment 2 encounter #12 trajectories.

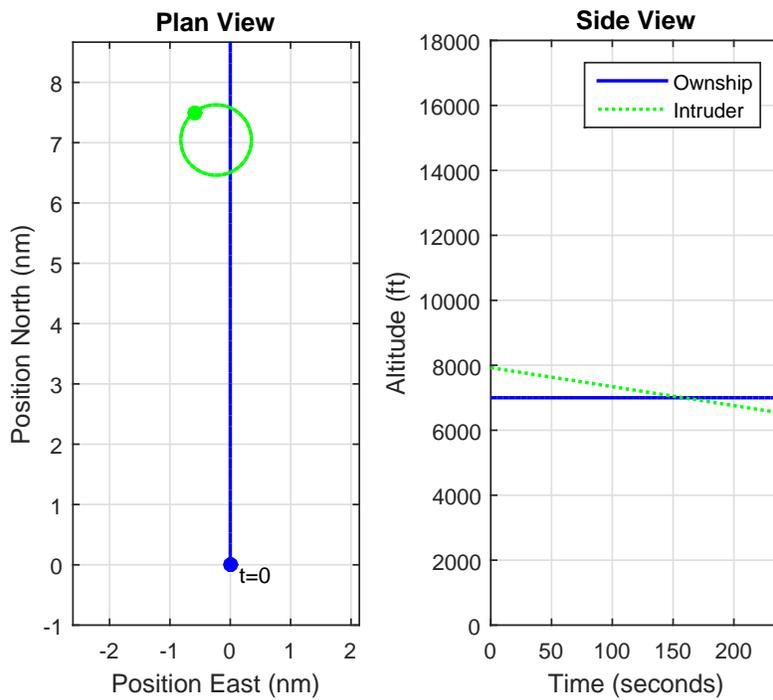


Figure 20. Experiment 2 encounter #16 trajectories.

model. If the subjects are unable to distinguish between the two types of encounter, the model may be considered valid by this test.

The second validation challenge is that the task being modeled is something that pilots do not do yet—it is an envisioned world problem [44]. Responses of pilots asked to use a prototype DAA system to fly an unmanned aircraft in a NAS environment may not be the same as the responses once a system is actually operational. To mitigate this issue, each pilot was extensively trained on use of the DAA system prior to participating in the experiments—about a third of the total time taken with each subject was spent on training. However, it is inevitable that as these systems are deployed preferences will evolve, and the validity of the model will need to be reexamined periodically.

This page intentionally left blank.

5. SAFETY ANALYSIS EXAMPLE

Validation showed that the model can generate a response similar to a known human response. The primary utility of the pilot model is to generate responses in large sets of encounters for which human response are unknown so as to assess the effectiveness of systems and/or procedures for safe flight. What follows here is a simple example of a safety analysis using this model.

The first step of the analysis is to generate encounters representative of the relevant airspace. This was done using the MIT Lincoln Laboratory Correlated Encounter Model for encounters in the NAS [45]. The encounter model produces representative encounters between two aircraft, at least one of which is receiving separation services from ATC. Rejection sampling was used to build up a set of one million encounters in which the trajectory of one of the aircraft was consistent with the following flight dynamics assumptions:

- Maximum turn rate: $3^\circ/\text{s}$
- Maximum rate of climb/rate of descent: 3000 fpm

In addition, the sampled initial conditions were propagated backwards to begin approximately 110 seconds from the nominal (unmitigated) closest point of approach so that the resulting trajectories begin prior to any alert.

Each encounter was simulated in CASSATT. No pilot interaction or any other safety system was used for this initial run. From the resulting trajectories, any instances of NMAC or LoWC were identified to establish a baseline, *unmitigated* probability of NMAC or LoWC given an encounter, $P(\text{NMAC}|\text{encounter})$ and $P(\text{LoWC}|\text{encounter})$.

Next, each of the encounters was simulated in CASSATT again, this time with the DAA system and pilot model in place. This example used a perfect surveillance system (no measurement error and unlimited detection range), DAIDALUS for alerts and suggestive guidance, and the pilot model in stochastic mode. Again, after simulating each of the encounters, instances of NMAC and LoWC were identified to get the *mitigated* probabilities of NMAC and LoWC.

Safety metrics were then computed from the two data sets, here the *risk ratios*. The risk ratio is simply the ratio of the probability of an event in the mitigated case to the probability of an event in the unmitigated case. NMAC and LoWC risk ratios are defined in Equations 10 and 11, respectively.

$$RR_{NMAC} = \frac{P_{mitigated}(\text{NMAC}|\text{encounter})}{P_{unmitigated}(\text{NMAC}|\text{encounter})} \quad (10)$$

$$RR_{LoWC} = \frac{P_{mitigated}(\text{LoWC}|\text{encounter})}{P_{unmitigated}(\text{LoWC}|\text{encounter})} \quad (11)$$

Use of the ratio eliminates the unknown probability of encounter. In this example, the NMAC risk ratio was 0.039 and the LoWC risk ratio was 0.132. These results make it evident that under these

circumstances, the DAA system and pilot are able to substantially reduce the number of NMACs (by about 95%) and LoWCs (by about 85%) that would otherwise have occurred. Of the NMACs that occurred with the pilot model, about 44% were *induced*: the pilot maneuvers caused NMACs to occur in encounters that would otherwise have been free of NMAC. Absent knowledge of the intruder aircraft's future maneuvers, the pilot may choose avoidance maneuvers that are consistent with the suggestive guidance but that nonetheless lead to conflict.

6. CONCLUSION

This product of this work is an empirical, rule-based, stochastic model of the maneuvers UAS pilots choose to remain well clear of other air traffic. Multiple human-in-the-loop experiments were conducted to explore UAS pilots' traffic avoidance maneuver decisions when using a DAA display showing traffic and maneuver recommendations, and the data collected in those experiments formed the empirical foundation of the model. The model is intended for use in studying the safety of unmanned aircraft flying in the NAS in the context of the RTCA Phase 1 DAA MOPS, where safe operation includes not only avoiding collisions with other aircraft but also retaining enough separation from other aircraft so as to be well clear.

Several generalizations can be drawn from the data collected in the experiments and are reflected in the model. First, contrary to the preferences of manned aircraft pilots and air traffic controllers documented in other studies, unmanned pilots often prefer to maneuver in the horizontal plane. As a caveat, this preference weakens when the unmanned aircraft is highly maneuverable in the vertical dimension or when the pilot is permitted to maneuver without an alert. Single-axis maneuvers heavily outnumbered multi-axis maneuvers, behavior possibly driven by the expedience of choosing and executing a single-axis maneuver. Within each maneuver plane, the preferred maneuver direction is often consistent with the direction of the smaller maneuver suggested by the guidance algorithm. It is not certain whether this is because the pilots are relying on the guidance, either consciously or unconsciously, or because the guidance mirrors the mental processes of the pilot. This behavior is worth more study, especially if this model were to be extended to situations without displayed maneuver guidance.

There were several relevant findings relating to the timing of avoidance maneuvers. First, if permitted to maneuver without an active alert, pilots frequently maneuvered well in advance of the earliest alert; consequently, the model's assumption that maneuvering does not begin until an alert makes it conservative for safety evaluation. In addition, these experiments demonstrated that the intended effect of recovery guidance, prompt maneuvering, was realized—pilots indeed responded very quickly after the appearance of recovery guidance. Finally, there was evidence that climbing or descending intruders caused increased cognitive load for the pilots and led to longer response times. This may be a result of the overhead perspective of the display that does not explicitly show the intruder's vertical trajectory. This final feature was not incorporated into the model but may be considered for a future version.

Based on characteristics of the data collected, the modeling approach selected first involved identification of various timing elements, the dependencies of those elements, and their distributions. Four timing elements emerged: an initial delay, an ATC coordination delay, an execution delay, and an update delay. These constitute a cycle that permits the model to select successive maneuvers, as pilots have been observed to do, until the conflict has been resolved. After establishing the timing framework, the approach chosen to model the maneuver direction choices was to derive rules from the data. The rules used by the model include a stochastic element: for each rule, an associated weight was identified reflecting the frequency it was followed by the pilots in the data collection experiments. Finally, analysis of the experimental data revealed that maneuver magnitudes, when considered relative to the suggestive guidance, are well represented by gamma distributions.

Using that model and the Lincoln Laboratory CASSATT aircraft encounter simulation, it has been demonstrated, as far as possible, that the model produces responses that are reasonably representative of observed data. Further, the stochastic nature of the model is consistent with the reality that there are many effective solutions when deciding how to avoid a nearby aircraft and that every pilot will likely choose a different one. Finally, application of the model to a simple safety analysis has been demonstrated, showing that with a DAA system using ideal surveillance (no errors and no detection constraints), pilots are able to reduce instances of loss of well clear by about 85% and near-collisions by about 95%.

The model may be used as is for safety analyses of DAA systems and can be readily tuned to explore minor changes in pilot behavior and operating procedures. Additional work remains to address the limitations detailed above. Of particular note, further exploration of the dependency of the pilot's maneuver choices on the time history of the displayed suggestive guidance is needed; the model's present use of only the current guidance may limit its validity in applications with large surveillance errors. Other useful improvements include modeling a pilot's decisions about when and how to return to his or her original course and modeling how intended future course changes influence traffic avoidance decisions. Additional validation work is also warranted, as well; in particular, a Turing-like test is recommended to explore whether pilots find the model-generated responses to be credible.

APPENDIX A: PILOT FEEDBACK AND HUMAN FACTORS CONCLUSIONS

After subjects completed Experiments 1 and 2, they were given the opportunity to provide feedback on their experience via survey. This appendix summarizes the data collected from those surveys and discusses conclusions drawn relating to human factors.

A.1 EXPERIMENT 1: NON-MANEUVERING INTRUDERS

Experiment 1 concluded with a survey intended to collect the participants' impressions of the experiment and the displays. Most significantly, we were interested in determining whether any of the participants had difficulties during the experiment that might be grounds for discarding that participant's data. In fact, all participants responded that they felt able to effectively complete the experiment, and no data was discarded.

One question explored the factors that contributed to the choice of a horizontal maneuver. Pilots were given a list of factors and asked to identify the most relevant. The results are shown in Figure A.1. An interesting observation of these responses is that the suggestive guidance ("fly/no fly bands") ranked near the bottom.

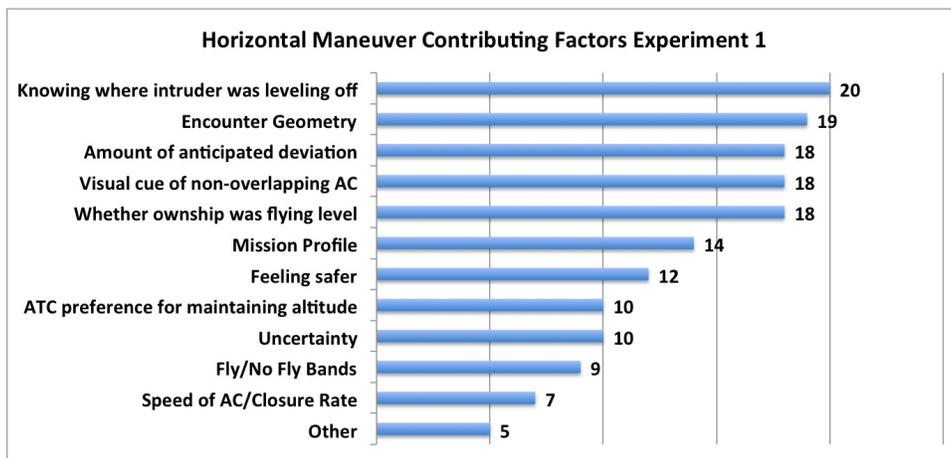


Figure A.1. Factors contributing to choice of horizontal maneuver in Experiment 1.

Pilots were invited to add comments further clarifying their responses. Verbatim comments included the following:

- Different situations are dictated by different aspects, such as, have I gathered intel to see if we are flying in a place with enemy aircraft or am I flying in a place where there are crop dusters or new pilots flying their Cessna. Also, if I am a slower aircraft and the intruder is going faster, I would make sure that I have more of a buffer zone, using horizontal deviation.

As my go to would be to deconflict by altitude well in advance, if they have more speed or are maneuvering to me, I would definitely be more apt to choose to horizontally maneuver.

- At that altitude, terrain avoidance would be a factor as well. Other traffic in the area would also be a factor if they are within +/- 5000 ft altitude.
- If I can easily lag him (conflict)[,] I will check turn to avoid the range requirement, that way I don't have to get permission to climb/descend as it is almost always harder to get climbs/descents rather than turns.
- Ownship performance (climb) and uncertainty of strange traffic, intention and performance (climb, turn).
- If already climbing/descending, considered level off for avoidance first.
- Intruder's track.
- Ownship vs. intruder speeds are significant factor (and geometry) on whether I will turn horiz. (my primary pref based on ATC Pref to maintain alt) or if I will climb/desc. Sometimes it's better (if mission allows) to maneuver even before bands become a factor, to alleviate uncertainty of intruder's intentions. Usually level state "encourages" horizontal maneuver as my 1st option (in most geometries).
- When a 30 degree or less heading change could cause my aircraft to pass behind the intruder, that would be my preference. I believe ATC would usually prefer a horizontal maneuver without altitude change as well. However, on several simulation runs the intruder changed heading to point at me, so this altered my choices[,] and I chose to use altitude split alone or in combination with horizontal offset.
- Mission requirements, i.e., priority for staying on track vs.[.] on altitude (vs. on speed, if it could change).
- ATC and mission profile under normal situations would play a major factor.
- Quickest way of deconfliction, often turning towards the tail of merging traffic.
- For the most part, I picked a heading that looked good to me in the green band. Sometimes it was 20 degree from the yellow, sometimes closer.
- Airspeed and descent rate deltas were the biggest influence on avoidance maneuver. Calculated closest point of approach.

The factors contributing to a vertical maneuver were explored as well. Figure A.2 summarizes those responses. Again, the suggestive guidance was ranked low. In this instance, pilots provided some specific remarks clarifying that ranking. Primarily, the pilots reported that the vertical guidance's position on the screen made it hard to pay attention to.

The verbatim free responses are as follows:

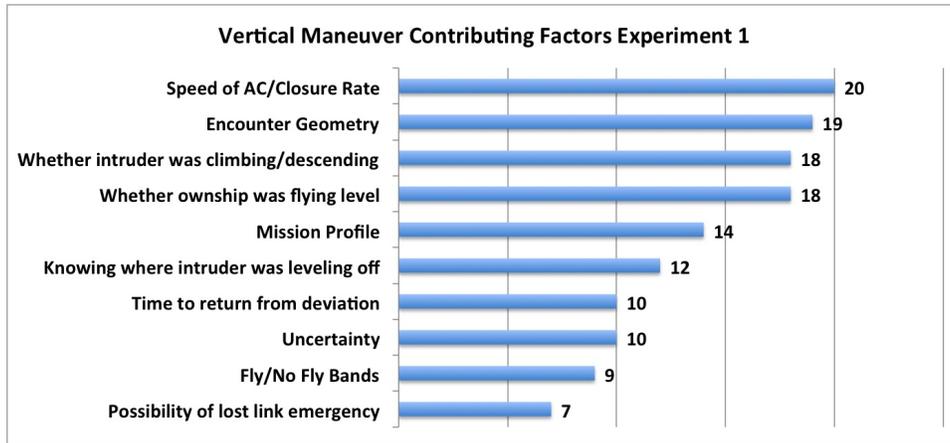


Figure A.2. Factors contributing to choice of vertical maneuver in Experiment 1.

- Didn't know where climbing/descending to[;] if ATC knows, easier.
- If I had a long time to deconflict and get approval with ATC, I would choose to deconflict by altitude first before deviating horizontally. Since Global Hawks fly to high altitude, I would rather do a climb and stay on mission profile, but I would want to make sure that the intruder is staying on heading/altitude and if not I would add additional layer of safety and deconflict horizontally.
- If climbing/descending into a flight path, I would hold or reverse altitude to avoid crossing intruder's or them crossing [me], especially if horizontal paths were converging too fast.
- Harder to get climbs descents but it is easiest geometry to solve. The biggest problem is uncertainty of targets a/c maneuver at end game.
- (Possibility of lost link emergency) Didn't consider this during simulation because it wasn't included. This would influence, but would depend on lost link logic of that system.
- Climbing was always preferred to descending[,] especially with low altitudes shown. Other factors that would have affected decision making would be reasons for being a certain altitude (wx, terrain, msn requirements).
- Overall [it was] easier to climb/descend for avoidance and minimize the time to get from pt A to B on mission profile.
- Lack of knowledge of where intruder was leveling off.
- Don't always know if intruder (climbing/descending) will level off[,] so this factors into Vert MNVR plan (and the need to possibly add a horizontal MNVR component as well if miss distance gets too close for comfort). EMER MSN (Lost Link) plan is small factor because typically encounter is a short event, and it's (dep. on the type of UAV) simple to adjust the lost link waypoints to follow the deviation. Vertical MNVR usually my 2nd option in close

(in the yellow) to increase miss dis, but my 1st option for longer range deconfliction if allowed (in the green distances).

- If my aircraft was climbing or descending and a level-off could provide altitude separation, that would be my preference. As noted previously, I tended to choose vertical offsets after several runs where the intruder turned toward me; in none of the scenarios did the intruders appear to start a climb or descent partway through the run[,] so altitude off set [sic] seemed more certain. In general, the Fly/No Fly Bands didn't offer any useful information (for horizontal or vertical maneuvers) that couldn't be determined from the informational displays. The real question in my mind as I tried to determine an appropriate response was what the intruder was going to do; if they were going to turn or climb/descend.
- Mission requirements. See comment for Q2 – Mission requirements, i.e., priority for staying on track vs[,] on altitude (vs. on speed, if it could change).
- I didn't feel like I noticed the vertical band when I was making a decision to maneuver vertically.
- I tried to maneuver vertically as a first plan because I didn't want [to] deviate from flight plan route, this may be weird assumption but it's how I would try to prioritize deviations when able.

Another question asked whether the pilots found any information unnecessary, confusing, or missing. The verbatim responses are as follows:

- Altitude bands were not utilized due to location on screen. Focus to heading bands is much easier.
- Target arrow up or arrow down < 500 fpm, trend arrow for target, color code this (trend)—maybe. Targets > 10K different altitude (unnecessary). When Target tags overlap (confusing).
- I thought this system was easy to use and was easy to learn.
- It seemed pretty clear once getting used to it after a few set ups [sic]. Things began to be easier to understand with practical use.
- I would like a flight path predictor and a temporal cue for yellow or red traffic for ownship so maybe a 30 s bubble where they could reach in 30 s the NM rings are nice, but temporal cues help when closure rates vary. Also I thought altitudes were unrealistically low for this UAV class. Lost link planning should be included in some way.
- More on instructions or displays. Would have like to see a displayed other aircraft heading displayed. I also would think Global Hawk pilots would almost never want to descend when already so low.

- Waiting for an alert was often unnecessary[,] especially when dealing with co-altitude aircraft. Early deviations made avoidance easier and allowed a greater buffer of safety for further maneuvers if non-participating [sic] a/c deviated from fight path/altitude.
- Flight level of conflict aircraft didn't seem to be at 100 ft increment (i.e., display FL029, actual 2950); actual altitude would be better (remove FL) background map on DAA was unnecessary.
- Extended vector line on ownship needed to visualize if current MNVR is "enough."
- Closest point of approach mark/line would help visualize intruders locations predicted @ passing ... whether I need to update my MNVR or not.
- All cues/info were useful. I did not use the altitude bands for controlling, but only as feedback, so it wasn't as useful or necessary. The heading bands were useful and I relied on them to minimize changes to aircraft track. With modified TAU in the horizontal plane only, this was simple. If there was a modified TAU-type of criteria for the vertical as well (altitude)[,] this would get much more difficult. As it was, the GUI was intuitive, easy to use[,] and I had plenty of space mental capacity to do other things or handle multiple conflicts.
- Not really for what information was desired.
- No[t] all info and displays were logical and informative.
- I think pilots need a bit of training (obviously) and a procedure for "when you see ___, do___". This made me start to treat this exercise more as science than art. I would tend to forget what an intruder alt had been then making me unsure if it was climbing or descending or up or down, an up arrow or down arrow may be helpful to show trend info. Any way the computer could give me closure speed may make it easier to do pilot math and ensure I can out climb or descent prior to converging. Good luck with your system development. Keep up the good work. We need this capability!
- Sometimes aircraft information was obscured by the range rings on the map. An altitude view/profile view could give pilots better situational awareness with the lateral/overhead view.
- Miss distance criteria not standard. Instructions as to priorities, minimal deviation of track vs. miss distance.
- Depiction of flight path and altitude prediction, based on current headings and VVI would be useful. All my altitude corrections were basically guesses.

A final question explored whether have to wait for an alert before maneuvering caused a conscious change in strategy. Responses to this question, as shown in Figure A.3, were sharply divided.

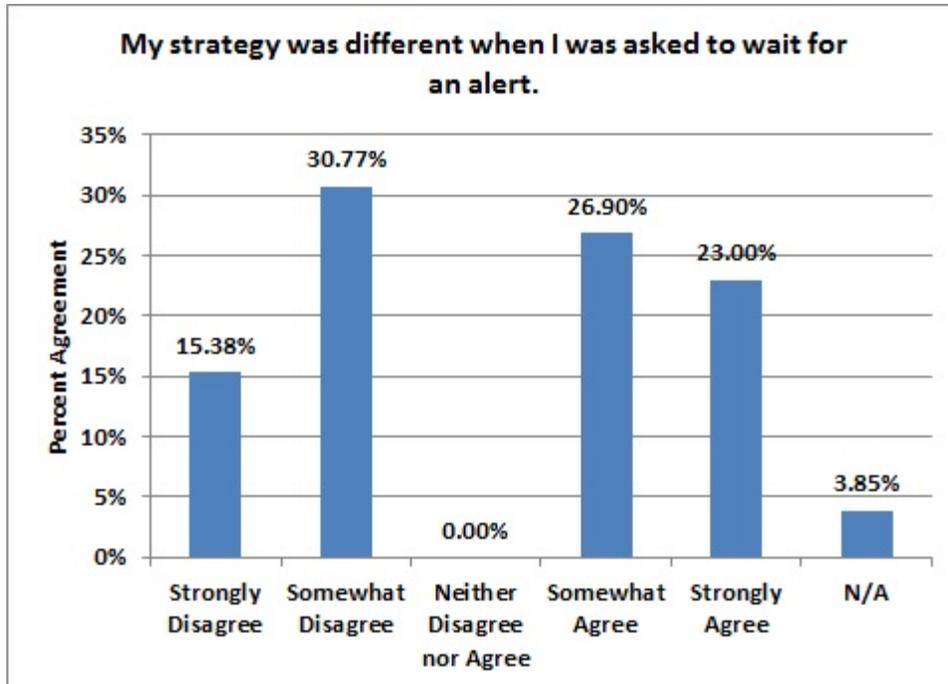


Figure A.3. Influence of maneuver protocol on maneuver strategy.

A.2 EXPERIMENT 2: MANEUVERING INTRUDERS

As in Experiment 1, Experiment 2 concluded with a survey to collect the participants thoughts on the experiment. Again, one of the main interests was to ascertain whether any of the pilots faced difficulties during the experiment that could invalidate his or her data. That was not the case and no data were discarded.

Questions about factors motivating horizontal and vertical maneuvers were included, with results shown in Figures A.4 and A.5. Of particular note, the suggestive guidance ranked higher compared to Experiment 1—despite no change to its appearance on the display. The cause of this is not known. Perhaps, as some pilots in this experiment also participated in Experiment 1, familiarity with the suggestive guidance improved its usefulness.

Concluding comments (verbatim) made by the pilots after the experiment are as follows:

- With the info that the system provided I would [try] to correct before the ATC clearance came back.
- Good facsimile of VGCS/OSGCS (Army) GCS command system.
- Seemed to be 30 sec of each run, at the end, that added little value.
- I felt the display gave me good situational awareness and enough information to avoid traffic[.]

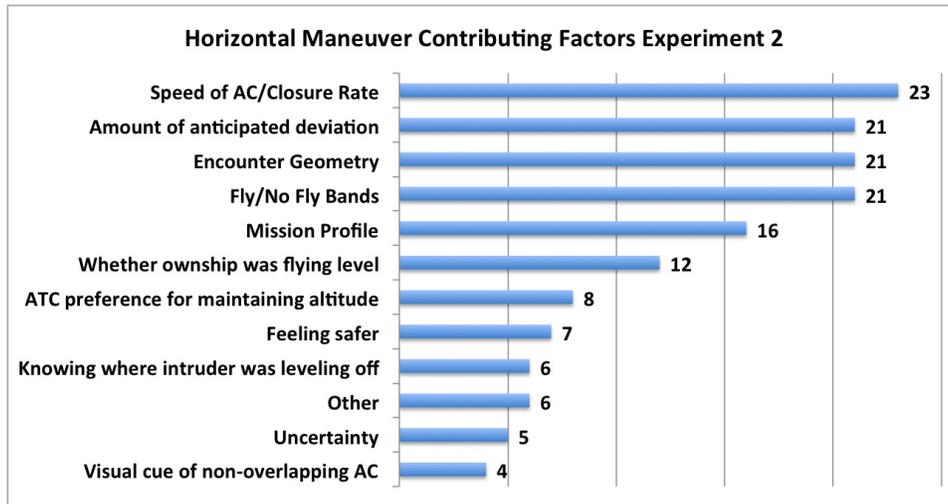


Figure A.4. Factors contributing to choice of horizontal maneuver in Experiment 2.

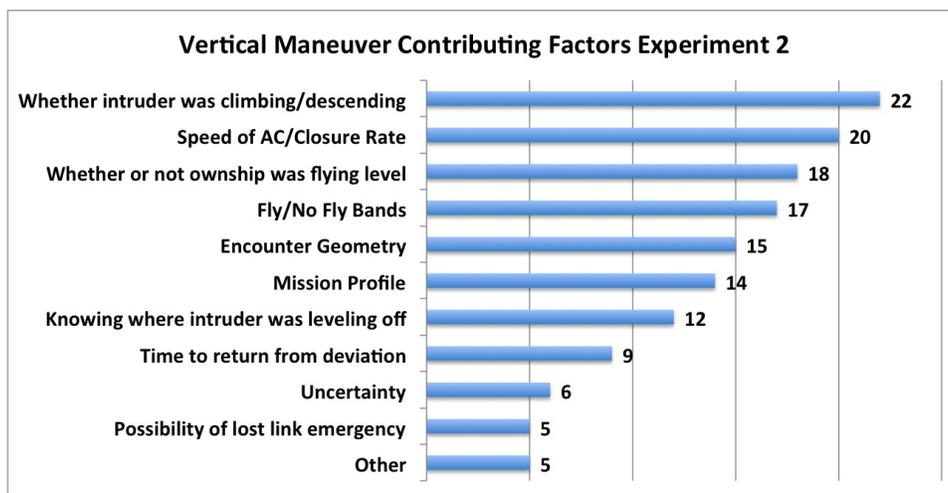


Figure A.5. Factors contributing to choice of vertical maneuver in Experiment 2.

- Small symbols or poor symbology make it hard to identify variations of yellow.
- Recommend adding other colors or changing shapes[.]
- More complex distractions could make the simulation more realistic[.]
- Very good TCAS-like display with great SA. Horizontal banding was great not available in TCAS II[.]
- Good Sim Representations and Data Collection, Thank you! Good Luck!, Well thought out + presented[.]
- It will be even more “fun”/interesting to work multiple conflicting targets.
- Vertical separation predictor requires higher workload in a red warning alert. Needs better visual cues.
- Not a good simulation for RQ-4 where you primarily minimize time in the terminal environment and you almost always prefer to climb above conflicts if possible[.]
- Moving map integration would be helpful. Good visual indications of avoidance possibilities provided by the GUI.
- I felt the A/C w/ degraded performance required more planning as it didn’t turn or climb well. Sometimes when the A/C info box crossed the mile rings I couldn’t read the info.
- Program was easy to understand and helpful to do the practice exercises first.
- Include abnormal responses to tests (red, failures, abnormal readings)[.]

A.3 EXPERIMENT 3: MANEUVER UPDATES

Experiment 3 did not include a survey. The participants were, however, encouraged to “think aloud” while completing the experiment. The key points gleaned from those remarks by the test directors are as follows:

- Many turned behind the intruder when maneuvering horizontally and leveled off for vertical maneuvers.
- Most thought heading bands were intuitive but altitude bands somewhat misleading.
- Aggressive intruders caused more than typical use of vertical maneuvers.
- Pilots appreciated green heading bands indicating that returning to course would be safe.
- Many learned to wait for the alerts and guidance in making both initial and updated maneuvers after observing if the prescribed maneuver was sufficient for safety.
- Some pilots verbally disagreed with type, direction, and magnitude of prescribed maneuvers.

- Some felt they needed to compensate a bit when surveillance was inaccurate.
- Some encounters were held to be unrealistically challenging with an aggressive intruder “like a magnet.”
- There was a tendency to maneuver more than necessary, but pilots learned more finesse with practice.
- Some pilots remarked about unrealistic flying characteristics of intruders and/or encounters.
- Sometimes pilots disagreed with the suggestive guidance bands, feeling that the suggestions were either unreasonable or lagging.
- Some felt the display encourages horizontal maneuvers even when vertical is preferred.
- Only some pilots considered the right-of-way rules; general aviation pilots were observed to be the most likely to comply.
- The appearance of recovery bands was generally held to be prominent, and most pilots responded quickly when recovery guidance appeared.

This page intentionally left blank.

APPENDIX B: MODEL PARAMETER SUMMARY

Values for each of the model parameters, as derived from the experimental data, are summarized here.

TABLE B.1
Pilot Model Parameter Values

Parameter	Description	Stochastic	Deterministic
initialDelayMin	Minimum time of initial delay from alert to to beginning of ATC coordination. [seconds]	0.0	5.0
initialDelayMu	Mean of exponentially distributed random component of initial delay before beginning ATC coordination. [seconds]	5.0	0.0
coordDelayMin	Minimum time of delay for coordination with ATC. [seconds]	0.0	11.0
coordDelayK	Shape parameter of gamma distributed random component of ATC delay.	5.5	0.0
coordDelayTheta	Scale parameter of gamma distributed random component of ATC delay.	2.0	0.0
executionDelayMin	Minimum time of execution delay from end of ATC coordination to transmission of maneuver command to own aircraft. [seconds]	0.0	3.0
executionDelayMu	Mean of exponentially distributed random component of execution delay. [seconds]	3.0	0.0
minUpdateTime	Minimum time to next maneuver decision by alert level. [seconds]	[12,12,9,6,6,0]	[12,12,9,6,6,0]
meanUpdateTime	Mean time in addition to minimum to next maneuver decision by alert level. [seconds]	[12,12,6,3,3,3]	[12,12,6,3,3,3]

Continued on next page

TABLE B.1 – continued from previous page

Parameter	Description	Stochastic	Deterministic
probFollowMinDev	Preference weight for maneuvering in the direction of the smaller suggestion for each maneuver plane.	0.8	1.0
probLeftTurn	Preference weight for choosing left turn over right if minimum suggestion is inconclusive.	0.6	1.0
maxRelativeHdg	Maximum heading difference for which the larger of the heading suggestions may be selected. [degrees]	40.0	40.0
probDescend	Preference weight for choosing to descend rather than climb if minimum suggestion is inconclusive.	0.5	0.0
probTurn	Preference weight for choosing to turn rather than climb or descend as a function of ownship maneuverability.	[0.5, 1.0, 1.0; 0.0, 0.5, 0.65; 0.0, 0.5, 0.4]	[1, 1, 1; 1, 1, 1; 1, 1, 1]
turnK	Shape parameter for gamma distributed horizontal maneuver magnitude relative to the minimum suggestion when direction selected complies with the minimum suggestion heuristic (primary turn direction).	6.21	0.0
turnTheta	Scale parameter for gamma distributed horizontal maneuver magnitude relative to the minimum suggestion in the primary turn direction	9.67	0.0
turnOffset	Offset of turn magnitude gamma distribution for turns in the primary direction.[degrees]	-30.0	30.0
turnK_alt	Shape parameter for gamma distributed horizontal maneuver magnitude relative to the minimum suggestion for turns opposite the primary direction.	5.47	5.0

Continued on next page

TABLE B.1 – continued from previous page

Parameter	Description	Stochastic	Deterministic
turnTheta_alt	Scale parameter for gamma distributed horizontal maneuver magnitude relative to the minimum suggestion for turns opposite the primary direction.	8.25	0.0
turnOffset_alt	Offset of turn magnitude gamma distribution for turns in the primary direction.[degrees]	-30.0	15.0
altitudeK	Shape parameter for gamma distributed vertical maneuver magnitude relative to the minimum suggestion in the selected direction.	9.73	0.0
altitudeTheta	Scale parameter for gamma distributed vertical maneuver magnitude relative to the minimum suggestion in the selected direction.	207.98	0.0
altitudeOffset	Offset of vertical maneuver magnitude gamma distribution. [feet]	-1500.0	0.0

This page intentionally left blank.

GLOSSARY

A/C	Aircraft
ADS-B	Automatic Dependent Surveillance-Broadcast
ATC	Air Traffic Control
CASSATT	Collision Avoidance System Safety Assessment Tool
DAA	Detect and Avoid
FL	Flight Level
GUI	Graphical User Interface
HALE	High Altitude Long Endurance
HITL	Human-in-the-Loop
HMD	Horizontal Miss Distance
IFR	Instrument Flight Rules
LL	Lincoln Laboratory
LoWC	Loss of Well Clear
MALE	Medium Altitude Long Endurance
MFD	Multifunction Display
MIT	Massachusetts Institute of Technology
MNVR	Maneuver
MOPS	Minimum Operational Performance Standard
NAS	(United States) National Airspace System
NMAC	Near Midair Collision
PFD	Primary Flight Display
SA	Situational Awareness
SAA	Sense and Avoid (equivalent to DAA)
TCAS	Traffic Alert and Collision Avoidance System
UAS	Unmanned Aircraft System
VMD	Vertical Miss Distance
VVI	Vertical Velocity Indicator

This page intentionally left blank.

REFERENCES

- [1] FAA Modernization and Reform Act of 2012, Title III, Subtitle B, Section 332.
- [2] Code of Federal Regulations, Title 14, Chapter 91 (2011).
- [3] RTCA, “Minimum operational performance standards (MOPS) for Detect and Avoid (DAA) Systems,” RTCA (2017).
- [4] A. Zeitlin, A. Lacher, J. Kuchar, and A. Drumm, “Collision avoidance for unmanned aircraft: Proving the safety case,” MITRE, McLean, VA, Report MP060219 (2006).
- [5] J.D. Griffith and J. Kuchar, “Evaluation of TCAS on Global Hawk with US airspace encounter model,” MIT Lincoln Laboratory, Lexington, MA, Project Report ATC-353 (2009).
- [6] S. Cook, D. Brooks, R. Cole, D. Hackenberg, and V. Raska, “Defining well clear for unmanned aircraft systems,” in *AIAA Infotech@Aerospace*, Kissimmee, Florida (2015).
- [7] D. McRuer and E. Krendel, “Mathematical models of human pilot behavior,” Advisory Group for Aerospace Research and Development, AGARDograph No. 188 (1974).
- [8] J.W. Andrews, “Air-to-air visual acquisition handbook,” MIT Lincoln Laboratory, Lexington, MA, Project Report ATC-151 (1991).
- [9] R. Walden and W. Rouse, “A queueing model of pilot decision making in a multitask flight management situation,” *IEEE Transactions on Systems, Man and Cybernetics* 8(12), 867–875 (1978).
- [10] K. Sarno and C. Wickens, “The role of multiple resources in predicting time-sharing efficiency: An evaluation of three workload models in a multiple task setting,” *International Journal of Aviation Psychology* 5(1), 107–130 (1995).
- [11] C. Lowe and J. How, “Learning and predicting pilot behavior in uncontrolled airspace,” in *AIAA Infotech@Aerospace*, Kissimmee, Florida (2015).
- [12] K. Ogaard and R. Marsh, “Discovering probabilistic models of pilot behavior from aircraft telemetry data,” in Y. Ohsawa and A. Abe (eds.), *Advances in Chance Discovery*, Springer, pp. 191–214 (2013).
- [13] N. Ernest, D. Carroll, C. Schumacher, M. Clark, K. Cohen, and G. Lee, “Genetic fuzzy based artificial intelligence for unmanned combat aerial vehicle control in simulated air combat missions,” *Journal of Defense Management* 6(1) (2016).
- [14] K. Virtanen, T. Raivio, and R. Hämäläinen, “Modeling pilot’s sequential maneuvering decisions by a multistage influence diagram,” *Journal of Guidance, Control, and Dynamics* 27(4), 665–677 (2004).
- [15] M. McLaughlin and A. Zeitlin, “Safety study of TCAS II for logic version 6.04,” U.S. Department of Transportation, Report DOT/FAA/RD-92/22 (1992).

- [16] E. Londner, “Interoperability of horizontal and vertical resolution advisories,” in *11th USA/Europe Air Traffic Management Research and Development Seminar*, Lisbon, Portugal (2015).
- [17] RTCA, “Minimum operational performance standards for Traffic Alert and Collision Avoidance System II (TCAS II),” RTCA (2008).
- [18] M. Kochenderfer, J. Holland, and J. Chryssanthacopoulos, “Next generation airborne collision avoidance system,” *Lincoln Laboratory Journal* 19(1), 55–71 (2012).
- [19] C. Lividas, J. Lygeros, and N. Lynch, “High-level modeling and analysis of TCAS,” in *IEEE 34th Real-Time Systems Symposium*, Phoenix, Arizona (1999).
- [20] S. Taguchi, S. Sekizawa, S. Inagaki, T. Suzuki, S. Hayakawa, and N. Tsuchida, “Stochastic modeling and analysis of drivers’ decision making,” in *IEEE International Conference on Systems, Man and Cybernetics*, Montreal, Canada (2007).
- [21] R. Lee and D. Wolpert, “Game theoretic modeling of pilot behavior during mid-air encounters,” in T. Guy, M. Kárný, and D. Wolpert (eds.), *Decision Making with Imperfect Decision Makers*, Springer, pp. 75–111 (2012).
- [22] D. Maki, C. Parry, K. Noth, M. Molinario, and R. Miraflor, “Dynamic Protection Zone alerting and pilot maneuver logic for ground based sense and avoid of unmanned aircraft systems,” in *AIAA Infotech@Aerospace*, Garden Grove, California (2012).
- [23] J. Hoekstra, R. van Gent, and J. Groenweg, “Airborne separation assurance validation with multiple humans-in-the-loop,” in *5th USA/Europe Air Traffic Management Research and Development Seminar*, Budapest, Hungary (2003).
- [24] S. Ellis and E. Palmer, “Threat perception while viewing single intruder conflicts on a cockpit display of traffic information,” NASA, Technical Memorandum 81341 (1982).
- [25] L. Thomas and C. Wickens, “Display dimensionality and conflict geometry effects on maneuver preferences for resolving in-flight conflicts,” *Human Factors* 50(4), 576–588 (2008).
- [26] C. Wickens, J. Hellenberg, and X. Xu, “Pilot maneuver choice and workload in free flight,” *Human Factors* 44(2), 171–188 (2002).
- [27] E. Rantanen and C. Wickens, “Conflict resolution maneuvers in Air Traffic Control: Investigation of operational data,” *International Journal of Aviation Psychology* 22(3), 266–281 (2012).
- [28] E. Rantanen, J. Yang, and S. Yin, “Comparison of pilots’ and controllers’ conflict resolution maneuver preferences,” in *Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting* (2006).
- [29] R. Rorie, L. Fern, and R. Shively, “The impact of suggestive maneuver guidance on UAS pilots performing the detect and avoid function,” in *AIAA Infotech@Aerospace*, San Diego, California (2016).

- [30] C. Santiago and E. Mueller, “Pilot evaluation of a UAS detect-and-avoid system’s effectiveness in remaining well clear,” in *11th USA/Europe Air Traffic Management Research and Development Seminar*, Lisbon, Portugal (2015).
- [31] W.C. Stirling, *Satisficing Games and Decision Making*, Cambridge University Press (2003).
- [32] H.A. Simon, *The Sciences of the Artificial*, Cambridge University Press (1996).
- [33] A. Rubinstein, *Modeling Bounded Rationality*, MIT Press (1998).
- [34] J.C. Pomerol, *Decision-Making and Action*, ISTE (2012).
- [35] D. Kahneman, *Thinking, Fast and Slow*, Farrar, Strauss and Giroux (2011).
- [36] R. Selten, “What is bounded rationality?” in G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox*, MIT Press, pp. 13–36 (2001).
- [37] D. Kahneman and A. Tversky, “Prospect theory: An analysis of decision under risk,” *Econometrica* 47(2), 263–292 (1979).
- [38] G. Gigerenzer, “The adaptive toolbox,” in G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox*, MIT Press, pp. 37–50 (2001).
- [39] M. Kuffner, R. Guendel, and S. Darrah, “Investigating traffic avoidance maneuver decisions of unmanned aircraft pilots,” in *16th AIAA Aviation Technology, Integration, and Operations Conference*, Washington, DC (2016).
- [40] C. Munoz, A. Narkawicz, G. Hagen, J. Upchurch, A. Dutle, M. Consiglio, and J. Chamberlin, “DAIDALUS: Detect and Avoid Alerting Logic for Unmanned Systems,” in *34th IEEE/AIAA Digital Avionics Systems Conference*, Prague, Czech Republic (2015).
- [41] L. Fern and R.C. Rorie, “An examination of unmanned aircraft systems pilots’ interaction with Air Traffic Control while responding to Detect and Avoid conflicts,” *Air Traffic Control Quarterly* 23(2/3), 113–135 (2015).
- [42] J. Kuchar, “Safety analysis methodology for Unmanned Aerial Vehicle (UAV) collision avoidance systems,” in *6th USA/Europe Air Traffic Management Research and Development Seminar*, Baltimore, Maryland (2005).
- [43] R. O’Keefe and S. O’Leary, “Expert system verification and validation: A survey and tutorial,” *Artificial Intelligence Review* 7(1), 3–42 (1993).
- [44] D. Woods and S. Dekker, “Anticipating the effects if technological change: a new era of dynamics for human factors,” *Theoretical Issues in Ergonomic Science* 1(3), 272–282 (2000).
- [45] M. Kochenderfer, L. Espindle, J. Kuchar, and J. Griffith, “Correlated encounter model for cooperative aircraft in the National Airspace System version 1.0,” MIT Lincoln Laboratory, Lexington, Massachusetts, Project Report ATC-344 (2008).

This page intentionally left blank.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 24-01-2017		2. REPORT TYPE Project Report		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE A Model of Unmanned Aircraft Pilot Detect and Avoid Maneuver Decisions				5a. CONTRACT NUMBER FA8721-05-C-0002 & FA8702-15-D-0001	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) R.E. Guendel, M.P. Kuffner, and D.E. Maki				5d. PROJECT NUMBER 1844	
				5e. TASK NUMBER 11	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02421-6426				8. PERFORMING ORGANIZATION REPORT NUMBER ATC-434	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFLCMC/WING Building 557 1640 Loop Road West Wright Patterson AFB, OH 45433-7106				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release: distribution unlimited; Case Number 88 ABW-2017-1669, 10 Apr 2017.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT For unmanned aircraft to share airspace with manned aircraft, extensive testing is first required to ensure that such vehicles can fly safely with manned traffic. Safe operation includes not only avoiding collisions with other traffic but also complying with the Federal Aviation Regulations to remain well clear of other traffic. One method for investigating the safety of unmanned aircraft operations is fast-time Monte Carlo simulation of encounters between unmanned and manned aircraft. As part of that simulation, one must model how the pilots of unmanned aircraft react to the encounters. To that end, an empirical, rule-based stochastic model of responses of unmanned aircraft pilots has been constructed based on data collected from a succession of human-in-the-loop experiments. This report details the main elements of that model and demonstrates its use in a safety analysis.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as report	18. NUMBER OF PAGES 82	19a. NAME OF RESPONSIBLE PERSON
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code)

This page intentionally left blank.