

Human Inspiration for Autonomous Vehicle Tactics

by

Jonathan Scott Beaton, ENS (USN)

B.S. Aerospace Engineering

United States Naval Academy, 2004

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of
Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2006

©2006 Jonathan Beaton. All rights reserved.


The author hereby grants to MIT permission to reproduce and to
distribute publicly paper and electronic copies of this thesis document
in whole or in part in any medium now known or hereafter created.

Author
Department of Aeronautics and Astronautics
May 26, 2006

Certified by
Christopher Dever
Charles Stark Draper Laboratory
Thesis Supervisor

Certified by
Brent Appleby
Lecturer in Aeronautics and Astronautics
Thesis Supervisor

Accepted by
Jaime Peraire
Professor of Aeronautics and Astronautics
Chair, Committee on Graduate Students



ARCHIVES

Human Inspiration for Autonomous Vehicle Tactics

by

Jonathan Scott Beaton, ENS (USN)

Submitted to the Department of Aeronautics and Astronautics
on May 26, 2006, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

Tactical control is needed in environments characterized by uncertainty and continuous, dynamic change. Given the likelihood of time constraints and high risks associated with poor tactical choices, current autonomous vehicles do not possess the decision making abilities to successfully perform in these environments. However, human experts frequently operate in these domains where they are forced to make quick, reactive decisions based on incomplete information. We propose, then, that the first step in augmenting autonomous vehicles (AVs) with improved tactical control capabilities is to learn, encode, and apply tactics exhibited by human experts. To test the method, five human subjects were given the task of performing an armed reconnaissance mission in a simulation environment over multiple cases with varying terrain and probability of enemy contact. By scoring the performance in each case, the best actions and decisions were filtered out and analyzed in depth to understand the strategies and tactics behind them. Human cognitive models and decision making theories were utilized to determine the cognitive processes underneath the decisions as displayed by the human subjects' think aloud reports and surveys. A baseline autonomous vehicle controller was designed independent of the human-in-the-loop experiments that could also perform the reconnaissance mission. After capturing the human tactics and encoding them into statechart form, a revised AV displayed a superior ability to engage enemy contacts uncovered during the reconnaissance when compared to the baseline AV. A final framework is presented that outlines how to learn and apply human-inspired tactics in future settings.

Thesis Supervisor: Christopher Dever
Title: Charles Stark Draper Laboratory

Thesis Supervisor: Brent Appleby
Title: Lecturer in Aeronautics and Astronautics

Acknowledgments

I would like to thank . . .

Chris Dever and Brent Appleby, my thesis advisors, for their support in this thesis. I would especially like to thank Chris for his in-depth reviews and comments throughout the entire project.

The Draper Fellow volunteers who were willing to give up their time to try their hand at the experiments.

My Draper Fellow cubicle pals, Brian Mihok, Stephen Thrasher, and Tom Krenzke, for the good times, great oldies, and late nights.

My roommate, Drew Barker, for helping out so much at the apartment and continually checking up on me to make sure I was still alive.

And most importantly, my fiance and soon-to-be wife, Kara Huffman, for her smile, joy, and encouragement through these past two years together at MIT. We made it through, by God's grace.

Contents

1	Tactical Decision Making: Humans and Automation	15
1.1	Motivation	16
1.2	Subset of Control Hierarchy	18
1.3	Existing Work	20
1.3.1	Computer Generated Forces	20
1.3.2	Artificial Learning	22
1.3.3	Human/Machine Collaboration	22
1.3.4	New Questions	22
1.4	Two-Stage Tactical Control Paradigm	24
1.5	Thesis Overview	26
2	Humans and Automation: Expertise and Reliability	29
2.1	Functional Task Allocation	30
2.1.1	Automation and Human Capabilities and Limitations	30
2.1.2	The Right Knowledge	32
2.2	Team Centered Automation	34
2.2.1	Unique Demands of Tactical Control Environments	34
2.2.2	Definition of Automation	37
2.2.3	Human-Centered Automation Design	39
2.2.4	Lessons Learned from Improper Automation	42
2.2.5	Characteristics of a Team	45
2.3	Implications and Closing Thoughts	47
3	Human Expert Performance and Cognitive Modeling Efforts	49
3.1	Components of Cognitive Model	50
3.1.1	Long-Term Memory	50
3.1.2	Working Memory	54
3.2	Human Problem-Solving	58
3.2.1	Information-Processing Theory	58
3.2.2	General Problem-Solver	58
3.2.3	Means-End Analysis	60
3.2.4	Common Characteristics	62
3.2.5	Problem Solving to Decision Making	63
3.3	Human Decision Making	63
3.3.1	Normative Decision Making	65

3.3.2	Bounded Rationality and Satisficing	67
3.3.3	Human Subjective Probability	68
3.3.4	Heuristics and Biases	72
3.3.5	Summary of Human Subjective Probability and Decision Heuristics and Biases	82
3.3.6	Naturalistic Decision Making	82
3.4	Other Cognitive Frameworks	87
3.4.1	Intentions, Plans, and Practical Reasons	88
3.4.2	Generic Error Modeling System	91
3.5	Expertise in Cognitive Processing	94
3.5.1	Experts versus Novices	95
3.5.2	Problems with Learning from Experience	96
3.6	Conclusions from Cognition	97
4	Procedures and Methodology	99
4.1	Experimental Framework	99
4.2	Human-In-The-Loop Experimental Setup	100
4.2.1	The Scenarios	100
4.2.2	The simulation environment	101
4.2.3	Display, Interface, Software	103
4.2.4	The Human Subjects	106
4.2.5	Human factors considerations	106
4.2.6	Time Line	108
4.2.7	Presentation of Experiments to Human Subjects	108
4.3	Autonomous Comparison	109
4.3.1	Vehicle Governing Behavior and Equations	110
4.3.2	Enemy Interactions	112
4.3.3	Parameters of Simulated Entities	114
4.3.4	Statecharts	115
4.4	Limitations of Experiments	126
4.4.1	Experts?	126
4.4.2	Real-time Interaction	127
4.4.3	Arbitrary Scaling	127
4.4.4	Wall-Clock Time versus Simulink Time	128
4.4.5	Varying Processing Rates	129
4.4.6	Human and Automation Differences	129
4.4.7	Learning, Training Effect, and Assumption Violation	130
5	Results of Human-in-the-Loop Experiments	131
5.1	Performance Metrics	131
5.1.1	Search Score	131
5.1.2	Engage Score	132
5.1.3	Summary and Bounds	133
5.2	Human and Baseline AV Scores	134
5.3	Finding Human Expert Performance	138

5.3.1	Fighting the Tank	139
5.3.2	Fighting the UAV	143
5.3.3	Fighting the SAM	150
5.4	Results of Human-Inspired Tactics	154
5.4.1	Against the Tank	157
5.4.2	Against the UAV	157
5.4.3	Against the SAM	160
5.5	Addressing Limitations by Human Cognitive Understanding	162
5.5.1	Harnessing Human Strengths	162
5.5.2	New Mental Model	164
5.5.3	Implications	165
5.5.4	Engagement Tactics Conclusions	166
5.6	Search Problem and Planning	166
5.6.1	Plans	167
5.6.2	Partial Plans and Intentions	174
5.6.3	Intention Examples	175
5.6.4	Intent Function	175
5.6.5	Rational Intent	177
5.6.6	Search Problem Conclusions	179
6	Conclusions and Tactical Framework	181
6.1	Human-in-the-Loop Experimental Design	183
6.2	Identifying Expert Performance	183
6.3	Training	184
6.4	Analysis through Cognitive Frameworks	185
6.4.1	Generic Error Modeling System	185
6.4.2	Recognition-Primed Decision Model	185
6.4.3	Belief, Desire, Intent Paradigm	186
6.5	Combining into Strategies	187
6.6	Translation	187
6.7	Encoding Tactics	188
6.7.1	Integrating Reactive and Planning Elements into Tactical Control	189
6.7.2	Missing Information?	189
6.8	Testing over Large Sample	191
6.9	Testing for Reliability	191
7	Future Work	193
7.1	Representing Intent	193
7.1.1	Fuzzy Logic	193
7.1.2	Strategies and Fuzzy Rules	196
7.2	Identification	196
A	Introductory Documents and Presentation of Scenarios to Human Subjects	199

B Surveys	205
C Think Aloud Reports	211
C.1 Round 2	211
C.1.1 S1	211
C.1.2 S2	222
C.1.3 S3	229
C.1.4 S4	237
C.1.5 S5	243

List of Figures

1-1	Spectrum of control levels for autonomous vehicles.	19
1-2	Unmanned laser designator statechart from Mark Hickie’s thesis. . . .	23
1-3	Systematic two-stage process of integrating tactical control logic into unmanned vehicles.	25
2-1	Flow chart showing application of the model of types and levels of automation [54].	41
2-2	Automation usage based on human factors [53].	45
3-1	British psychology students’ probability of recalling American presi- dents and probability of attributing the above quote to an American president [57].	54
3-2	General Problem Solver [47].	59
3-3	Steps of a heuristic search method [47].	61
3-4	Key processes and components involved in an information processing model of decision making [87].	64
3-5	Human distortion of probability [87].	68
3-6	Human biases in estimating variance [87].	69
3-7	Hypothetical aircraft crashes modeled by a Bernoulli process with con- stant probability $p = 0.01$	71
3-8	Displaying two aspects of prospect theory [15].	78
3-9	Recognition-primed decision model [38].	84
3-10	Rasmussen’s performance level framework and the Generic Error-Modeling System [57].	93
3-11	Luchins’ water jars problem [58].	94
4-1	Two-window visual display of simulation.	104
4-2	Visual cues.	105
4-3	Human-in-the-loop experimental setup.	105
4-4	The AV’s top level nominal and reactive superstates.	116
4-5	Search superstate - the search behavior of the AV.	117
4-6	Move state - diagonal searching between waypoints.	118
4-7	Enemy_Contact superstate - detection of enemy contact.	119
4-8	Enemy_Contact.ENEMY - three engagement modes.	121
4-9	Enemy_Contact.ENEMY.AVOID - maneuvering to avoid enemy. . . .	122

4-10	Enemy_Contact.ENEMY.ENGAGE_TA - maneuvering to engage the enemy with a tactical advantage.	124
4-11	Enemy_Contact.ENEMY.ENGAGE_TD - maneuvering to engage the enemy with a tactical disadvantage.	125
5-1	Human and baseline AV performance for the first round of experiments.	135
5-2	Human and baseline AV performance for the second round of experiments.	136
5-3	Human performance against enemy tanks.	140
5-4	Human-inspired statechart to engage enemy tank.	142
5-5	Human performance against enemy UAVs.	144
5-6	Human-inspired statechart to engage enemy UAV.	149
5-7	Human performance against enemy SAMs.	151
5-8	Top-level of human-inspired statechart to engage enemy SAM.	154
5-9	AVOID state of human-inspired statechart to engage enemy SAM.	155
5-10	ENGAGE_TD state of human-inspired statechart to engage enemy SAM.	155
5-11	Monte Carlo simulation to test human-inspired tactics against enemy tank.	158
5-12	Monte Carlo simulation to test human-inspired tactics against enemy UAV.	159
5-13	Monte Carlo simulation to test human-inspired tactics against enemy SAM.	161
5-14	Very good chance of enemy contact in both the air corridor and the critical area.	168
5-15	Possible chance of enemy contact in the air corridor and a very good chance in the critical area.	169
5-16	Slim chance of enemy contact in both the air corridor and the critical area.	171
5-17	Possible chance of enemy contact in the air corridor and a very good chance in the critical area.	172
5-18	Possible chance of enemy contact in the air corridor and a slim chance in the critical area.	173
6-1	Framework for learning and applying human-inspired tactics.	182
6-2	Simulink diagram of reactive and planning tactical control for the AV.	190
7-1	Overlapping membership functions.	194
7-2	Membership functions that always result in 1.0 degree of membership.	195

List of Tables

2.1	Comparison between human and machine.	32
2.2	Air Force Research Laboratory levels of autonomous control	38
2.3	Army Future Combat Systems levels of autonomous behavior	38
2.4	Levels of Automation of Decision and Action Selection	39
3.1	The central executive	56
4.1	Historically Validated Unmanned Aerial System (UAS) Roles [50]	102
4.2	Independent variables.	109
4.3	Defining parameters of simulation platforms.	115
5.1	Performance Metrics.	133

Chapter 1

Tactical Decision Making: Humans and Automation

The purpose of this thesis is to address the following two related questions:

1. How can the tactical decision making capabilities of human experts be learned and transferred over to an autonomous vehicle?
2. Can a human expert learn how to exploit vehicle-specific dynamics in tactical scenarios to achieve high levels of performance for goal-oriented missions?

Autonomous vehicles (AVs) currently perform reconnaissance missions [50]. How can they be improved to perform armed reconnaissance missions? What are the major limitations in control systems that exclude AVs from participating in such tactical scenarios?

This research proposes that the missing piece is the ability to generate appropriate decisions when the environment is both uncertain and dynamic. Human experts operating in tactical environments must make decisions all the time based upon information that is uncertain and/or missing. In fact, these human experts train so much for these environments because of the uncertainty, the decision making process becomes “intuitive.” Selecting courses of action based on “intuition” is extremely hard to quantify. Furthermore, human experts operating in tactical environments must quickly generate decisions due to the dynamic change in the environment. This time-constrained, on-line computation problem is very difficult to implement. Finally, human experts operating in tactical environments must respond appropriately to novel situations. To write software code for unknown and unforeseen circumstances is an enormously tough challenge. All of these difficulties for AVs are an accepted element in the human expert’s task. Therefore, this research proposes that the first step in creating better decision making AVs is to learn from the human expert.

1.1 Motivation

AVs in the battlefield do not yet have the control capabilities to replace pilots in tactical situations. The major advantage of integrating AVs into today's air forces is that the probability of losing human pilots decreases, especially in extremely risky and dangerous missions. Military commanders are less and less willing to send pilots into situations where the expectation of loss of life is high compared to the perceived benefit of destroying a specific target. This reluctance is because we place a high value on human life. Military commanders would be more willing to risk losing an AV in exchange for destroying a high value target.

Because the theory of decision making is a central theme in this research, it is important to show at the outset the mathematical validity of the above reasoning. The military commander's choice to risk an AV rather than a human life can be shown to be mathematically valid, or rational, using techniques from classical or normative decision theory. According to classical decision theory, a choice is *rational* if it meets the following three criteria [15]:

1. It is based on the decision maker's *current* assets.
2. It is based on the possible consequences of the choice.
3. When these consequences are uncertain, their likelihood is evaluated without violating the basic rules of probability theory.

Mathematician John von Neumann and economist Oskar Morgenstern published *Theory of Games and Economic Behavior* in 1953 in which they presented the principle of expected utility in accordance with normative decision theory [83]. This principle guided decision makers in how to make rational decisions based on the three criteria above.

Mathematically, computing the expected utility of a decision is exactly identical to computing the expected value of an event. The principle of expected utility states that with each possible outcome, x_i , there is both an associated probability of its occurrence, $P(x_i)$, as well as a subjective personal utility (a gain or loss), $u(x_i)$. In terms of discrete probabilities, the expected utility of a choice is the sum of each possible gain or loss multiplied by its probability of occurrence, which is given by Equation 1.1.

$$U[X] = \sum_{i=1}^n u(x_i)P_X(x_i) \quad (1.1)$$

Note that $E[X]$ is the standard notation for expected value. Here $U[X]$ emphasizes that this is expected *utility*. Expected utility theory attempts to mathematically explain why different individuals make different decisions based on their own personal values.

Here, the military commander must compare the expected utility (personal value) of sending a human pilot into a risky and uncertain mission against sending an AV. Equation 1.2 describes the expected utility of sending the human pilot.

$$U[\text{sending in pilot}] = \alpha_1 P(\text{loss}) + \alpha_1' P(\text{life}) + \alpha_2 P(\text{success}) + \alpha_2' P(\text{failure}) \quad (1.2)$$

where α_1 , the loss associated with the pilot's death ($\alpha_1 < 0$), and α'_1 , the gain associated with the pilot living ($\alpha'_1 > 0$), are utilities according to the military commander's personal belief of the value of human life. Likewise, there is a gain and loss associated with the mission success or failure ($\alpha_2 > 0$ and $\alpha'_2 < 0$, respectively), which, for example, is the destruction (or not) of a high value target.

Because both loss (death) and life as well as success and failure are mutually exclusive (assuming no outcome of a damaged state), $P(\text{loss}) + P(\text{life}) = 1$ and $P(\text{success}) + P(\text{failure}) = 1$. Equation 1.2 can then be rewritten as following:

$$U[\text{sending in pilot}] = (\alpha_1 - \alpha'_1)P(\text{loss}) + \alpha'_1 + (\alpha_2 - \alpha'_2)P(\text{success}) + \alpha'_2 \quad (1.3)$$

Furthermore, there is an associated expected utility with sending in an AV, given by Equation 1.4.

$$U[\text{sending in AV}] = (\alpha_3 - \alpha'_3)P(\text{loss}) + \alpha'_3 + (\alpha_2 - \alpha'_2)P(\text{success}) + \alpha'_2 \quad (1.4)$$

where α_3 , the loss associated with losing the AV ($\alpha_3 < 0$), and α'_3 , the gain associated with the AV living ($\alpha'_3 > 0$), are utilities according to the military commander's personal belief of the value of the AV. Likewise, there is a gain and loss associated with the mission success or failure ($\alpha_4 > 0$ and $\alpha'_4 < 0$) for the AV. If we assume that the human and AV have the same performance skills in relation to the same enemy, then they have equal probabilities of loss, life, success and failure. Furthermore, since the destruction of the target is independent of who destroyed it, the gain and loss associated with $P(\text{success})$ and $P(\text{failure})$ are the same. Finally, if the military commander has an equal gain associated with bringing back all assets from the mission, then $\alpha'_1 = \alpha'_3$, the gains associated with $P(\text{life})$.

As discussed above, the military commander would rather risk losing an AV rather than a human, which is equivalent to $U[\text{sending in AV}] > U[\text{sending in pilot}]$. Substituting in Equations (1.4) and (1.3), canceling all like terms, and solving for the loss associated with $P(\text{death})$, we see that $\frac{\alpha_3}{\alpha_1} < 1$. Thus, all things being equal, the expected utility of sending in the AV is greater than that of sending in the pilot if and only if the military commander places a higher value on human life. Here, assigning a higher value to human life is a larger loss associated with $P(\text{death})$ of human life, α_1 is more negative than α_3 . Assuming the military commander does place a higher personal value on human life, the decision is both rational and mathematically sound. It is desirable, then, for AVs to perform the more dangerous and risky missions. However, the reality of today's battlefield is that humans and AVs do *not* have the same performance skills. Thus, there will be different probabilities of loss, life, success, and failure, and the military commander may or may not choose to send in the AV. This then, represents somewhat of an ultimate goal of this thesis. We desire to bring the performance skills of AVs in dangerous, risky scenarios up to par with human experts so that the military commander can make decisions according to the above framework.

The dynamic, uncertain, and multi-dimensional environment of tactical scenarios have all combined to limit AVs from performing these missions. On the other hand,

humans are trained to become experts in these operating environments. The disparity highlights the following strengths of humans: first, the ability to filter massive amounts of information and make quick decisions only on relevant cues [38]; second, the ability to estimate uncertain and missing information which affects the possible decision choices [36]; third, the ability to adapt to new circumstances and find creative solutions [18].

Therefore, the goal of this research is to systematically interrogate, represent, and encode military tactics, providing AVs with expert decision making capabilities. This level of expertise will complement manned platforms in two significant ways. First, it will provide more autonomy and more flexibility in the range of executable missions for AVs. Second, by bringing a human expert into the design loop of an AV's decision making ability, we hope to make AVs more predictable, trustworthy, and better understood by their manned counterparts. As AVs are integrated into today's battlefield where manned and unmanned platforms operate together, both of these outcomes are necessary so that humans and AVs can interact synergistically to enhance the overall team performance.

1.2 Subset of Control Hierarchy

There is a spectrum of decision-making levels or control hierarchy inherent in each AV-based mission carried out by the military. Figure 1-1 depicts this hierarchical structure extending from the high-level planning of a mission down to the lowest-level actuation of control devices to follow a desired trajectory. The middle layer of tactics, as defined for this research, exists somewhere between the operations research and optimal control problems. Before defining these tactics more specifically, a couple of assumptions must be stated. It is assumed, first of all, that the higher-level mission planning has already occurred, i.e. - the number of allocated resources (including AVs), the objectives of the mission, and the proposed routes and waypoints have already been decided. Second, it is assumed that the AV has a low-level trajectory generation algorithm and closed-loop control which executes dynamically feasible motion for the AV in response to its tactical decisions. Note that all of these control levels are tightly interwoven. A tactical decision made in response to a pop-up threat results in a replanning of the vehicle's desired route which can only be implemented by an inner-loop controller thereby utilizing all levels of control.

Tactical decision making must also be defined along with the scenarios and problem constraints. For this research, a scenario is the dynamic environment a human expert must operate in to carry out an assigned mission with specific objectives and goals. Examples include reconnoitering an air corridor in the desert to allow safe passage of troops to designated landing zones, launching off an aircraft carrier and searching for a downed pilot at sea, helping special operations forces laze a target for a Hellfire missile launch, and battle damage assessment after engagements. In each of these scenarios, the human has a set of constraints that limit the range of available tactical options: vehicle dynamic capabilities, rules of engagement, out-of-range communication limitations, survivability instincts, weapons sensor field-of-view, fuel

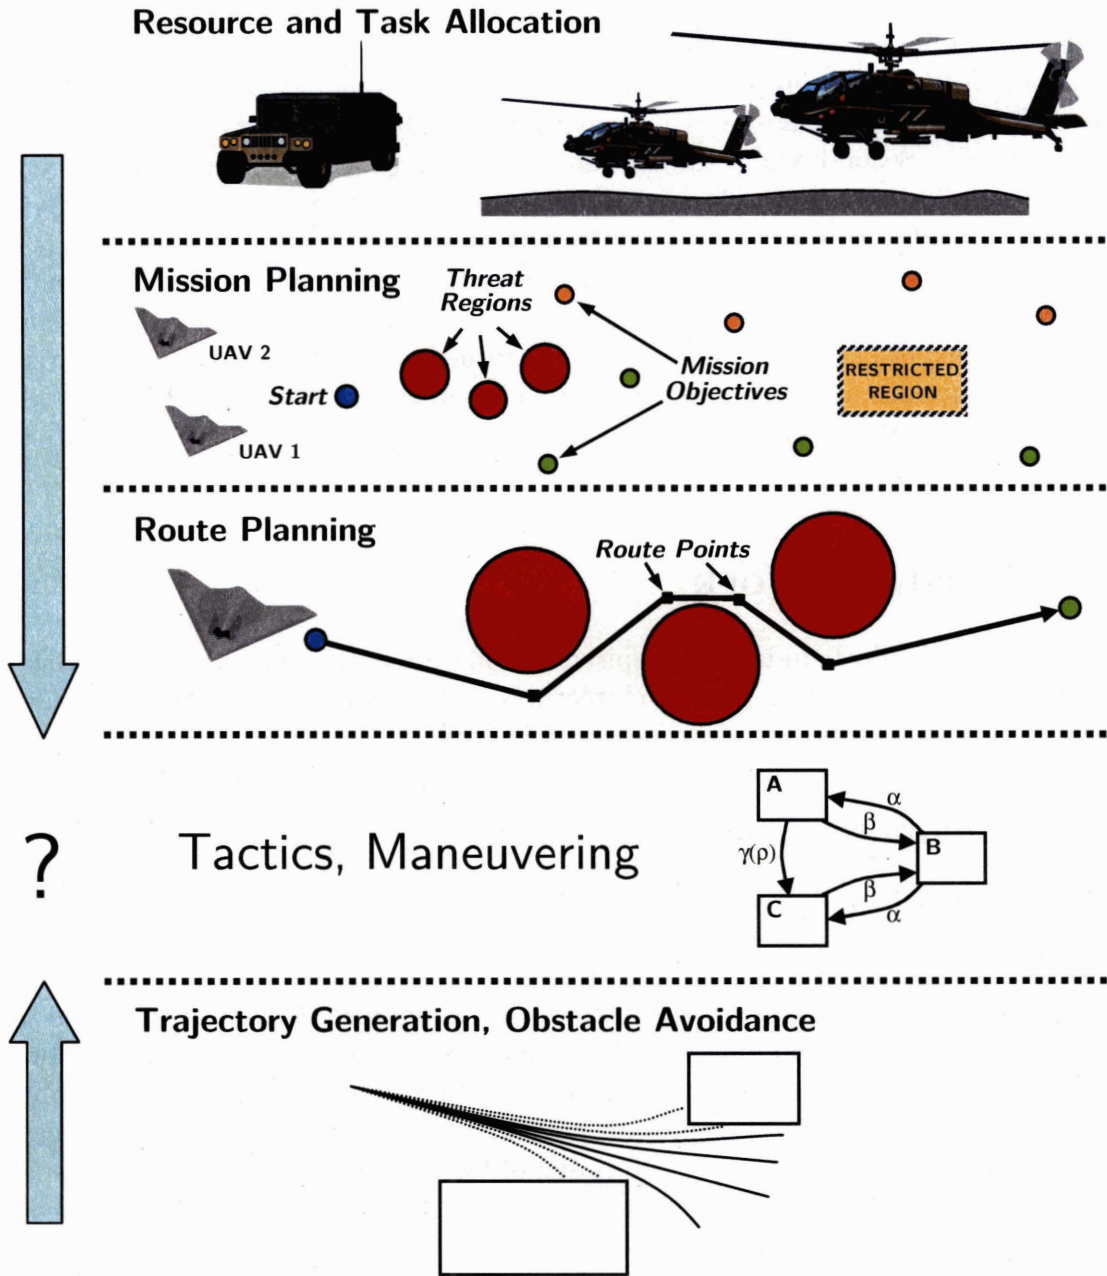


Figure 1-1: Spectrum of control levels for autonomous vehicles.

and ammunition remaining, etc. Therefore, the tactics the human employs in these scenarios are the decisions made, techniques employed, and actions taken to successfully carry out the mission in the face of dynamically changing environments while satisfying constraints.

By this definition, tactics can be treated in several ways. First, these tactics can be as simple as general rules-of-thumb. For example, if a threat of ground fire exists, say from rocket propelled grenades (RPGs) or small-arms fire, the helicopter pilot helping to scout out the urban environment must never hover, even if for just a few seconds. Second, these tactics can be strict protocol, such as flying specially designed profiles in and out of green zones so friendly forces can rapidly identify the vehicle as friendly irrespective of radio communication or identification-friend-or-foe codes. However, in their hardest form to quantify and represent are the human's reactive decisions to uncertain and dynamic environments. It is only through years of training and experience that humans form the intuition and skill to choose consistently appropriate actions and reactions in these environments. Thus, these tactics are the high risk, real time decisions humans make under time pressure. This third category of tactics is what we seek to understand and encode.

1.3 Existing Work

Current research efforts in human-inspired automation design have focused on the following three areas: creating agent architectures and human behavior models, artificial learning techniques, and human/machine collaboration.

1.3.1 Computer Generated Forces

By far, the largest amount of effort in human-inspired automation design is creating software agents or computer generated forces (CGFs) through human behavior models (HBMs). One of the most well-known HBMs is SOAR. This cognitive architecture uses the concepts from Newell and Simon's book *Human Problem-Solving* called universal sub-goaling and chunking [47, 55]. Universal sub-goaling is the continual break down of a goal into subgoals until a solution path is found by means-end analysis. Chunking is the recording of a solution path in memory. In the SOAR architecture, every problem-solving situation is composed of four elements: a goal, problem space, state, and operator. If the SOAR agent is given a problem that matches a previous chunk, it implements the already recorded solution. If not, the SOAR agent creates a problem space and searches for a solution path to the goal by sequentially applying operators to current states. After a problem has been solved, the SOAR architecture records the goal, problem space, states, operators, and the solution path in memory. This is a primitive cognitive architecture that has the potential to create a large knowledge base of chunks. In fact, the primary application of SOAR has been TacAir-SOAR, which has been used to model human pilots for large distributed military simulation exercises, and in 1991 it already contained over 5,200 rules or chunks [35].

Other more recent agent architectures have been implemented using the Belief, Desire, Intent (BDI) paradigm [7] and the Recognition-Primed Decision (RPD) model [38]. Both of these models will be discussed in detail in Chapter 3. BDI describes the human behavior to plan and coordinate as the result of the human's existing beliefs, desires, and intentions. The attractiveness of BDI is that it achieves a useful balance between planning and reaction-based behavior. The BDI paradigm has been applied in object-oriented models as the fundamental agent architecture for both distributed multi-agent systems [37] and single simplified agent design called "poor man's BDI" by the designers [3]. BDI has also been used to build Quake 2 agents using the programming language JACK [48]. As a final example, the SOAR architecture has been redefined in terms of BDI which has opened up new possibilities of interaction between combined SOAR and BDI architectures [27].

RPD describes the mechanisms of expert decision making in real world settings. It has been widely accepted because of its simplistic but accurate understanding of human decision making. A large effort in building CGFs with a RPD framework has been focused on evaluating RPD agent performance in the OneSAF Test Bed and air traffic controller environments [86]. Other work includes building agents to perform simulated driving tasks [71]. All of these agents have applied Hintzman's multiple-trace memory model [32] as the basis for storing goals, cues, expectancies, and actions that can be recalled by recognition routines. A final effort to mention is the application of RPD to a composite agent network, where separate agents are responsible for the four recognition by-products of goals, cues, expectancies, and actions [70].

CGF applications like TacAir-SOAR and BDI and RPD agents have one main requirement which differ from this research. The agent must appear human. The cognitive architecture, then, must incorporate variability in terms of decisions made, actions taken, and goals pursued both within and across simulated entities [75, 85, 91]. The implications, then, are twofold. First, it is not always desirable to choose the optimal solution for a given situation. Humans tend to either make emotional, irrational decisions or rational but suboptimal decisions. Second, human experts are mainly employed for critiquing the behavior of the simulated entity, i.e. - is it realistic? Therefore, the initial creation of the knowledge base, like TacAir-SOAR, arises mainly from military field manuals and tactics, techniques, and procedures documents which describe the normative process that should occur in tactical scenarios. Human experts become consultants to validate the "obvious" response a simulated entity should choose [35, 84]. However, this thesis focuses on forming a tactical knowledge base directly from the human expert. In this way, the human expert's role is not to be a consultant but a trainer. Furthermore, this research seeks to find the best decisions made and strategies used by human experts. In no way does it desire to create human-like AVs. Rather, the whole purpose is to extract only the human strengths of tactical decision making strategies and build them into AVs.

1.3.2 Artificial Learning

Another major area of research in human-inspired automation design is artificial learning. This machine learning method typically takes the form of a human expert performing a task in a simulator where data can quickly be collected and analyzed by the automated observer. In this area of synthetic learning, researchers have successfully implemented both a hybrid neural networks structure [65] and a combination of genetic programming and context-based reasoning [22] to learn motor vehicle control skills. In game theory, Bayesian networks have allowed automated players to learn tactics in a football simulation and perform consistently better against stronger opponents [28]. Lent and Laird describe a hierarchical operator structure based upon SOAR to learn to fly a racetrack pattern [82]. While artificial learning is certainly a time-saving approach to learning expert knowledge, it is not sophisticated enough to be applied to tactical environments characterized by uncertainty and dynamic change. There are many aspects of human decision making in these tactical environments that are difficult to quantify and can only be subjectively interpreted by another human. Therefore, it is important to emphasize that the purpose of this research is not artificial learning. This research recognizes the importance of a human interpreter because the goal is not how fast can tactics be learned and applied to an AV, but how plausible is the idea.

1.3.3 Human/Machine Collaboration

A final area of research that requires note is human/machine collaboration. The main focus of this research is how to find the optimal balance of decision making authority as appropriated between humans and machines. Malasky et al. simulated a command and control environment for planning and resource allocation and varied the levels of human/machine interaction [40]. Forest looked at human input in algorithmic design by varying when the human guided the design process [24]. Fan et al. created a multi-agent system where humans interacted with RPD-based agents [21]. This allowed for adaptive decision making between humans and agents where the more-experienced agents, as determined by recognition capability, had more decision making authority. Human/machine collaboration, as discussed in these sources, is an extremely important step in the philosophy of team-centered automation design. The only way that automation will be tightly integrated into all aspects of the battlefield is if it exhibits reliability in its decision making and engenders trust by putting the team's goals above its own. However, human/machine collaboration is the next step beyond this research. Once the AV possesses greater tactical control and decision making abilities, the question will then be how to pair up the AV with a manned asset, for example as a pilot's wingman [46].

1.3.4 New Questions

Existing research in human-inspired automation design addresses issues in this thesis - the need for a human cognitive architecture, the ability to learn from a human,

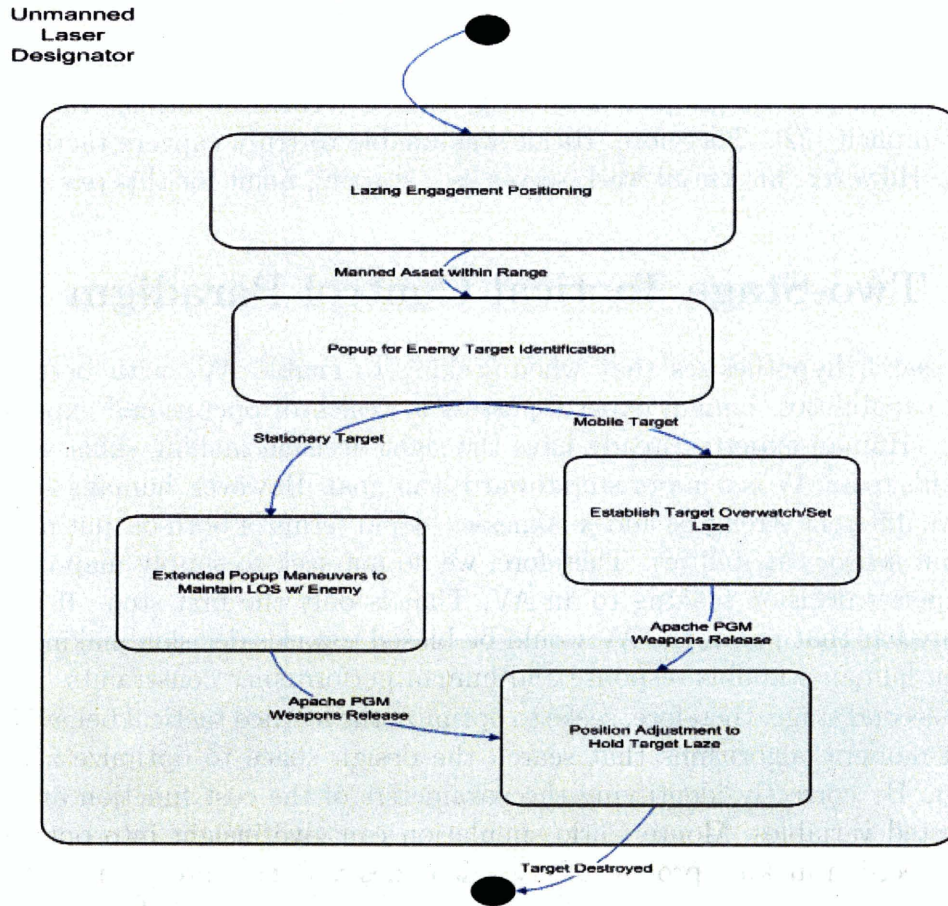


Figure 1-2: Unmanned laser designator statechart from Mark Hickie’s thesis.

a human/machine collaborative mindset to keep a team-centered automation design approach - but does not address the fundamental concern of how to discriminate and learn the best human tactics and apply them to an AV. As will be seen, part of answering this question is in answering another. The related issue is: can a human apply his or her tactical expertise to a given platform and use the platform’s capabilities as both a medium and springboard for tactical decision making? These questions have not been tackled in literature. Yet, an initial effort was made by Mark Hickie in his Master’s thesis [31]. By interviewing Army helicopter pilots, consulting field manuals, and running simulations in the U.S. Army’s force-on-force simulation tool, One Semi-Automated Forces (OneSAF) Testbed Baseline 2.0 (OTB 2.0), he identified, encoded, and validated military tactics for rotary wing AVs. He proposed statechart diagrams as the representation of tactical knowledge. For example, Figure 1-2 depicts his statechart representation of the decision-making process for an AV acting as an unmanned laser designator for a manned platform. However, Hickie was only able to show a small improvement in performance. One issue was the complexity and opacity of the OTB 2.0 software. The other issue was that his tactical knowledge base was too constrained. An expert’s knowledge can be described as either explicit

or implicit. While this is an overly-simplified categorization, it serves to differentiate between the explicit knowledge a human expert can consciously and verbally relate to an interviewer and the implicit knowledge that is termed intuition. Tactical knowledge is implicit [22]. Therefore, Hickie was unable to truly capture tactical decision making. However, his initial work serves as a starting point for this research.

1.4 Two-Stage Tactical Control Paradigm

This research hypothesizes that when seeking to enable AVs with better decision making capabilities, human experts possess a wealth of operational experience and training. Human experts already have the right decision making skills, and to map these skills to an AV is a major step towards this goal. However, humans and automation have different strengths and weaknesses [88] in terms of both output performance and input sensor capabilities. Therefore, we do not seek to simply map over the human expert's decision making to an AV. That is only the first step. If the process terminated at that point, the AV would be biased towards decision making strategies based on human stimulus response and human performance constraints.

The second stage, therefore, seeks to optimize the learned tactical behavior through modern numeric algorithms that search the design space to optimize a given cost function. By correctly identifying the parameters of the cost function and perturbing selected variables, Monte Carlo simulation can give insight into optimizing the tactical decision-making process of AVs. The desired end-state, then, is to develop a systematic way to integrate tactical behavior into unmanned vehicles in two steps. This two-stage tactical control paradigm is displayed in Figure 1-3. First, we learn from subject matter experts (SME) and arrive at a suboptimal solution derived from a finite number of cases. Second, we optimize the human-derived solution for the specific application and/or desired behavior for the unmanned vehicle.

As will be discussed later, the tactical knowledge learned from human experts will be encoded in statechart form [30]. In statechart form, discrete states and transitions drive the continuous vehicle motion. Thus, the optimization problem is a hybrid control problem consisting of discrete and continuous variables. Hybrid control is still a relatively new area of research and presents many complex challenges [6]. Different nonlinear optimization methods are being compared to help address this problem. In particular, the evolutionary genetic algorithms (GAs) offer new techniques [33]. Many algorithms require linearity, convexity, continuity, or only discrete variables. However, GA has no such requirement. If the problem can be encoded into a GA chromosome, GA can search nonlinear, hybrid, high-dimensional functions and find decent solutions. Ultimately, the goal is to find optimal solutions while guaranteeing performance and algorithm completeness. GA does not guarantee even a good solution, but it does perform well in many cases of interest. Furthermore, GA can be fast. Therefore, GA performance depends on the problem and on how it is applied to the problem, but it appears promising [78].

Alan Schultz and John Grefenstette showed how genetic algorithms can improve tactical plans [61]. They tested how well genetic algorithms could improve a set of

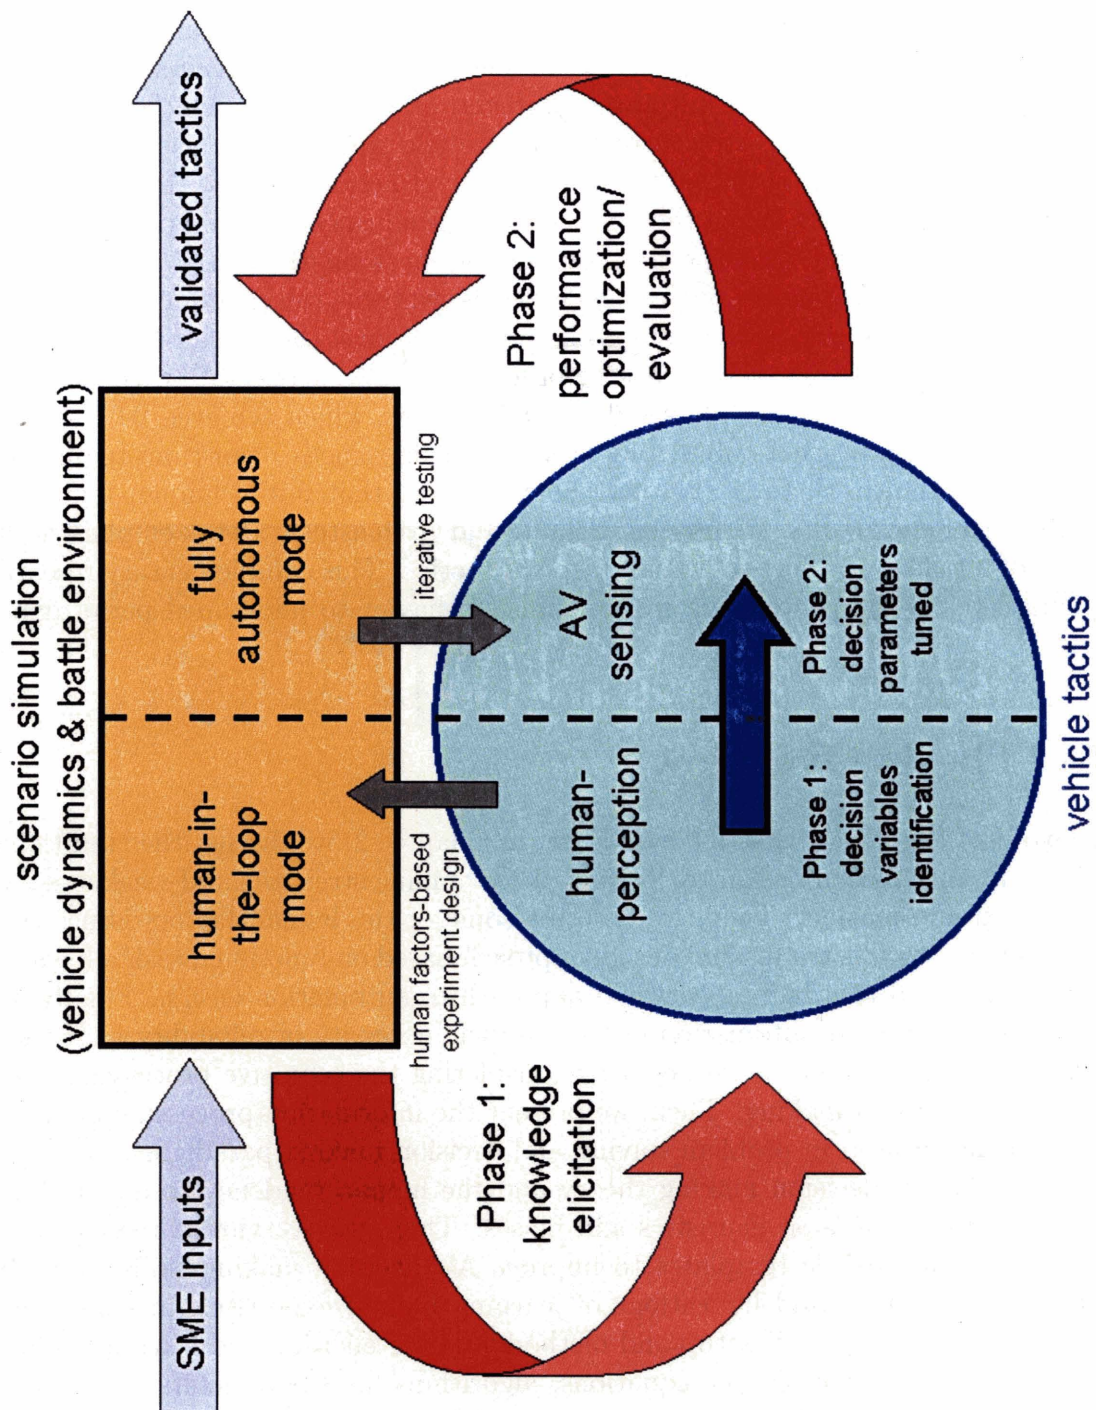


Figure 1-3: Systematic two-stage process of integrating tactical control logic into unmanned vehicles.

conditioned-action rules for an airplane evading a missile. Not only did they show the ability to learn decision rules and tactical plans that drastically improved the airplane's success rate in evading a missile, they compared the learning ability when initialized with different sets of rules. Specifically, they placed humans in charge of maneuvering the airplane away from the missile, and used their demonstrated rules as one initial set. The second initial set consisted of a single rule which stated that for any sensor inputs (i.e. - the airplane and missile states), take action X , where X is a random selection of one of the possible set of actions. Thus, this initial rule executed a random walk. When these two sets of initial rules were compared, Schultz and Grefenstette found that the performance of the human-inspired rules quickly rose to 95% success rate after only 50 generations, whereas the random walk rule set only achieved a little over 80%. Even after 100 generations, the human-inspired rules combined with genetic algorithm learning produced a 98% success rate with a smaller variance in performance than the random walk rule set which achieved 94% success rate. Therefore, there is not only strong reason for using genetic algorithms to improve tactical performance, there is also experimental proof that initial human input can produce better end results. Ultimately, Schultz and Grefenstette's work emphasize the benefits of a follow-on effort to human-inspired tactics. This complementary research will complete the transformation from manned behavior to unmanned performance [78].

1.5 Thesis Overview

The layout of this thesis is as follows. First, in Chapter 2, we discuss the traditional approach in automation design by looking at the unique strengths and weaknesses of humans and automation. Then, we describe some of the human factors issues that have arisen due to this traditional design approach based on functional task allocation and how they lead to a human-centered approach to automation design. Finally, we describe team-centered automation design and the necessity of reliability in tactical environments. In Chapter 3 we begin by considering the cognitive processes underlying human decision making. Then, we present the information-processing model of human cognition used in problem solving and decision making paradigms. Next, we analyze normative decision making theory and the human tendency to depart from the theory due to decision heuristics and biases. Then, we overview three cognitive frameworks that provide templates to improve AV decision making skills. Finally, we discuss the nature and limitations of learning human expertise. In Chapter 4, we discuss the experimental setup and methodology for eliciting and learning human strategies. We then present the equations, algorithms, and parameters that govern the simulated entities' behavior in the simulation. Next, we analyze in detail the baseline, untrained AV behavior. We conclude the chapter by discussing the limitations of the experimental method. In Chapter 5, we present the results of the human-in-the-loop experiments, the process of learning and encoding the human-inspired tactics into an improved AV, and the performance increase of the improved AV behavior over the baseline. In Chapter 6, we summarize our conclusions by proposing a frame-

work for learning and applying human-inspired tactics. Finally, we offer future work possibilities in Chapter 7.

In short, this research will show the ability to learn reactive engagement tactics from human experts and apply that knowledge to improve performance. A simulated baseline autonomous vehicle was developed that exhibited simple but logically coherent and reasonable behavior. This baseline behavior was designed independent of any human-in-the-loop testing. The baseline autonomous behavior was strictly reactive. Then, two rounds of experiments were designed to compare the performance of the human subjects and baseline behavior. The first round was to acclimate the human subjects to the simulation environment and provide enough training for them so that they could exhibit expert behavior. In the second round, the human subjects' performance was scored, and they clearly showed superior performance over the baseline. Through surveys, recorded verbal data, and observations of the human subject's actions, successful strategies and tactics to evade and engage pop-up threats were encoded in statechart form. This improved autonomous behavior was then compared to the baseline autonomous behavior through Monte Carlo simulation. This step was necessary to verify that the learned tactics were truly superior or if the small number of cases used to derive the tactics would limit the scope of their applicability.

This research will also present in detail how to formulate search strategies by focusing in on the humans' objectives, goals, and intentions. Reacting to enemy contacts and engaging them was only one part of the mission objectives given to the human subjects. They also had to divide their time between searching through different sections of terrain, with one part of the terrain being more important to completely search through than the other. Furthermore, the human subjects were given a probability of enemy contact in each of the terrain sections. Finally, by adding a time constraint to the scenario, the human subjects were forced to carefully consider how they could best accomplish the mission objectives of searching through as much terrain as possible, being especially cognizant to completely cover the more critical terrain, and reacting to enemy contact that could pursue and destroy the humans' vehicle. Therefore, we conclude with a framework of how to learn and apply human-inspired tactical knowledge that incorporates both reactive tactics and searching strategies.

Chapter 2

Humans and Automation: Expertise and Reliability

The familiar saying that unmanned aircraft are better suited for “dull, dirty, or dangerous” missions than manned aircraft presupposes that man is (or should be) the limiting factor in performing certain airborne roles. Although any flight can be dull or dangerous at times, man continues to fly such missions, whether because of tradition or as a substitute for technology inadequacies . . . The attributes that make the use of unmanned preferable to manned aircraft in the above three roles are, in the case of the dull, the better sustained alertness of machines over that of humans and, for the dirty and the dangerous, the lower political and human cost if the mission is lost, and greater probability that the missions will be successful. Lower downside risk and higher confidence in mission success are two strong motivators for continued expansion of unmanned aircraft systems. [50]

The above quote, from the Pentagon’s Unmanned Aerial Systems (UAS) Roadmap 2005, makes it clear that a major reason for investing in autonomy is human limitations. Autonomous vehicles (AVs) have longer endurance for dull missions, such as around-the-clock surveillance. They are also cheaper and more acceptable to risk losing when performing dirty missions, such as flying into the cloud of a dirty bomb explosion to determine its chemical makeup, or dangerous missions, such as destroying multiple enemy air defenses. Finally, AVs are more agile and more precise in control, and thus they can sustain more aggressive maneuvering that a human could not. Therefore, the author(s) of this quote argue that the AV is functionally better suited for these “dull, dirty, and dangerous” missions.

The first section of this chapter explores the traditional approach to automation design by functional task allocation, which helps define natural boundaries of expertise between humans and automation. The second section of this chapter discusses how this traditional automation design approach has failed to adequately integrate humans and automation into a cooperative system. It then highlights the need for team-centered automation design and its unique applicability to tactical environments.

2.1 Functional Task Allocation

Since the 1950s, many automation designers have relied on a functional allocation of tasks to determine the relationship between man and machines. Designers begin with a set of requirements that are typically high-level performance goals, and they desire to build in automation to the system. By identifying where the human has failed before and confirming that automation is better suited for these particular tasks, a design decision is made to replace the human. This is functional task allocation [63]. If automation can better perform a task (where “better” equates to more precise, faster, and/or less costly), than the human should be replaced.

2.1.1 Automation and Human Capabilities and Limitations

Fitts List

In 1951, Fitts et. al. were tasked in addressing the relationship between humans and automation in the future of air traffic control [11]. As a starting point, they asked the following two question: what can men do better than machines and what can machines do better than men? The answers to these two questions formed the famous Fitts List and the basis of functional task allocation. The comparison between humans and automation still hold true today. The basic strengths and weaknesses of humans have not changed in fifty years. Automation is certainly more mature now than in Fitts’ day, and vision-based control systems, pattern-recognition algorithms, improvements in sensor technology, and adaptive learning systems are all blurring the lines between human and machine strengths. Yet, Fitts List addresses fundamental dichotomies that will for the near future not change.

Humans possess five functional characteristics that elevate them over the machine [11]. First, the amazing auditory and visual acuity sensory functions allow for extremely low stimulus thresholds. For example, a human eye can detect the flare of a match that is lit fifteen miles away on a dark night. A trained human ear is so sensitive that it can almost detect random collisions of molecules of air. Note, though, that artificial sensors allow detection of energy wavelengths outside the human eye’s and ear’s bandwidth. Second, perceptual abilities allow a human to abstract a pattern into long-term memory. For example, a human can recognize an uncle that has not been seen for a few years and has changed from always being clean-cut to a grown beard. The abstraction of the uncle into long-term memory allows rapid retrieval and recognition even if the uncle has grown a beard, lost some weight, etc. Consider also the qualities of squareness, roundness, and triangularity which can be easily understood and recognized even though square faces, round edges, and triangles exist in an abundance of forms. Pattern-recognition routines are still a long way off from achieving this level of perceptual ability. Third, the flexibility of humans enable them to tackle old problems in new ways or to simply improvise. As Fitts notes, “the machine will attempt as many different kinds of solutions as its designer planned for and no more.” Fourth, after attaining a level of situational awareness, humans can selectively recall previous experiences from long-term memory storage and judge

how best to proceed in the current situation. Fifth and finally, inductive reasoning is unique to humans. Of all these strengths, only the sensory functions and judgment and selective recall abilities (long-term memory storage retrieval in combination with an inference engine) appear potentially executable by maturing technology.

In contrast, what can machines do better than humans? First, the processing speed and power of machines greatly exceed human capabilities. The fastest reaction human time from stimulus appearance to input response is 0.1 seconds. An average response time in the cockpit is 1.5 seconds for the pilot. Second, machines can perform routine tasks quicker and more accurately than a human. No human enjoys “busy work.” Third, the computational capabilities of a computer far exceeds that of a human. Fourth, machines maintain more efficient use of working memory. Sometimes humans have difficulty erasing information from working memory (“I cannot get that song out of my head.”) which takes up necessary storage space for other problems. Finally, through partitioning, a computer can perform several simultaneous activities at the same time. Other than for extremely basic functions such as breathing or walking, humans are serial processors. Try taking the inverse of a matrix while carrying on a conversation. Attention has to be continually diverted to the problem, then to the conversation, but not both at the same time. Out of all of these strengths, a human expert may be able to gain greater computational skills and more efficient working memory usage, but only through years of practice. Yet, even then, computational ability would not be close.

The strengths of humans do not lie in manual labor, whether that is the construction worker welding beams together, the high school student taking a pre-calculus test, or the pilot making continuous adjustments to the control inputs to maintain steady, level flight. Automation can perform these sorts of tasks more accurately and more efficiently. The strengths of humans lie in cognitive judgment processes, concept abstraction, and creative, inductive reasoning skills. Before addressing how these strengths contribute to learning tactical knowledge, we discuss one other more recent comparison of human and machine expertise.

Other Comparisons

Table 2.1 displays another listing of the strengths and weaknesses of humans and machines [9]. The top left set of characteristics in the table describe the strengths of human expertise in problem-solving. Humans are creative and adaptive. Humans have the ability to try completely new ways to solve a problem, and they learn from their mistakes and successes to better position themselves for the future. Human perception of the environment is also vastly superior to a machine’s because of sensory experience, allowing complex problems such as pattern recognition to be an unconscious, inherent part of everyday human life. Also, by utilizing parallel strategies and lines of thought, humans can maintain a broad focus or a so-called “big picture” view that helps guide the overall problem-solving process while taking time to solve narrow-focused problems. Humans can hierarchically tackle a problem. Finally, humans can extrapolate their experiences and knowledge to many other areas in life by using common sense.

HUMAN EXPERTISE	MACHINE EXPERTISE
The Good News	The Bad News
Creative	Uninspired
Adaptive	Needs to be Told
Sensory Experience	Symbolic Input
Broad Focus	Narrow Focus
Common-sense Knowledge	Technical Knowledge
The Bad News	The Good News
Perishable	Permanent
Difficult to Transfer	Easy to Transfer
Difficult to Document	Easy to Document
Unpredictable	Consistent
Expensive	Affordable

Table 2.1: Comparison between human and machine.

The bottom right set of characteristics in Table 2.1, describe the strengths of machine expertise in several particular areas. First, machine memory is permanent. In contrast, if humans do not consciously revisit and rework areas of expertise that have not been used over long periods of time, they risk losing that expertise. Second, machine expertise can be easily transferred to other machines with common formatting architectures, such as would be found in mass production, by downloading and uploading. However, it takes a long time for a human expert to teach an apprentice all of his or her expertise. Third, though debugging is extremely painful, in theory, machine expertise is easy to document because it's already recorded in software lines of code and written to hard drives. One of the main weaknesses of humans is the inability to be consciously aware of their own cognitive processes. Human expertise is difficult to document because many times experts declare that they simply acted out of intuition. Fourth, it is true that if a machine is given the exact same set of operating conditions, it will act in the exact same way as before. The problem is that because software has become so complex, there is a transparency issue where the user and even computer programmer can easily become confused as to its behavior. Humans, on the other hand, are more subject to unpredictability, primarily because of their emotions. Finally, machines are economical and affordable, otherwise humans would not have continued to design and build them. It is very expensive, though, to replace a human expert, not only because humans have more life-sustaining needs than machines, but because human expertise, as described above, is very difficult to transfer.

2.1.2 The Right Knowledge

Examining the strengths and weaknesses of humans and automation is necessary in attempting to learn human tactical knowledge and implement that knowledge in AVs. We are not interested in designing AVs that make decisions exactly like humans. It

is not a one-to-one mapping of human cognition to software design. Rather, we want to learn the best knowledge and the best decisions made and implement that subset of human knowledge into AVs. To do that, we have to understand what should be carried over from the human, and where the AV is already superior. For example, the AV's reaction time to a quantifiable sensory input, such as radar, will always be superior to a human's reaction. However, the question of how to interpret that input in a context of dynamic environments and mission objectives is something an AV lacks.

Consider an intelligence, reconnaissance, surveillance (ISR) mission. The task is to search through a section of terrain and report any enemy contacts. Given a terrain geometry, a section of the terrain to be searched, a time limit, vehicle equations of motion, and sensor specifications, a computer can use heuristic searching algorithms to find a solution path that maximizes the amount of terrain seen within constraints. On the other hand, when this section of terrain is only one piece of an entire path to be searched, when there is a very good but still uncertain chance of enemy contacts along this terrain, and when there is a more important piece of terrain still to be reached with time running out, it takes the creative reasoning, adaptability, and broad focus of a human expert to effectively accomplish the ISR mission goals.

This example also highlights the need to learn *strategies* not actions, to understand why the human expert chose their actions, not just the actions themselves. We can observe the human expert utilize creative reasoning to find a new solution path to a problem, and then simply note the actions taken and the environmental variables that were present and encode this information as a new rule. However, this would fall far short of understanding why the new solution path was chosen. If the strategy was known, it could be used as a template or higher-level goal that would help solve future variations of the problem. Learning strategies and not just actions are important for at least three reasons. One, it is impossible for the human expert to participate in a full factorial search of the multi-dimensional environment to create a complete rule set. Two, the continual change of the battlefield requires decision making skills that go beyond preset actions. It also can make any such full factorial effort as irrelevant. Three, understanding strategies and not just actions is the only way to extrapolate lessons learned in a simple simulation environment to higher fidelity exercises, and ultimately real life.

Therefore, it is not enough to observe a human subject matter expert solve a problem which is completely quantifiable. A computer can do that. Remember that tactics were defined as the decisions made, techniques employed, and actions taken to successfully carry out the mission in the face of dynamically changing environments while staying within constraints. Thus, when we seek to learn human tactical knowledge, there must be uncertainty in the scenario, a hierarchy of objectives the human is attempting to accomplish, and an allowance of flexibility and creativity so that the human can learn effective ways to solve each new problem. In fact, the conclusion from the Fitts list section of the 1951 report is as follows: [11] (*italics for purposes of this research*)

In summary then, we can see that the human carries within him some

remarkable powers that cannot yet be duplicated by machines, especially abilities needed to deal with *changing* situations and *unforeseen* problems.

Tactical environments are changing and uncertain and yet human experts operate successfully within them. These, then, are the strengths of humans and where we seek to improve AV capability.

2.2 Team Centered Automation

In Chapter 1, we stated that by improving AV tactical control, we hope to make AVs more predictable, trustworthy, and better understood by their manned counterparts. The above section discussed what aspects of human tactical capabilities can help augment an AVs tactical control. The question still remains of whether this improved tactical control will either compete with current human experts or cooperate and support them.

2.2.1 Unique Demands of Tactical Control Environments

Because we are considering tactical control in high pressure, high risk, high tempo, dynamic environments, there are unique demands placed upon any AV that autonomously makes decisions in these battlefield situations. These demands are due to the unpredictability and the consequences of battlefield operation, which both point to a necessary element of trust between members of a team operating within tactical environments. This need for trust underscores AV design that is also team-centered.

High Tempo and Uncertainty

The tactical knowledge we seek to learn from human subject matter experts is reactive to changing environments. Battlefield operations are the quintessential environment of uncertainty and continuous change with high-stakes outcomes. There is never enough intelligence for any mission because too much is unknown and too much will change [46]. In addition to this unpredictability, there are other sources of noise and error. These are the AV's own noisiness in sensing the state of the environment, the possibility of internal system malfunctions, and the ever-insidious presence of software bugs. Therefore, all these error sources combined with the sheer uncertainty and change of the battlefield guarantees two things. First, there will always be situations in which the AV chooses the wrong decision [54]. Probability theory dictates there will always be some chance of failure or loss. What was a right decision initially may quickly become the wrong one during its execution. Second, there will always be situations in which the AV has no experience. Expert systems only display expertise for those conditions in which they have been programmed. AVs cannot be required to display expert decision making for a set of conditions not included in their rule base. Artificial learning and adaptive systems seek to address this problem, but as Canning says, "machines lack knowledge of the world *context* that they are in, something that people learn from birth" [9]. Without a world context that exists beyond a rule

base, the AV must rely on a master switch return-to-base or hover-and-wait function that turns on in case of confusion. Therefore, the probability of a wrong decision or the failure to make any decision forces AV designers to consider the consequences of such outcomes. Wrong decisions can be tolerated to the degree of the negative cost incurred.

High Pressure and High Risk

The life and death nature of battlefield operations also places unique demands on reactive decision making AVs. Both the Law of Armed Conflict (LOAC) as well as the Geneva Conventions provide rules and guidelines that attempt to minimize the casualties to civilians during warfare [10]. Unfortunately, recent conflicts reveal how enemies try to exploit friendly forces attempts to abide by LOAC by intentionally placing civilians in harm's way to gain tactical advantage. Consider, for example, mass uprisings against United States (U.S.) peace-keeping forces in Somalia where insurgents either hid between women or children to fire at U.S. ground forces or simply placed semi-automatic weapons into the hands of their children to fire at U.S. ground forces [5]. Even if AVs possess the tactical decision making capabilities to successfully provide ground support for friendly forces, the consequences of missing targets and hitting civilians or even friendly forces are a tremendous hurdle. Note, however, that there could be situations in which fully autonomous reactive decision making is appropriate. For example, a new feature being considered for the Lockheed Martin Joint Strike Fighter's flight control system is the auto-eject for the U.S. Navy variant [49]. Pilots do not have the reaction time to eject while being catapulted from the deck of a U.S. aircraft carrier if there is a major failure during takeoff. Therefore, auto-eject is appropriate to save the pilot's life if it can be proven robust and reliable. Unfortunately, no matter how many metrics of reliability the AV passes, no matter how great the end-to-end testing program, accidents will occur [9]. If the consequences of incorrect decisions are great, then Parasuraman et. al. recommended the following: [54]

Giving the pilot the opportunity to review the decision choice and forcing a conscious overt action, provides an "error-trapping" mechanism that can guard against mindless acquiescence in computer-generated solutions that are not contextually appropriate.

For instance, if tactical decision making on the battlefield includes the use of weapons, some sort of "error-trapping" gate will have to be in place. Note that any level of "error-trapping" in AVs is appropriate, regardless of whether it's armed or not. The point of this section is to underscore that arming an AV provides a greater need for reliability. Therefore, though the AV could be the weapon delivery platform, the consequences of battlefield operations require a team effort centered on trust and reliability to ensure the minimization of civilian or friendly force losses.

Lack of Trust

The attitudes of current pilots highlight the need for trust in tactical situations. Morales and Cummings investigated how pilots responded to the use of AVs as “wingmen,” where wingmen are the elements in an aircraft formation [46]. In their study, pilots could vary the level of control given to and task three AV wingmen through a cockpit interface. The pilots who participated included four A-10 pilots, two F-16 pilots, two crewmembers in a multi-crew AC-130 cockpit, and two ground operators of the Predator AV. The scenarios included target acquisition, AV assignments, battle damage assessment, and secondary strikes. The research objective was to answer the following three questions:

1. What levels of pilot control and AV/human interaction do pilots think are appropriate?
2. What is the relative importance of different display characteristics?
3. Should AVs play the role of a “wingman?”

Analysis indicated that pilots generally agreed on the following points. AVs should be allowed to defend themselves with complete autonomy and should automatically collect images of targets and transfer that data to manned assets. On the other hand, AVs should not be allowed to designate a target and should not perform any kind of battle damage assessment. Interestingly, the AC-130 crewmembers and Predator operators were more open to AVs performing combat offensive missions than the A-10 and F-16 pilots who actually train for those operations. The two F-16 pilots declared that AVs should never operate in the same airspace as manned tactical aircraft. For example, one of those pilots cited his own eye-witness account of a software malfunction causing a Predator to drift into the path of a group of fighters, which almost resulted in a mid-air collision. One A-10 pilot described his relationship with his human wingman as, “one of trust and loyalty.” They trained together, worked together, and fought together, and therefore a AV could never replace a human wingman.

Though the typical response from an AV designer is to brush off the pilots’ remarks as arrogance, bias, and fear of being replaced, the force of the above comments emphasizes how absolutely crucial is the need for trust. The F-16 pilot may have seen other close mid-air collisions between manned aircraft, but the pilot would still rather trust a human to make that mistake than an AV. Why? The A-10 pilot viewed the human as a team member and thus, trusted him more. Therefore, tactical decision making places unique demands on the decision maker because of the high risk and dynamic battlefield environment. The life-and-death consequences of wrong decisions create a need for all decision making assets to be on the same team. That team must exhibit a culture of trust and loyalty. What, then, is team-centered automation? To answer this question we first define automation more precisely, discuss human-centered automation design, briefly mention the lessons learned when automation has failed to be human-centered, and finally review what characteristics define a team.

2.2.2 Definition of Automation

There are many different definitions of automation in the literature. Researchers Raja Parasuraman, Thomas Sheridan, and Christopher Wickens define automation as “a device or system that accomplishes (partially or fully) a function that was previously, or conceivably could be, carried out (partially or fully) by a human operator” [53]. Dr. Charles Billings, Retired Chief Scientist of the National Aeronautics and Space Administration (NASA) Ames research center in Silicon Valley, defines automation as “a tool, or resource, that the human operator can use to perform some task that would be difficult or impossible without machine aiding” [56]. The Autonomy Levels for Unmanned Systems (ALFUS) working group at the National Institute of Standards (NIST) defines autonomy in two ways [34]. Autonomy is:

- (A) The condition or quality of being self-governing.
- (B) An unmanned system’s own ability of sensing, perceiving, analyzing, communicating, planning, decision-making, and acting, to achieve its goals as assigned by its human operator(s) through designed human-robot interaction (HRI). Autonomy is characterized into levels by factors including mission complexity, environmental difficulty, and level of HRI to accomplish the mission.

The definition from Parasuraman et. al. emphasizes that automation is not “all or none” but can be conceived of as ranging across a continuum of levels. Dr. Billings emphasizes a human-centered definition in that the role of automation is to be a tool or resource that aides the human operator. The ALFUS working group’s definition from NIST attempts to capture the roles of intelligence and capabilities of the AV while characterizing the human operator as assigning goals to the AV. The ALFUS working group also recognizes that autonomy should be characterized by levels, including the level of interface between the human and the robot. Therefore, automation is a tool for the human operator, can be characterized by levels, and exists within a complementary system that includes both the human and the automation.

Both the U.S. Air Force and Army have proposed classification levels of autonomy [9]. The Air Force Research Laboratory has defined ten levels of autonomous control, as depicted in Table 2.2. At level 4 autonomy of onboard route replan, reactive decision making is needed. How else will the AV know when to replan its intended route and what the new route should be unless there is a reactive decision making process onboard? By the definition in this thesis, onboard route replan represents a form of tactical control. However, note that AFRL does not explicitly specify tactical control unless existing within a group. This thesis does not address group tactical environments, but rather describes experiments which extracted single vehicle tactics. Group tactical control is left for future work. However, the explicit grouping of manned and unmanned assets in tactical control until level 10 autonomy is reached serves to emphasize the importance of team-centeredness in tactical environments.

Table 2.3 displays the levels of autonomous behavior for unmanned ground vehicles (UGVs) as defined by the U.S. Army and its Future Combat Systems (FCS) initiative. Here, reactive decision making occurs at level 7 autonomy with auto ne-

LEVEL	DEFINITION
1	Remotely Guided
2	Real Time Health/Diagnosis
3	Adapt to Failures and Flight Conditions
4	Onboard Route Replan
5	Group Coordination
6	Group Tactical Replan
7	Group Tactical Goals
8	Distributed Control
9	Group Strategic Goals
10	Fully Autonomous Swarms

Table 2.2: Air Force Research Laboratory levels of autonomous control

LEVEL	DEFINITION
1	Remote Control/Tele-operation
2	Mission and Task Planning
3	Improved Route Following on Paved Roads
4	Unimproved Route Following Dirt Roads
5	Off-Route Mobility No Roads
6	Obstacle Detection and Alert Operator (> 0.2 meter obstacles)
7	Obstacle Detection and Auto Negotiation (> 0.2 meter obstacles)
8	Tactical Payload Mission Behaviors
9	Cooperative Behaviors with Manned and Unmanned Systems
10	Reactive Intelligent Tactical Behaviors

Table 2.3: Army Future Combat Systems levels of autonomous behavior

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

Table 2.4: Levels of Automation of Decision and Action Selection

gotiation of obstacles. Again, it is important to note that a cooperative team of manned and unmanned systems needs to exist before reaching level 10 autonomy of reactive intelligent tactical behaviors. Therefore, both the Air Force and the Army characterize tactical control levels for AV only in co-existence with human assets. It is not enough to provide AVs with a greater level of expertise and then expect to replace the human. By definition, both in academia and in the military services, autonomous tactical decision making only takes place in a system of manned and unmanned systems. Learning the best tactical knowledge is important, but it cannot occur outside of a team-centered context and be indiscriminately applied to an AV. This is why the human expert must be so directly involved in the design process.

2.2.3 Human-Centered Automation Design

The traditional context of a complementary human and autonomous decision making system casts the human in the role of supervisor/operator and the automation in the role of an aid. In this system, a design decision must be made about how much authority should be given to the automation to make its own decisions. Sheridan proposed ten levels to describe the various levels of interaction between the human and machine for decision and action selection [62], as depicted in Table 2.4. An example of this system context is the air traffic controller (ATC). It is the responsibility of the ATC to direct the flow of traffic in and out of airports to ensure efficiency and the safety of all involved. In doing so, the ATC has the authority to issue headings, velocities, and holding patterns to all aircraft. In major airports, this is an extremely demanding task, and the potential benefits of workload reduction by automation are tremendous. The focus of automation decision aids has been the ability to predict and project the current tracks of aircraft some time into the future so that the ATC can resolve possible collisions between aircraft. Suppose the computer calculates a collision course between two aircraft. At level 3 automation, the automated decision aid recommends a few courses of action to resolve the projected situation. The ATC chooses one and relays the information to the pilots. At level 7 automation, the decision aid automatically uplinks new course and heading information to the aircrafts' onboard computers and then informs the ATC of the action. In both scenarios, the

ATC is the ultimate supervisor and operator of the system. The difference is in how much authority has been given to the automation.

As a parallel to the ATC, consider the battlefield commander waiting to move in a company of troops by helicopter to landing zones. That battlefield commander wants to ensure that a specific proposed air corridor through which the troop-carrying helicopters will fly is free from enemy contact. The commander is the supervisor for the improved AV who will perform the reconnaissance. In this case, as the AV performs a reconnaissance mission, the interaction between the AV and the battlefield commander could easily resemble that between the ATC and the decision aid. For example, the AV, trained by human expertise, encounters a pop-up threat, reacts, and now prepares to engage the target, *if* the battlefield commander approves. From Table 2.4, this would be level 5 automation. The question, then, that the battlefield commander, the ATC, and most importantly the automation designer must answer is how much decision making autonomy is appropriate.

Figure 2-1 displays one proposed method of designing human-centered automation [54]. In this flow chart, it is assumed that the human brain can be viewed as an extremely sophisticated information processing system. Then the mapping of input cues to output actions as experienced by the human can be represented by the following four stages: sensory processing, perception/working memory, decision making, and response selection [87]. Admittedly, these four stages are an oversimplification at this point, and they will be more fully dealt with in Chapter 3. Parasuraman et. al. proposed four classes of functions that are roughly equivalent to these four stages of human information processing [54]. These are information acquisition, information analysis, decision and action selection, and action implementation. For each function class there exists levels to describe how fully automated the system is within that particular class, as in Table 2.4. Therefore, a specific unmanned system can be described using this method by its level of automation along all four dimensions. These four dimensions are pictured in the flow chart as the four parallel blocks of acquisition, analysis, decision, and action.

In answer to the top question of Figure 2-1 of what should be automated, this research seeks to automate tactical control in AVs. Tactical control falls into the branch of decision automation. Next, from the ten levels in Table 2.4, the human user, such as the battlefield commander, would help decide an appropriate level of automation decision making capability. Say level 6 automation was initially chosen where the commander has a limited amount of time to veto any action automatically chosen by the AV. Then the primary evaluative criteria of whether level 6 automation is appropriate or not is the human performance consequences of that design choice. The secondary criteria is to evaluate the level of automation reliability required, the costs of action, etc. These criteria force an iterative refining of the automation. Parasuraman et al. emphasize that Figure 2-1 is a framework to help provide guidelines for automation design. They also recognize that this framework may be more useful in helping to define upper and lower bounds of automation rather than a specific level.

The foundation of human-centered design captured by the flow chart is that an appropriate level of autonomy is chosen primarily on how it affects the human that

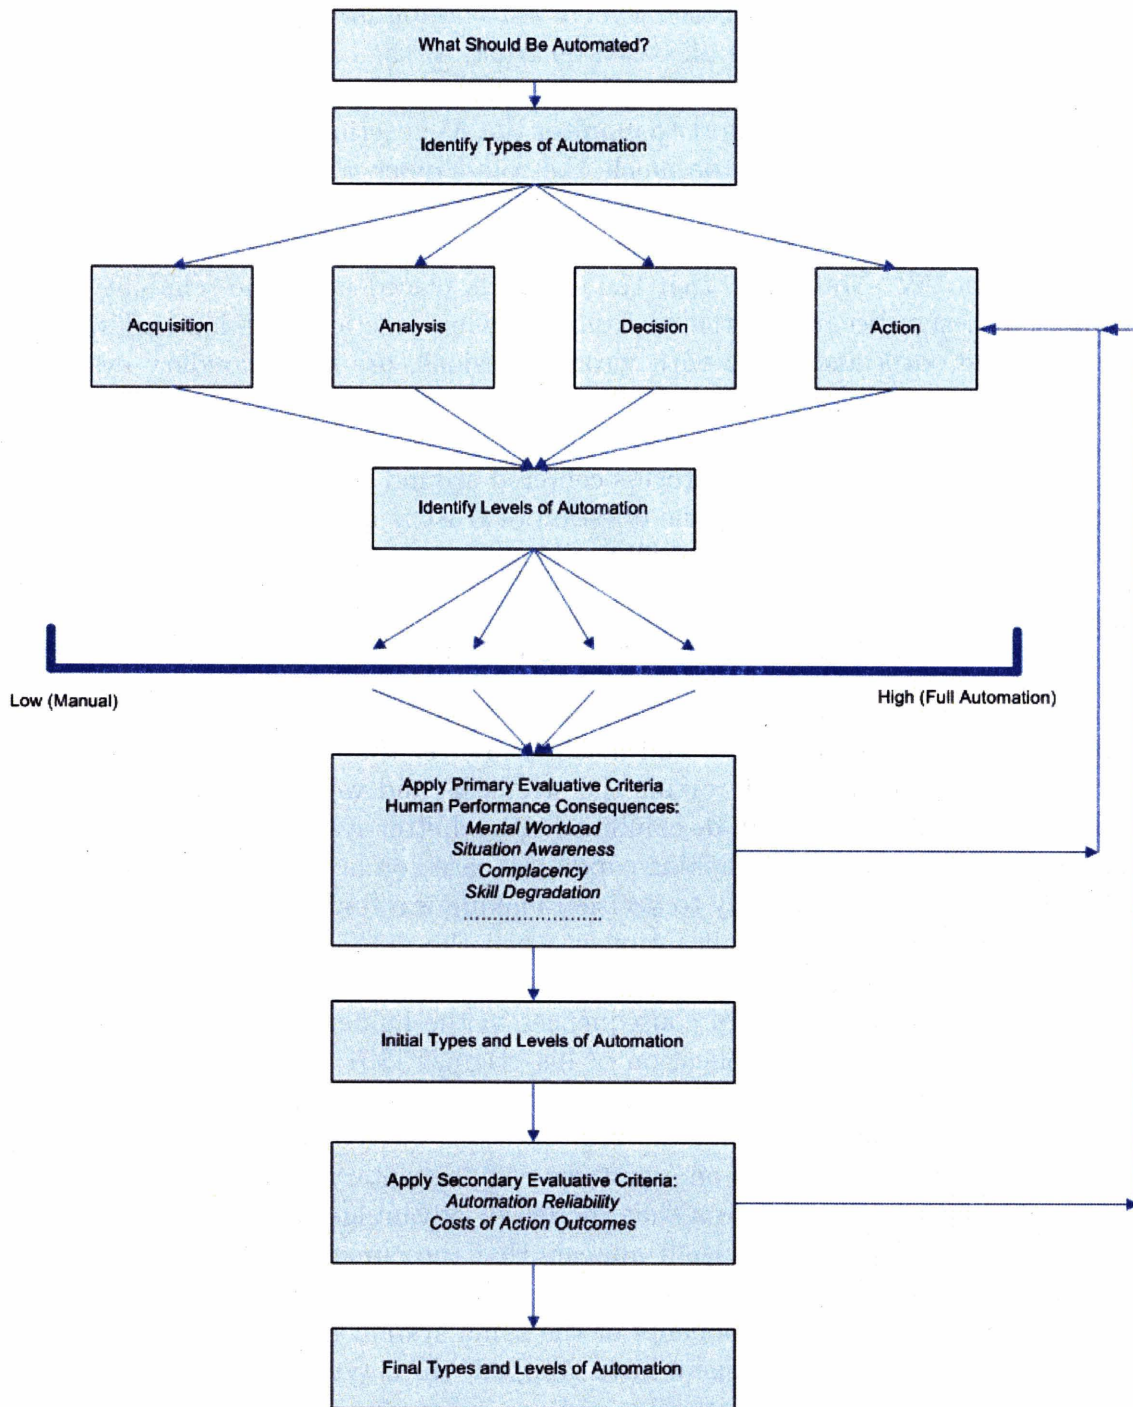


Figure 2-1: Flow chart showing application of the model of types and levels of automation [54].

is still a part of the system. Thus, the designers must ask what are the consequences to the battlefield commander's mental workload, situation awareness, complacency, skill degradation, etc., at the chosen level of automation for tactical decision making. Interestingly, these criteria are all based on trust. If the commander does not trust the AV, the commander will have a higher mental workload because the commander will devote a lot of cognitive effort to monitor the AV's actions. On the other hand, if the commander trusts the AV too much, the commander might lose situation awareness by not carefully monitoring the AV's actions. The commander may also become complacent and not react timely to the AV's decisions and actions if the commander over-trusts the AV. Note then, that the demands placed upon the reliability of the automation design because of the environment where tactical decision making takes place and the costs associated with wrong decisions, are the secondary evaluative criteria in human-centered design. Those demands pointed to the need for trust between all decision making assets, both manned and unmanned, in tactical situations. Therefore, beginning the design process centered around the human, trust becomes a primary criteria. Unfortunately, the question of trust is not just whether the human operator will use the AV, but also will the human operator use the AV correctly.

2.2.4 Lessons Learned from Improper Automation

A functional task allocation, as discussed in Section 2.1 asks the question what can humans do better than machines and machines better than humans? The benefit of asking this question is that it reveals the strengths and weaknesses of humans and automation so that the two can be combined into a better system. If we are seeking to improve autonomous decision making capability based on human performance, knowing what human strengths apply to decision making is a fundamental step. However, the failure in functional task allocation is when the automation designer identifies weaknesses, improves the autonomy, replaces the human as much as possible, and never stops to consider the final consequences to the human. In fact, Dr. Fitts admitted ten years after the publication of his original 1951 list of the superiority of man over machine and vice versa was misleading [23]. He declared that he had fallen into "a trap" with that list, and the real question was not of allocating functions based on superiority but based on a systems complementary approach [63].

From Table 2.1, it can be seen that the major advantage of automation is not to replace the human simply because it appears that the human is limited or expensive to maintain. Rather the automation should be complementary to the human [9]. However, complementary components of the same system require an interface, and this is where automation designers have stopped. In a typical manned/unmanned system, both the human and machine aid are subsystems, and the display is the interface. Furthermore, it is well known in design that next to requirements, designing the proper interfaces between complementary subsystems is extremely critical [12]. Consider, for example, an interface failure when the O-rings in the space shuttle Challenger's solid rocket boosters failed due to extreme cold weather. Hot gas leaked past the O-rings, the failed interface, and ignited the entire stack [13]. Though this is an extreme example, it shows that interfaces cannot be taken for granted. Even

machine aiding and the simple presentation of several options to the human requires considerable thought on how to display the options to the human user [87]. As a case in point, a whole academic field of human factors and ergonomics has arisen to research how to create a complementary system composed of both humans and automation because it is such a difficult problem [1].

In the human factors literature, there is a term for when designers automate, partially or fully, functions within a system without due consideration of the consequences to the human operator of that system. That term is “automation abuse,” and it has been a direct cause of serious incidents and accidents [53].

During the 1970s and early 1980s . . . the concept of automating as much as possible was considered appropriate. The expected benefit was a reduction in pilot workload and increased safety . . . Although many of these benefits have been realized, serious questions have arisen and incidents/accidents have occurred which question the underlying assumptions that a maximum available automation is ALWAYS appropriate or that we understand how to design automated systems so that they are fully compatible with the capabilities and limitations of the humans in the system. [56]

This quote from the Air Transport Association of America (ATA) Flight Systems Integration Committee in 1989 underscores how automation abuse caused an entire system of pilots, engineers, designers, and safety controllers to step back and think through the issue of automation design. Automating everything as much as possible was found to be no longer appropriate, and the question was then what should even be automated.

There are two ironies associated with automation abuse in which the designer wishes to replace the human operator with automation because of the human’s tendency to make errors [1]. First, the designer has now simply replaced the human operator with himself or herself. Now the system is prone to the errors in design, which are still human in origin. Second, the designer who tries to eliminate the human operator still leaves the operator to perform certain tasks which the designer cannot think how to automate. More and more research has proven that automation does not supplant human activity. Rather, it changes the nature of that activity, often in ways unforeseen and unintended by the designer, which has led to several problems in real world applications [1, 53, 54, 56, 88, 90].

As briefly discussed with the human-centered flow chart, human operators can either over-trust or under-trust the automation. Over-trust or over-reliance on automation has been termed “automation misuse.” Numerous accidents have occurred due to misuse, such as the crash of Eastern Flight 401 in the Florida Everglades, when the crew failed to notice that the autopilot had been disengaged [53]. They were not monitoring the aircraft’s altitude while diagnosing a possible problem with the landing gear. Not only does over-trust lead to reduced situation awareness, it also leads to skill degradation. Pilots who tend to always let the automation fly the airplane lose some measure of their skills as a pilot. Only in emergency situations or very complex tasks do most pilots tend to disengage the autopilot. Ironically, these are the times when the pilot’s skills should be the sharpest, but they have degraded

because of misuse. On the other hand, some pilots begin to distrust their own skills due to over-reliance on automation, and they rely on the autopilot for safety. This caused one pilot to crash short of the runway at Columbus, Ohio, in 1994 when he relied heavily on the autopilot to land during a nighttime snowstorm [53]. Over-reliance may also result in the human operator becoming complacent in monitoring the automation. Complacency becomes life-threatening when combined with automation that fails silently. For example, if the failure of an autopilot results in large, unexpected banks, the automation failure is obvious. On the other hand, if the autopilot fails silently and the airplane begins to roll ever so slightly due to a slightly unstable roll mode, the pilots may not recognize the failure due to complacency. This situation occurred in the 1985 China Airlines incident when the wings were almost vertical before corrective action was taken [53]. Over-trusting the automation, then, leads to reduced situation awareness, skill degradation, and complacency.

Under-trust or under-reliance is termed “automation disuse.” Automation disuse is usually related to the false alarm problem. For example, early versions of the Ground Proximity Warning System produced so many false alarms that pilots stopped trusting its warnings. Operators may also use “creative disablement” to turn off the alarming system, such as the Conrail train accident in 1987 [53]. Investigators found that the loud buzzer in the train cab that alerts of high speeds had been taped over. Under-trust also results in increased mental workload. Because the pilot does not trust the automated solution, the pilot must spend some significant “cognitive overhead” to create his or her own solution, then compare the two, and finally choose one. Under-trusting the automation leads to disuse of the automation or increased mental workload.

Over-reliance and under-reliance represent issues of trust between the human and the automation. Automation abuse also results in other problems. First, automation surprises result from the complexity of modern systems and algorithms. Because closed-loop control is so tightly optimized for fuel efficiency, there are many times when the pilots get very confused as to the purpose behind the airplane’s behavior. It is an issue of transparency versus opacity. Second, “clumsy automation” is a term coined by Wiener that describes automation that reduces workload when the workload demands are already low and increases them when attention and resources are needed elsewhere. An example is the flight management system (FMS) that performs waypoint following [56]. During the transit flight phase between destinations, the FMS reduces workload when it’s already low. However, during descent when the co-pilot should be scanning for other aircraft, the co-pilot has to spend time reprogramming the FMS to change the plane’s descent path. A third problem is silent failure of automation, as discussed above. Graceful degradation of performance is a human strength over machine [1]. However, it is not something to be carried over to automation. Automation should fail obviously, especially when that failure has tremendous negative costs associated with it.

As a final summary chart, Figure 2-2 displays the theoretical and experimentally verified human factors variables that picture why a human either uses automation or not [53]. The dotted lines represent theoretical relationships or relationships that depend on the system in question. The solid lines represent relationships supported

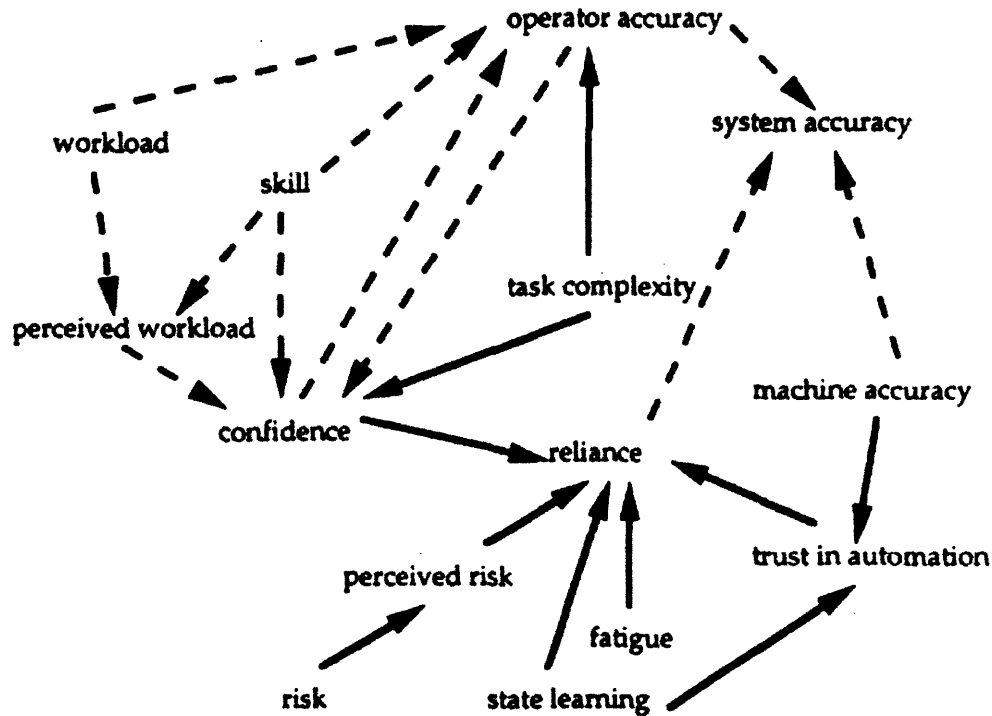


Figure 2-2: Automation usage based on human factors [53].

by experimental data. The important note to take away from Figure 2-2 is that the factors involved in a human choosing to use automation are complex and not completely understood. Therefore, it is very difficult to predict how a human will choose to use automation or not. However, this is not to discourage a human-centered automation approach. In fact, note the importance of reliance and trust in automation as *experimentally* verified relationships. A human-centered automation approach is crucial, both to avoid the myriad of lessons learned as described above as well as to increase the acceptance of automation in the field.

2.2.5 Characteristics of a Team

In learning tactical knowledge from human experts, the human in the system is not a supervisor, operator, or commander, but rather a team player. This is a subtle distinction. In a reconnaissance mission, the performance and capability of the AV to effectively search, and if necessary to engage enemy contacts, directly affects the company of troops that are waiting to pass through the scouted terrain. In this scenario, we are not designing the AV with the battlefield commander/supervisor exclusively in mind. Rather, we are designing the AV with the company of troops in mind that desire safe passage to their destination. Those troops are hoping the AV was trained properly and knows how to prioritize its searching/evading/engaging efforts in order to strengthen the overall team position. This is not going to happen

unless we can completely understand the human previously performing this task with that team mindset.

Human-centered automation for this research focuses on the human expert from whom we learn tactical knowledge for two reasons. First, the human expert is our primary source of designing the tactical playbook. The AV is not designed to make up for human errors or weakness. We are admitting the AV is lacking and that much can be learned from the human expert. Secondly, the human expert can trust the AV more because he or she has trained it. In the same way that the flight instructor teaches the student how to fly in formation, which is by necessity a cohesive, team-centered unit, the human expert trains the AV to make correct decisions so that the team will benefit. Furthermore, the human expert is not training their replacement. Even when AVs with tactical decision making capabilities are fielded, human subject matter experts such as fighter and attack pilots will still be very much needed. Thus, the human expert has a vested interest in making sure the AV works toward team goals of which the human expert is still a part.

Does this negate the flowchart in Figure 2-1 or lessons learned in the previous section, which were presented with the human assumed to be the operator of the system? We argue that it does not. Because of the consequences of wrong decisions in battlefield environments, there has to be a human supervisor of any decision making AV. Therefore, the traditional context of a system composed of a human operator and machine aider still holds. Furthermore, the lessons learned of reduced situation awareness, complacency, increased mental workload, automation surprises, clumsy automation, and silent failures are even more important in the battlefield context because of those same huge negative costs of wrong decisions. Therefore, the lessons learned still hold. However, we argue that there should be a second set of primary evaluative criteria in Figure 2-1 of Team Performance Consequences. How is team performance in a battlefield context measured? The chain of command issues a set of orders that filters down through the ranks until everyone understands their mission specific objectives. Also, those orders describe the desired end-state of the commanders so that soldiers can grasp how their mission specific objectives contribute to the high-level goal [81]. Therefore, one proposed way of evaluating Team Performance Consequences is to measure how well the AV meets the specific mission objectives that the battlefield commander desires.

For reference, there are five characteristics of good team players [56]. First, team players are reliable. If the AV has been given a list of mission objectives, how reliably are those objectives met? Moreover, knowing that uncertainty characterizes the tactical environment, how robust and flexible is the AV to meet those mission objectives in unexpected situations? Second, team players communicate effectively with each other. Does the AV's performance reveal automation surprises and silent failures? Is the complexity of the decision making algorithm so dense that the AV cannot effectively relay to other team players its intentions and actions? Third, team players coordinate activities with each other. Do the choices made by the AV impede the performance of other team players? Note that coordination can only occur in dynamically changing environments if effective communication is present. Fourth, team players monitor each other in order to "back each other up." Does the AV have

the capability to monitor the other team players? Does the AV have the intelligence to shift its priorities if they need the AV's help? Or does the AV simply perform its nominal mission until a help signal has been received? Fifth and finally, team players are guided by a coach. Does the AV respond appropriately to changes issued by the battlefield commander? Team players, then, are characterized by reliability, flexibility, effective communication, coordination, monitoring, re-tasking, and re-prioritizing all for the sake of improving the team's position to meet the team's goals. This list of team player characteristics is a daunting yet necessary summary of how AVs must be designed in order to integrate properly into today's battlefield.

2.3 Implications and Closing Thoughts

This research proposes that by explicitly designing the automation based on human inspiration, automation will be more acceptable as a team player with the human users. The reasoning is as follows: the human expert has essentially trained the automation in the best interest of the team. Therefore, human-centered automation as described in this research naturally results in team-centered automation. After discussing automation abuse, misuse, and disuse, Parasuraman's and Riley's first conclusion is that, "better operator knowledge of how the automation works results in more appropriate use of automation" [53]. There will always be a period of time for new AVs to prove themselves reliable to domain experts. One of this research's goals is to encourage human acceptance of AVs by making AVs more predictable, trustworthy, and better understood by their manned counterparts through a design effort centered on human inspiration. Only then will the other team members have a better knowledge of how the automation works and be more open to its usage in the field.

Yet, there is at least one word of caution. Because the second step in the two-stage tactical control paradigm (see Section 1.4) is to optimize the AV's behavior for the specific mission, this will result in atypical reactive decisions as perceived by the human. This will have to be addressed in training. This will also provide more emphasis on the team player characteristic of AVs communicating effectively with manned assets. Again, the limitation is that we learn from human experts who naturally sense the state environment differently than AVs. Therefore, what is a right action for the human, may not be right for the AV.

In conclusion, one of the failures of automation designers in the past was to identify the human weakness, build the right automation, and replace the human. Our design method is to identify the human strength and improve the AV's decision making capabilities. This research is an initial step in how to use human-inspired tactics to ultimately achieve level 10 autonomy as described in Tables 2.2, 2.3, and 2.4. Furthermore, we believe that centering the automation design around the human expert naturally follows the human-centered automation design approach in Figure 2-1. This thesis addresses the decision automation branch in the flow chart and proposes the addition of another set of secondary criteria termed Team Performance Consequences. We only seek to improve AV tactical decision making in a way that

strengthens the position of the entire team. It's a fine line, but the motivation will drive the reliability, trust, and acceptance of the system.

Chapter 3

Human Expert Performance and Cognitive Modeling Efforts

Accurate and reliable observation and interpretation of a human expert's tactical knowledge requires a decision making framework. Though interpretation implies that this is, at least, a partially subjective procedure, it is a necessary one if autonomous vehicles (AVs) are to make tactical decisions in future battlefields, as described in Chapter 1. To review, Chapter 2 discussed the interaction between expertise and reliability in tactical environments composed of both humans and AVs. There were two major conclusions in Chapter 2. First, analyzing the functional strengths and weaknesses of humans and automation helped determine how to design experimental scenarios to exploit human strengths in tactical decision making. Second, the lessons learned from indiscriminate application of automation in the past and the unique challenges of real world tactical environments underscored the necessity that the design approach be team-centered. The five characteristics of reliability, effective communication, coordination, monitoring, and being guided by a coach describe what it means to be part of a team, and the smartest tactical AV will not be trusted by humans and thus not integrated into the battlefield if it fails to behave as a team player. Now, in Chapter 3, we present the underlying mechanisms of human expertise. Understanding these mechanisms is crucial in making objective interpretations of human decision making. We propose that there are three levels in learning human tactics - actions, strategies, and cognitive mechanisms. These correspond to answering the three questions of what, why, and how. We observe the "what," and we desire to know the "why." Therefore, we must also understand the "how" of human decision making by understanding human cognition and decision making frameworks.

Before beginning the discussion on cognitive theories, we wish to present one other prefatory note on the specifics of this chapter. Cognitive science is an extremely diverse field, and the following sections only touch upon a few concepts. Yet, the rise of human factors research in response to the greater reliance on automation has brought the fields of engineering and psychology closer. This chapter, then, aims to familiarize the reader with some of the main theories. However, we explicitly state that only a few of the concepts presented in this chapter were actually applied in the experimental method of learning human-inspired tactics. Namely, these are the following decision

making frameworks: Recognition-Primed Decision model, Generic Error Modeling System, Belief-Desire-Intent model. We also touch upon human interpretation of probability and some decision heuristics and biases in decision making that appeared in the experimental results. Every section in this chapter will finish by summarizing its specific contribution to understanding human decision making so that follow-on work in human-inspired tactics can be even more grounded in objective theory.

3.1 Components of Cognitive Model

To begin understanding the underlying mechanisms of decision making, it is important to first discuss the underlying cognitive structure, namely long-term and working memories. Human experts make tactical decisions by drawing upon past experiences in long-term memory and combining these lessons learned with the present state in working memory to form a set of alternative choices for action.

3.1.1 Long-Term Memory

A human's long-term memory is crucial to living life efficiently. It is in the long-term memory where intuitive, every day actions, such as the motor skills necessary to brush teeth, are embedded. If there were no long-term memory, the ability to brush teeth would have to be relearned every day (or twice a day depending on hygiene). The organized structure, large capacity, and retrieval mechanisms of long-term memory all contribute to how humans make decisions.

Storage Structure

Experiments in accessing long-term memory have continued to confirm the very interesting fact that the storage of long-term memory is not random. In fact, long-term memory storage is highly structured [89]. In 1996, Lipshitz and Bar Ilan analyzed retrospective case reports of low-to-middle tiered managers and their success and failure in problem solving in the work place. Lipshitz and Pras in 2000 questioned these findings because, "it is not clear if their findings pertain to how problems are actually solved or to a cognitive schema that drives the reconstruction of problem-solving processes from long-term memory" [39]. Through a series of experiments in which subjects were asked to think aloud as they solved one well-defined and one ill-defined problem, Lipshitz and Pras verified the existence of a reconstruction that had occurred in long-term memory. They found that "elements in a story that are in a purposeful (in-order-to) relation, such as consecutive elements in an action plan, are remembered better than elements not related in this fashion." Long-term memory storage is structured because the human cognition can more efficiently recall particular facts about past events when the information exists in an organized fashion. This structure is termed *cognitive schema*.

In 1932, Bartlett first proposed that the knowledge embedded in long-term memory were in the form of schema [2]. Definitions of schema emphasize that it is a

structure or model of data or objects in a database. Bartlett found that humans tended to recall memories in more organized, meaningful, and systematic ways than the actual occurrences of them. Humans tended to not remember odd or uncommon details so that the memory's retrieval conformed more to the person's present expectations. For example, consider parents who are recalling the mischievous behavior of their young adult son or daughter when he or she was a child. How often does the young adult disagree with the parents' stories saying, "I never did that! I was never that bad!"? Bartlett proposed that humans were unconsciously attempting to organize those memories in knowledge structures. In his famous phrase, he called this an "effort after meaning." Bartlett defined schema as the following:

[Schema is] an active organization of past reactions, or of past experiences, which must always be supposed to be operating in any well-adapted organic response. That is, whenever there is any order or regularity of behavior, a particular response is possible only because it is related to other similar responses which have been serially organized, yet which operate, not simply as individual members coming one after another, but as a unitary mass. [2]

Note first that long-term memory composed of active knowledge structures rather than passive images reconstructs past experiences rather than reproduces them. Furthermore, though schema is "serially organized," retrieval of knowledge is not a serial search. As will be discussed in the next section, long-term memory retrieval is very rapid. If it was activated by a serial search, the time taken to retrieve a particular knowledge structure would be proportional to the total number of knowledge structures encoded as schemata. Rather, schema operates as a "unitary mass," which allows quick retrieval. Finally, note that knowledge storage is associational in long-term memory. Certain sensory inputs evoke similar responses already stored in long-term memory. Retrieval of knowledge highlights this notion of associativity. Therefore, schemata are unconscious, active mental structures composed of organized past experiences.

Research in the last half of the twentieth century has contributed further concepts to this long-term memory storage structure called schema [57]. First, schemata are high-level knowledge structures. Consider a human subject presented with a picture of a typical living room for only a brief amount of time and then asked to describe all features of the room. If the view of the room contained a wall clock and if the human subject was pressed to describe the clock, the human subject would probably be prone to say the clock had hands. This is because a high-level schema exists that contains knowledge of a prototypical living room with a clock that has hands. Second, each schema accepts only specific information or data. A schema can be considered, then, an "expert" in whatever field of information it requires. If these informational "slots" are not being filled by present inputs, they take on default values from previous experiences. Therefore, when sensory inputs to humans trigger memory recall but only provide partial information, the default values of the past enable humans to infer, either rightly or wrongly, about the present. Finally, there is no known limit to

the number of schemata that can be stored in long-term memory. It is assumed that there is infinite capacity for the storage of knowledge structures [47].

James Reason describes how three errors can arise from inference involving schemata. First, humans fit data to the wrong schema. Second, in an effort to efficiently use memory recall for present actions, humans fit partial data to the right schema, but do not seek further information to fill in the gaps. Rather, they rely on “best guesses” from past experiences. Third, humans tend to rely more on active, presently-invoked schemata and salient, attention-getting schemata. Reason summarizes the good and the bad of schemata in the following:

The very rapid handling of information characteristic of human cognition is possible because the regularities of the world, as well as our routine dealings with them, have been represented internally as schemata. The price we pay for this largely automatic processing of information is that perceptions, memories, thoughts, and actions have a tendency to err in the direction of the familiar and expected. [57]

The tendency for information processing to “err in the direction of the familiar and expected” gives rise to predictable biases which will be discussed later on in this chapter.

Storage Retrieval

Long-term memory retrieval typically occurs very rapidly and is associational both in terms of similarity and frequency. Experiments show that the time to read a knowledge structure in long-term memory is of the order of hundreds of milliseconds [47]. This time is derived from presenting a stimulus to a subject and requiring the subject to respond to with some sort of verbal description, such as naming the color presented. The subject must first recognize the stimulus, which is equivalent to reading a knowledge structure in long-term memory, and then respond appropriately. The entire recall and react time is typically half a second to one second, and by having a good idea of human reaction time, the reading of long-term memory can be backed out to the order of hundreds of milliseconds or less than half a second.

Note that the terms *recognize* and *recall* are used interchangeably above to describe reading knowledge structures from long-term memory. These terms are actually differentiated in literature [89]. To recall is to verbalize knowledge in the head, such as recalling a home address. To recognize is to verbalize knowledge in the world, such as recognizing the sound of an ambulance siren. When asked about a particular event, a human may not be able to recall certain facts, but once presented with the information, a human quickly recognizes it. (“I can’t remember her name, but I would know the face.”)

The association of present inputs to stored information in long-term memory is a combination of similarity-matching and frequency-gambling. The idea of similarity-matching is simply that certain cues active the retrieval of specific information. However, the set of present cues do not typically match a set of schemata completely or

perfectly. Therefore, frequency-gambling and inference exist so that stored information from schemata that are only partially matched can be retrieved and combined to generate some appropriate response. The concept of frequency-gambling is that not only does the schema store specific information, but the schema maintains a trace of its past activation. This is Hintzman's multiple-trace theory [32], where the long-term memory exhibits a sort of "frequency map." The more times a given input is encountered, the larger the "pile" of traces becomes for that particular schema. Furthermore, and more importantly, the more times a given input is encountered, the higher the probability that in the face of partial similarity-matching, long-term memory retrieval will be biased towards this higher frequency content [57]. When a human is faced with a partially novel situation, then, and is unable to generate a novel solution, that human tends to revert back to a somewhat similar situation and proclaims, "I might as well try it. It's worked before."

Inference is the sum total of similarity-matching and frequency-gambling. A set of sensory inputs activates retrieval from a set of knowledge structures based on the similarity and frequency of past encounter of those same cues. The human must then sort through, combine, and manipulate that information to generate a response. Inference occurs in the working memory, which has a limited capacity for storage and manipulation of information. Before moving on to discuss working memory, an example of a long-term memory experiment ties these three concepts of similarity-matching, frequency-gambling, and inference together.

J. Reason asked 126 British psychology students the following question: "Who said (or, more accurately had a sign on his desk saying), 'The buck stops here'?" The answer is President Harry S. Truman. Now the word "buck" in this context came from the slang phrase "pass the buck," which meant to hand over responsibility to another. How would British psychology students, assuming they could not associate the phrase directly with President Truman, infer the identity of the speaker? Their line of thought could be as follows. Quotations typically involve famous people. The term "buck" suggests an American. The most famous Americans are presidents. There are a total of thirty-nine (at the time of the experiment) presidents. Based on the process of long-term memory retrieval, if similarity-matching of the quotation does not partially match any schema, the students would be forced to infer the identity of the speaker only through frequency-gambling. Therefore, Reason also asked the British students to recall as many presidents as possible in five minutes.

Figure 3-1 displays two discrete probability distributions. Reason terms the rear distribution the "salience gradient" dominated by the first eight presidents listed. In descending order, they are Reagan (the incumbent president at the time, and thus the most salient), Kennedy, Carter, Nixon, FDR, Lincoln, Washington, and Ford. President Truman is number twelve on the graph. This distribution shows that almost every British student knew Reagan was an American president, but only 13% could name Truman as a president. The front distribution is how likely each student attributed the quote, "the buck stops here" to a particular president. As the distribution depicts, there was no agreement among the British students over who said this quote. There are two important notes. First, over 80% of attributions were made to the five most frequently remembered presidents. This is shown by the first five bars of

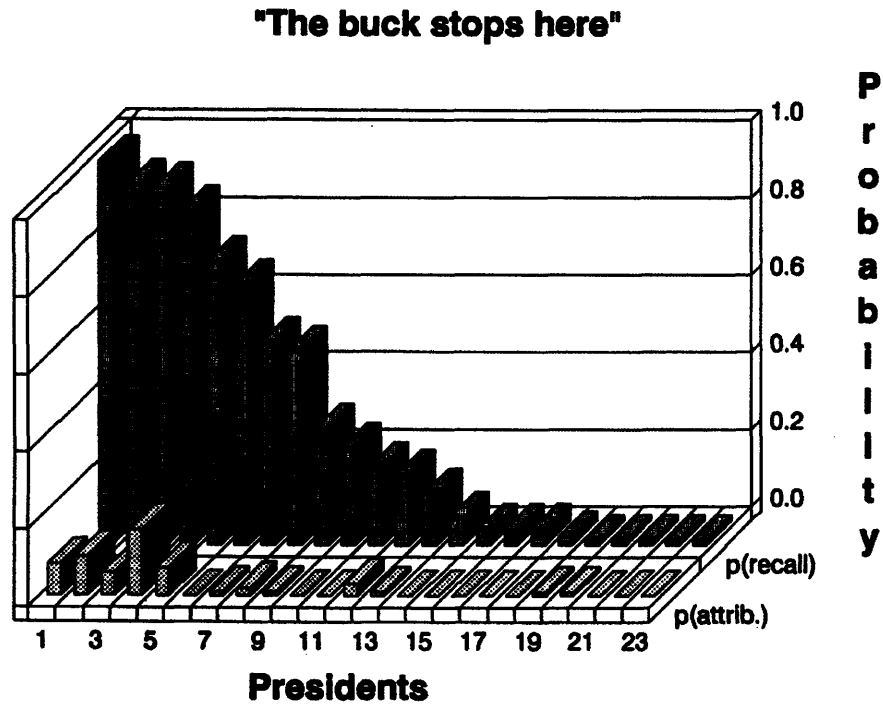


Figure 3-1: British psychology students' probability of recalling American presidents and probability of attributing the above quote to an American president [57].

the front distribution and confirms the presence of frequency-gambling in the face of failing similarity-matching. Because the average British student could not associate the quote with any president, he or she was most likely to attribute that quote to the most easily remembered, the most salient president. Second, President Nixon, the fourth on the list, dominated the attributions, as seen in the front distribution. Upon questioning the students after analyzing the data, Reason found that many matched the word "buck" to mean a dollar rather than responsibility. Out of the presidents most likely to remember, British students knew Nixon was involved in some scandal that presumably included money. Therefore, they inferred that the buck, meaning money, stopped here, in Nixon's pocket. Similarity-matching, then, dominates inference when the cues can be directly associated with stored knowledge; they fill the informational slots of particular schemata. Frequency-gambling dominates when the cues are ambiguous and there is little contextual knowledge in the required area.

3.1.2 Working Memory

While long-term memory storage has presumably infinite capacity and long-term memory retrieval occurs very rapidly independent of its size, working memory is limited both in the quantity of knowledge it can contain as well as in the temporal preservation of knowledge. Working memory is where all actions that require conscious thought are processed. Consider a mother of two children who runs over sharp

metal with her van and receives a flat tire. If the mother has had to fix a flat tire before, she can associate the present circumstance with that past experience stored in long-term memory. However, the remembrance of what occurred during the last flat tire is not sufficient to fix this flat tire. It takes conscious thought to properly position the jack underneath the van, unscrew the lug nuts, etc. This occurs in working memory.

Capacity

The capacity of working memory is limited to a small set of symbols. George Miller in his famous 1956 paper showed that the total number of symbols that could be stored in working memory, what is referred to as the memory span, is “the magical number seven plus or minus two” [44]. (Note that Miller refers to immediate memory in his paper and Newell and Simon refer to short-term memory [47]. The term working memory encompasses both of these.) These experiments which test memory span are simple in that a string of digits is provided to the subject, and the question is how many can the subject remember. For example, an average human can remember the last seven digits of a telephone number in working memory. As long as the area code is familiar enough that it can be represented in long-term memory, then this presents no problem for working memory. However, if the area code is new, it is unlikely that the human will remember all ten digits because all ten digits must be present in the working memory capacity.

Chunking

A successful strategy for overcoming the limitations of working memory capacity is chunking. A chunk of information is a set of data that is held together in working memory by associations in long-term memory [89]. Newell and Simon defined chunks as recognizable stimulus patterns [47]. One example of chunking is asking a subject to remember seven three-letter words. As long as these words are familiar and thus stored in long-term memory - cup, can, car, dog, eye, etc. - the entire word is represented as one chunk in working memory. Thus, Miller’s number of seven does not have to refer to seven letters, but could be seven chunks. In this example of three-letter chunks, the human could store up to twenty-one letters in working memory. Another example is to take a seven letter sentence whose words are combined with rule patterns or known associations in long-term memory - America’s national anthem is the *Star-Spangled Banner*. This sentence then is one chunk of information. A final example of chunking is to take a string of alphanumeric characters and parse them into chunks. For example, the string FB IJF KTV cannot be as easily remembered as FBI JFK TV.

Duration

Not only is there a limited quantity of information that can be stored in working memory, but there is a temporal decay of information that resides in working memory. In 1959, Brown, Peterson, and Peterson conducted a simple experiment to test the

FUNCTION	EXAMPLE
Coordinate performance on multiple tasks	Stockbroker must converse with a client while checking current prices of a volatile stock
Temporarily hold and manipulate information stored in long-term memory	Physician may compare the set of symptoms presented by a patient with those of previous patients
Change retrieval strategies from long-term memory	Scientist struggling to solve a problem must consider a variety of approaches
Attend selectively to stimuli	Sailor must attend only to the radar display and ignore nearby conversations

Table 3.1: The central executive

duration of information in working memory [89]. In this Brown-Peterson (typically shortened, but representing both Peterson and Peterson) paradigm, human subjects are presented with three random letters and are told to remember them. In order to keep the subjects from verbally rehearsing the letters, they are then told to count out loud backwards from some starting value in increments of three's. At some point, an auditory cue tells the subjects to stop counting and to recall the random three letter sequence. These experiments found after twenty seconds of counting backwards, almost no one could remember the original letter information. In fact, literature suggests that without continual rehearsal, little information is retained after ten to fifteen seconds [89].

Working Memory as a System

Working memory is a system of three components [57, 89]. The first is the verbal component of the system, which contains both a phonological store and the articulatory loop. The phonological store is a passive storage of information in linguistic form. The articulatory loop is where the verbal information is rehearsed. The second component is the visuospatial sketchpad where information is passively stored and actively rehearsed primarily in visual form. For both of these components, rehearsal loops account for the necessity of humans to continually repeat information so that it does not decay from working memory. The third component is the central executive which controls the information processing in the working memory and assigns attentional resources to subsystems. Table 3.1 displays the four main functions and examples of the central executive [89]. Note that the central executive for the stockbroker, physician, and sailor has to assign resources to both the articulatory loop and visuospatial sketchpad. This is not surprising because humans are auditory and

visually stimulated creatures. What is surprising is how these two components can function in parallel under the central executive.

In 1968, Brooks performed two experiments that investigated the relationship between spatial and verbal working memory [89]. In the first, he asked subjects to imagine a capital letter, such as F. Then he asked the subjects to mentally “walk” around the edge of the letter, and every time they approached a corner, he asked them a yes-no question about the orientation of the corner. For example, was the corner facing lower right? The subjects answered this question either by a vocal yes-no or by pointing to a column of Y’s and N’s. In this experiment, then, Brooks forced the subjects to work in spatial memory (walking around the letter), and then he required a response either in verbal memory (vocal yes-no) or in spatial memory (pointing to the column of Y’s or N’s). He found that subjects performed better with a verbal response, suggesting that subjects could divide resources up between spatial and verbal working memory without loss of performance. However, when the subjects were required to tax the spatial working memory by responding to another visual stimulus (the columns of Y’s and N’s), there was an overload or an interference in the visuospatial sketchpad.

For the second experiment, Brooks reversed the working memory component from spatial to verbal for the subjects’ primary task. In this experiment, he gave the subjects a familiar sentence, such as “The quick brown fox jumped over the lazy dog.” He then asked them to identify each word in the sentence as an adjective, noun, verb, etc. Again, they responded either verbally or by pointing to a table that displayed the words “adjective, noun, verb,” etc. This time, because the primary task was in verbal working memory, the subjects performed better by pointing to the correct answer. Thus, they again divided their tasks between the articulatory loop and the visuospatial sketchpad.

The conclusion then is that working memory is composed of two subsystems which contain and operate on different types of information using different resources. However, the input of new stimuli which operate in the same working memory subsystem as the one already in use causes disruptions. Therefore, the stockbroker in Table 3.1 does not converse with the client and check current stock prices through a playback machine. Rather, he or she converses with the client through the articulatory loop and processes the current stock prices through ticker tape or displays thereby operating in the visuospatial sketchpad. The central executive controls this process. Finally, the limitations of the working memory, namely its memory span of 7 ± 2 symbols, its decay after tens of seconds with no rehearsal, and its susceptibility to disruption by additional stimuli, all add to the human’s cognitive load in solving problems. Therefore, just as there are predictable biases resulting from long-term memory storage in the form of schemata, the constraints of working memory force humans to find heuristics in decision making so that they can reduce or minimize the “cognitive strain” involved in processing too much information in too noisy an environment.

The implications for learning tactical knowledge from the previous sections are two-fold. First, we must assume that all human subjects will take shortcuts. The question of expertise, then, may entail whether the shortcuts are valid and appropriate

based on past experience. Second, battlefield environments are extremely complex. If too many variables are given to the human expert to model complexity and force decisions, the human will naturally not attend to some variables, and the results may appear to have gaps.

3.2 Human Problem-Solving

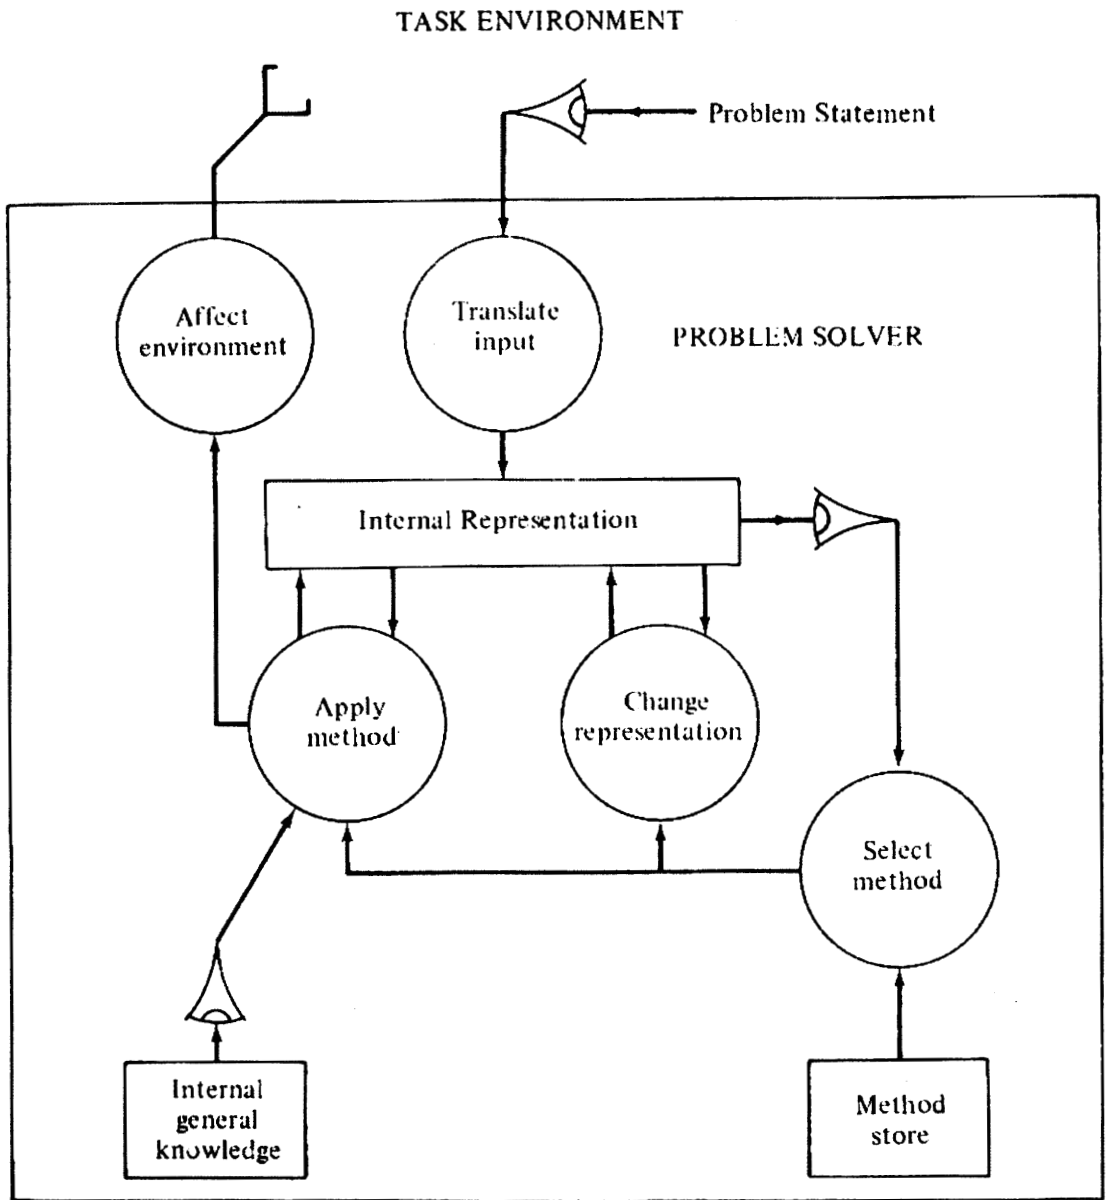
The work of Newell and Simon in examining human problem-solving laid the foundation for the artificial intelligence (AI) community to begin constructing human-like agents. This research borders very closely to work in the AI community without crossing over, and thus a general understanding of the beginnings of AI research is important. In the following discussion we present only a few important points from Newell and Simon's theory of human-problem solving relevant to this research. These are the modeling of a general problem-solver, the definition of a means-end analysis by heuristic search methods, and the description of a goal-oriented system.

3.2.1 Information-Processing Theory

Newell and Simon proposed that the thinking, or more specifically, the problem solving of humans could be characterized as information processing [47]. This is not to say that man should be modeled as a computer. Rather, it describes how "man processes task-oriented symbolic information" in problem solving. Newell and Simon challenged human subjects with problem solving tasks of cryptarithmic, logic, and chess problems. During each task, the subjects "thought aloud" as they mentally searched through their own problem space to find the solution. From these verbal reports, Newell and Simon constructed problem behavior graphs, which pictorially represent the states of knowledge and operators the subjects stepped through on the path to a solution. These problem behavior graphs allowed Newell and Simon to construct their theory of information processing behavior of human problem solving. They found certain information processing system characteristics that were common to all problem solvers and tasks. Probably more than any other, their work has shaped the notion of human problem solving and how it carries over to the artificial intelligence community.

3.2.2 General Problem-Solver

Figure 3-2, reproduced from Newell and Simon's book, depicts the general process of a human solving a problem [47]. It begins with the task environment and the problem statement. At this point, the problem solver must translate the problem statement into an internal representation. This internal representation includes choosing a problem space. The internal representation and problem space are crucial to the ability to solve the problem. They are the framework within which the problem solver works, and they have the ability to "render problem solutions obvious, obscure, or perhaps unattainable." Next, a problem solving method is chosen and applied. The method




Note: the eye  indicates that input representation is not under control of inputting process

Figure 3-2: General Problem Solver [47].

may be terminated during the process of application depending on its performance of achieving a solution. After termination, the problem solver may select and try a new method, look for a different internal representation to help reformulate the problem, or simply quit trying to solve the problem.

Note that this description of the general problem-solver is one of serial processing. A method is selected and applied one at a time. This does not mean the perceptual and sensing processes have to be serial as well. Indeed, a major human advantage is the parallel processing of sensory inputs. As Newell and Simon say, “the problem solver may see many things at once; it only does one thing at a time about them.” For example, suppose the task is to divide a number by an integer, say seven. If the number is exactly divisible by seven, the subject is asked to reply with a simple yes, else the subject should reply no. For the first problem, the human subject is given the number 35642. Is it exactly divisible by seven? A second problem consists of two numbers: 35642 and 69416. If information processing was parallel, it would take an equal amount of time to solve both problems. For a human, it does not. As such, the general problem-solver will have to take an iterative approach towards selecting an appropriate goal given the task environment and problem statement, selecting a method given the internal representation and problem space, evaluating the results, selecting another goal, etc. By working through the problem, new subgoals may surface. These subgoals may branch in different directions, some taking immediate precedence, some to be returned to later. Therefore, the general problem-solver must possess the ability to maintain a goal stack or hierarchy.

3.2.3 Means-End Analysis

The problem solving task is a means-end analysis accomplished by selecting and applying a search method. The goal of the search has been articulated, and the problem is to search through the problem space to achieve the goal. As Newell and Simon write, “a method can be understood only in reference to its goal.” Furthermore, the internal representation of the problem statement constrains the means of searching. This is why the internal representation and choice of problem space is so important. Also, note from Figure 3-2 that there is a method store. The method exists independently of the problem formulation. Therefore, there must be an interpretation capability to apply the method to the specifics of the problem space. Figure 3-3 depicts the heuristic search method from Newell and Simon’s book. Every search method has its specifics, but Figure 3-3 is, in its abstract form, *the* heuristic search method (italics are from Newell and Simon’s book) [47]. The heuristic search method begins with the initial element of the search space. Then, a current element and operator are chosen. The operator is applied to element, and the outcome is tested for the solution. If it is not the solution, there must be evaluative criteria as to whether to accept or reject the new element. If accepted, it is inserted into the solution path. If it is rejected, the new element is not inserted into the solution path. The next step is to decide to continue to applying new operators to the current element, to advance by replacing the current element with the new element (typically if the new element was accepted), or to completely abandon the current path and try again. As

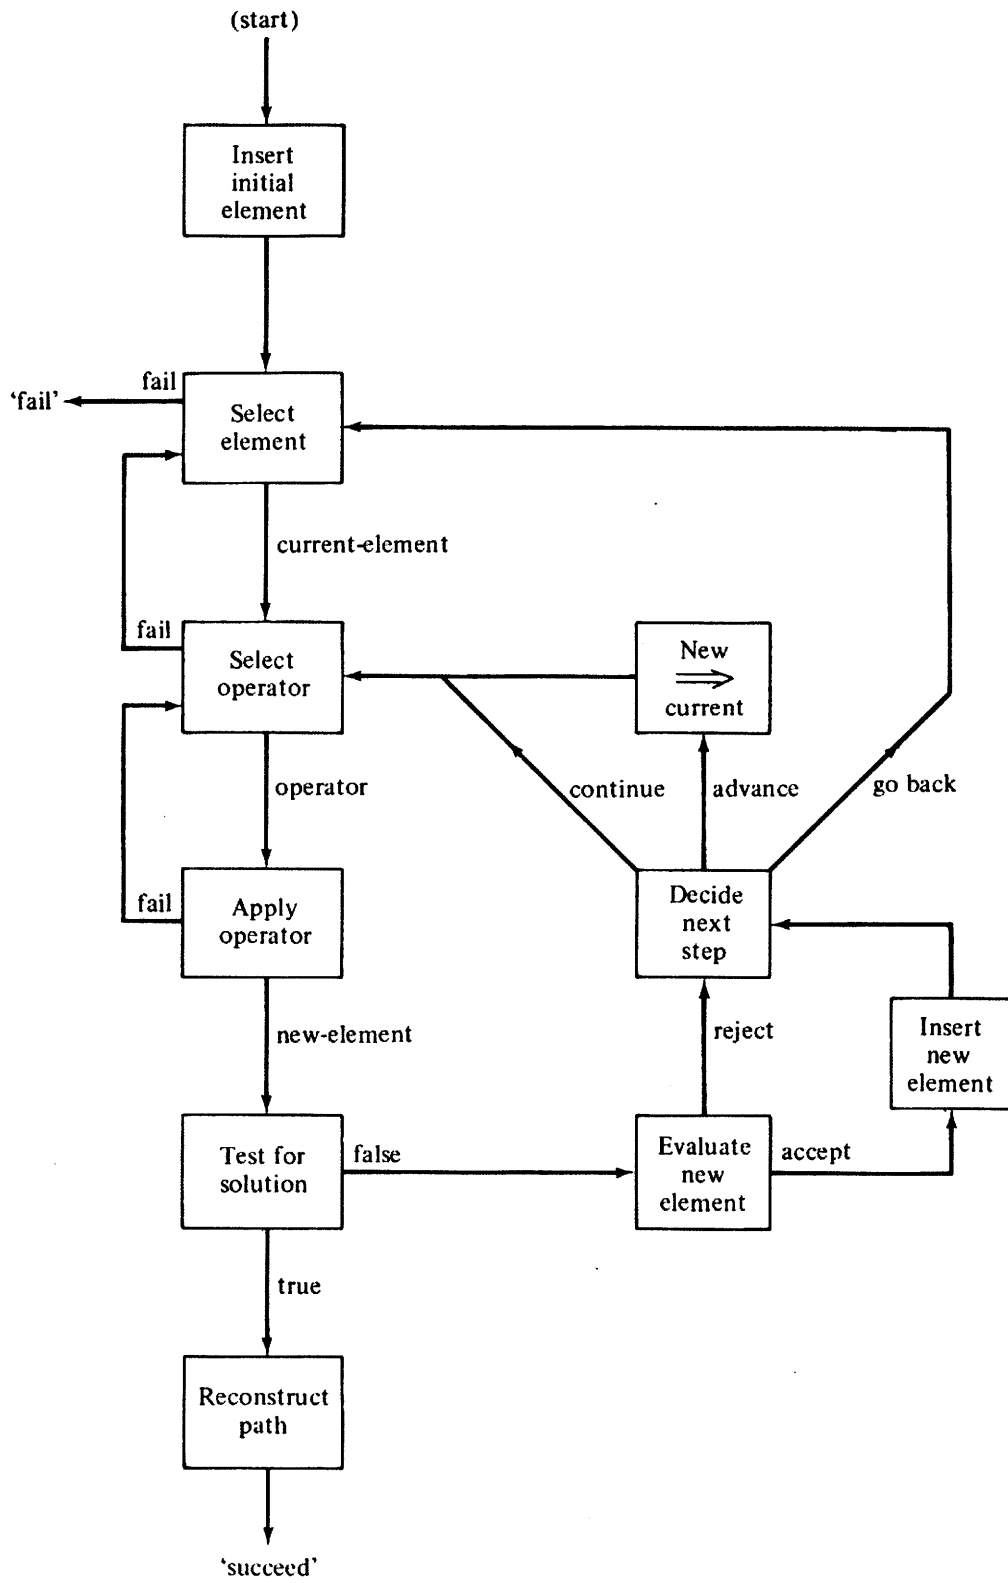


Figure 3-3: Steps of a heuristic search method [47].

long as the previously accepted elements along the path have been stored, the option to go back can start from any of these stored elements. When a solution is found, the path is reconstructed, and the problem is solved. Note that the specifics of a particular heuristic method will be found in how it selects elements, what operators are available, how it selects the operator, and the criteria for accepting or rejecting the new element.

3.2.4 Common Characteristics

After constructing and analyzing the problem behavior graphs from the cryptarithmic, logic, and chess tasks, Newell and Simon collected a list of common characteristics for the information processing system of human problem solvers. The first is the size, access characteristics, and read and write times for both long-term and working memories (see Section 3.1). The second is that information processing occurs serially. Perception of the task environment is no doubt parallel, but choosing and selecting methods is serial. The third invariant characteristic is production-like and goal-like organization. (Newell and Simon compared humans to production systems, but felt that the argument was not conclusive.) The goal-driven behavior, demonstrated by human problem solvers, is so crucial to modern attempts at agent modeling that we present the following list of six criteria. These criteria are taken directly from Newell and Simon, and they specify a goal-directed information processing system (IPS) with six characteristics:

1. *Interruptibility.* If the IPS is removed or distracted from a situation, it later returns to directed activity at the same point.
2. *Subgoalting.* The IPS itself interrupts its activity toward a goal to engage in an activity that is a means to that goal, and then returns (often after considerable time lapse) to the activity directed toward the original goal, making use of the means produced by the subgoal.
3. *Depth-first subgoalting.* When the subgoalting behavior indicated above occurs to a depth of several goals, the evidence is particularly conclusive [that the agent's behavior can be considered goal-oriented].
4. *Equifinality.* If one method for attaining a goal is attempted and fails, another method toward the same goal, often involving quite different overt behavior, is then attempted.
5. *Avoidance of repetition.* More generally, the system operates with memory of its history of attempts on goals, so as to avoid repetition of behavior.
6. *Consummation.* If the goal situation is attained, effort is terminated with respect to the goal.

The importance of the criteria is that it serves as a benchmark for distinguishing goal-oriented behavior from other forms of behavior, such as reactive behavior. The major theme in this criterion is that the long-term behavior of the system should be cohesive and not random, always traveling along a path or branch to the desired end goal.

3.2.5 Problem Solving to Decision Making

Every decision can be posed in the form of a problem. Newell and Simon focused exclusively on problem solving. For the human subject given a logic task, there must be a sequence of decisions made in searching the problem space. These decisions are exactly what specify a heuristic search method, namely what element to choose next, what operator to select, and how to know if the result should be kept or not. These are the decisions the human problem solver must make. Yet, there are two major differences between Newell and Simon's tasks and the task a pilot in a tactical situation would face. One, Newell and Simon's problem solvers could pace themselves and they had all the information in front of them. They were under no time pressure. Two, there was no uncertainty in the information. By taking their time, they could process the information deliberately but slowly. As Newell and Simon write, "since the problem solving activity is self-paced, the problem solver can adjust his rate and style of processing information so that he does not appear either to have a rapidly decaying short-term memory or to rehearse" [47]. Here, short-term memory is synonymous with working memory. The self-determined pace of problem solving helped these humans to adjust processing rates and styles to ease the cognitive strain of working within cognitive limitations. However, they are still subject to the same cognitive limitations as battlefield decision makers. The difference is, again, both the time pressure of a decision and the uncertainty in the information. Time pressure and uncertainty, when combined with extremely limited working memory capacity, force decision makers to turn to decision heuristics that reveal a number of biases humans are prone to possess. These decision heuristics are decidedly different from Figure 3-3 in that they do not systematically search the problem space for a correct path. Rather, these decision heuristics match the starting conditions, the desired solution, and any relevant, predicted subgoals that will branch off along the solution path to earlier experiences. If a match is made, the previous solution path is simply chosen and applied. Of course, during application the results will be monitored, and the method can always be changed. This sort of decision heuristic will provide solutions very quickly. Yet, the solution path will never be optimal, it might be satisfying, and at worst, it could be completely wrong. This is not to suggest that the information-processing theory that Newell and Simon proposed should be abandoned. Indeed, engineering psychologists have abstracted it to a higher level so that experimentally determined human decision heuristics and biases can be modeled.

3.3 Human Decision Making

Figure 3-4 depicts an information processing model of decision making [87]. More specifically, this model is based on signal detection theory. Signal detection theory states that there are two discrete states in the environment, the signal and noise, and it is difficult to discriminate between them. Beginning at the left, cues in the environment become sensory inputs to the system. These cues are numerous, changing, and uncertain. Thus, there must be a filtering process that selects cues and then passes

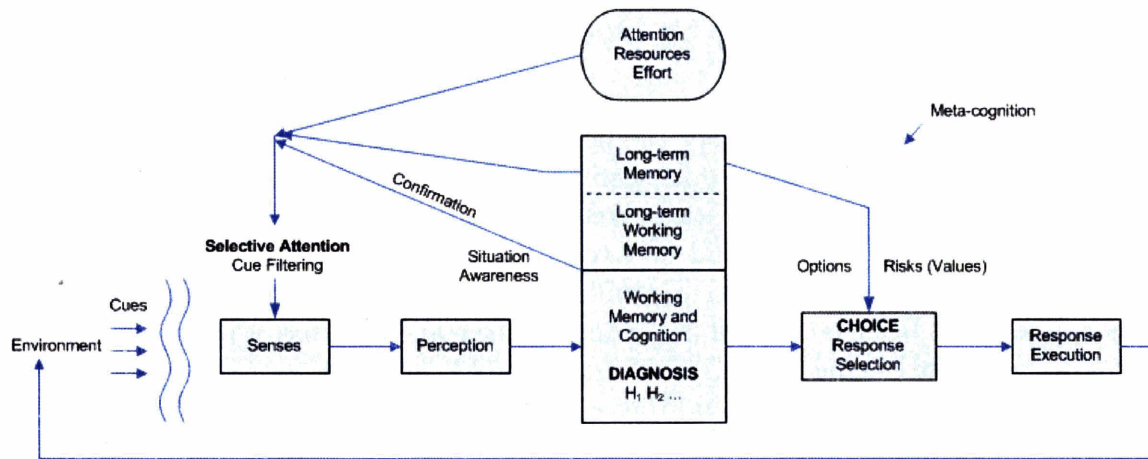


Figure 3-4: Key processes and components involved in an information processing model of decision making [87].

them to a perception process that estimates the importance of each cue to help build a situation awareness of the environment. At this point, the human information processor must integrate the cues from the perception process and the past experiences from the long-term memory. (Long-term working memory is a newer concept of another mechanism inside the working memory that uses long-term memory for skilled performance. It will not be discussed here, but see [19] for more information.) Then, the human information processor must operate on the combined information in the working memory to form a diagnoses or hypothesis of the situation. Once reaching a diagnosis, the system must then choose an action. This choice can be based on options that are either recalled from long-term memory and the associated risks, on novel options that have not been encountered before, or the combination of the two. Upon execution of the response or selected action, the environment will necessarily be changed and the process repeats itself. Furthermore, note that the environment is constantly changing during this whole process. Thus, as the system is forming a working hypothesis of the situation and choosing appropriate actions and reactions, new cues could enter that might either confirm or contradict it. Both the long-term memory and the working hypothesis act as feedback loops to the cue filter to help selectively focus on specific cues. Also, in choosing the correct response, the diagnosis must not only result in a list of options, but the corresponding risks associated with each one. Finally, meta-cognition is the system's awareness of its own limitations in seeking to find and choose correct decisions.

The following discussion of decision making will fall into two broad categories. The first category is the normative, rational theory of decision making. The second category is the human decision heuristics and biases in decision making. After noting the significant divergence of normative and human decision making, we summarize the implications for the information-processing model of human cognition, namely how to judge the appropriateness of the decision. If the outcome of the decision scored well, but the basis of the decision was a human bias, how well should it be

received?

3.3.1 Normative Decision Making

Estimating Uncertain Cues

Normative decision making begins with the processing of information to create an accurate situational awareness and to form a proper diagnosis. The first step is to estimate the cue worth [87]. Each cue can be characterized in three main ways: diagnosticity, reliability, and salience. Cue diagnosticity is the probability that given a hypothesis, the cue should be present. Thus, if the hypothesis is that it is raining, there is 100% probability that raindrops should be present. If the hypothesis is that there is 50% chance of rain showers, the presence of dark clouds will be somewhat diagnostic, but not completely so. If the decision maker is considering several alternative hypotheses, the diagnosticity of the cue will vary across each hypothesis. Cue reliability is the probability of whether the cue can be believed or not. The reliability of raindrops would be 100% unless someone was spraying water out of a window. An eyewitness account may point to a suspect as guilty (cue diagnosticity), but the reliability of the eyewitness must still be determined. More importantly, cue reliability is independent of cue diagnosticity. Therefore, because the intersection of two probabilistic independent events is the product of their individual probabilities, cue reliability and diagnosticity are then multiplied together to form the information value of the cue. The final characteristic of cues is their salience, or ability to draw attention to themselves. This salience is important in two ways. First, the salience of a cue draws attention to itself so that it can be analyzed in terms of its diagnosticity and reliability. For instance, the cold wetness of a raindrop on the skin is very attention-getting. Second, those cues which are not salient may be filtered out from the perception process. This is why debugging software code is so time-consuming. There is no salient feature of the bug; it must be sought after. As will be discussed later, certain decision heuristics and biases are a direct result of cue salience. Finally, note that the long-term memory plays an important role in estimating the cue worth. Frequently encountered salient cues are easy to remember, and the working memory is not needed to help the filtering or estimating process. For example, raindrops, ambulance sirens, and the smell of freshly baked cookies do not require intense processing time to acquire accurate situation awareness. It is when long-term memory interacts with ambiguous and uncertain cues in the working memory that difficult decisions arise.

The information value of the cue allows the decision maker to integrate all existing cues into a few alternative hypotheses that these cues reflect. The integration could simply involve summing the information values over all cues for each hypothesis. Taking it one step further, the information value could be combined with prior beliefs to form conditional probabilities about the existence of cues and hypotheses and thus follow a Bayesian approach. The decision maker can choose to continue to seek more information from the environment to confirm one hypothesis over another. Otherwise, the decision maker merely selects a hypothesis and now must choose a

correct response. Rational choice must adhere to the following three criteria [15]:

1. It is based on the decision maker's *current* assets.
2. It is based on the possible consequences of the choice.
3. When these consequences are uncertain, their likelihood is evaluated without violating the basic rules of probability theory.

Notice once again the importance of probability theory. From estimating cue worth, to integrating information values, and now to evaluating the consequences of choices, probability theory is necessary to deal with the inherent uncertainty. The theory of subjective expected utilities addresses how to choose rationally.

Subjective Expected Utility Theory

Mathematician John von Neumann and economist Oskar Morgenstern published *Theory of Games and Economic Behavior* in 1953 where they presented the principle of expected utility [83]. The theory is based on a list of axioms. Von Neumann and Morgenstern mathematically proved that if a decision is made in accordance with the axioms, the decision maker will be able to define utilities or personal values. The decision maker can then choose between the probabilistic consequences of the list of alternative choices based on whichever choice maximizes the expected utility. The implications of this theory are amazing. By taking the time to analyze choices made in the past, present personal values, and probabilities in hypothetical contexts all in the context of satisfying the axioms, an expected utility analysis can reveal the exact nature of the decision maker's personal values or utilities. Now the decision maker is armed with a list of his or her own personal values to judge future choices that involve probabilistic consequences. Herbert Simon, who co-authored *Human Problem Solving* as discussed above, described this theory as, "a beautiful object deserving a prominent place in Plato's heaven of ideas" [67]. The theory is beautiful, but seemingly unrealistic.

Many objections quickly arise to this theory of subjective expected utility [15, 57]. The first objection is that it is "dehumanizes" life's important decisions by degrading it to numbers and rules. The second objection is that the theory seemingly purports that mathematics will reveal a human's personal values better than how the human knows his or her personal values. The third objection is that values change, and thus certain decisions that maximized expected utility in the past will no longer do so in the present. The fourth objection is that it is extremely difficult to assign subjective values and probabilities to the consequences of choices. The fifth objection is that probability theory invokes the notion of a closed set of events, and how can every possible outcome be accounted for?

In Chapter 1, we presented an example of calculating the utility of a military commander choosing to assign a dangerous mission to either a human or an AV. In order to initially carry out the calculations, we assumed that the human and the AV had an equal probability of life, loss, success, and failure in the mission. Now, this is not true as stated in the summary of that section, and subjective probabilities and their associated weights have to be determined. This seems too tedious and general

to be helpful. However, if we could take the time to complete a subjective expected utility analysis of human experts, there might be a chance to learn useful insights into the experts' value systems. It is intriguing, but it appears to rigid to be applied to the dynamic environment of the battlefield. Robyn Dawes, a strong proponent of subjective expected utility theory, agrees that it is difficult to determine whether a choice has satisfied all of the axioms or not [15]. She concludes, therefore, that we should not be bound by the axioms, but should consider them as helpful in generating an alternative list of decisions.

3.3.2 Bounded Rationality and Satisficing

Normative theories describe how decisions should be made in rational and mathematical manner. The three criteria given by Dawes and the subjective expected utility theory are the benchmarks of this approach (see also [8] for a Bayesian inference approach to naturalistic decision making). However, human decision making frequently contradicts such normative propositions. Human decision heuristics and biases reveal the bounded rationality and satisficing nature of human decision making. Simon first proposed the expression *bounded rationality* and stated:

The capacity of the human mind for formulating and solving complex problems is very small compared with the size of the problems whose solution is required for objectively rational behavior in the real world - or even for a reasonable approximation of such objective rationality. [66]

Simply stated, bounded rationality is a combination of the human cognitive limitations and the enormous problem space in the real world. When this problem space also includes uncertainty and the decision maker is under time pressure, humans are forced to use shortcut methods that seek a satisficing solution rather than an optimal one. Decision heuristics are not altogether bad. In fact, humans employ them for two main reasons. First, they do a pretty good job; else they would not be used. Most of the time, they find a satisfactory answer. Second, under extreme time pressure, humans do not have the option of evaluating a set of alternative choices. However, a decidedly wrong reason for using heuristics, but common to humans, is the desire to minimize cognitive strain. In a word, we simply do not want to expend the mental effort to sort through information and find the best solution. In the following sections, we first present the human inability to follow the principles probability theory. Then we discuss well-known decision heuristics and biases and provide examples. Each section concludes with implications for decision making.

The following discussion of human subjective probability and human decision heuristics and biases relies entirely on three sources: Wickens and Hollands' textbook *Engineering Psychology and Human Performance* [87], Robyn Dawes' book *Rational Choice in an Uncertain World* [15], and James Reason's book *Human Error* [57]. Also, note that the names Tversky and Kahneman will be mentioned frequently throughout this discussion. Through many, many experiments, they pioneered the research of drawing out human heuristics and biases. See [36] for a recent compilation of decision theorists research edited by Tversky, Kahneman, and Slovic.




Figure 3-5: Human distortion of probability [87].

3.3.3 Human Subjective Probability

Mean, Variance, Conservatism

In analyzing a set of data or observations, humans exhibit some biases in predicting basic statistical parameters. First, the good news is that humans can estimate the mean value of observations pretty well. Consider simply estimating the midpoint of this page. Second, in a binary set of observations (faulty vs. normal parts, true vs. false, etc.), humans can estimate the total ratio of occurrence (percentage of normal parts in set, the ratio of true to false answers) fairly well, given that the ratio falls between extreme values. However, the bad news is that near the extreme values, the proportions are distorted. At the lower extreme of 0, humans tend to overestimate the frequency of occurrence. Near the upper extreme of 1.0, humans underestimate the frequency of occurrence. This phenomenon is depicted in Figure 3-5 where the human distortion of actual probability causes the function to both curve away from the origin and to fall short of reaching 1.0. Thus, humans are conservative, a sort of “never say never” behavior. The downwards offset of the curve is a result of a framing effect. This framing effect will be discussed more thoroughly. Third, humans do a relatively poor job of estimating the variance of a set of observations due to two biases, as shown by Figure 3-6. First, people tend to underestimate the variance if the mean of the observations is higher. Thus, for Figure 3-6(a), humans would estimate the variance in the length of the line segments on the left as higher than on the right. In fact, the variance is equal. The second bias of estimating variance is the influence of extreme values in the set of observations. For Figure 3-6(b), people would estimate the variance in the position of line segments on the left as higher than the right. This variance is also equal. Fourth, when predicting the future behavior of a nonlinear trend, the prediction tends to be a linear extrapolation. Therefore, if data indicate exponential growth, humans will underestimate the future growth using a linear prediction. Part of this bias appears to be, again, human conservatism. However, another underlying reason for this bias is the human long-term picture of events, where there tends to




Figure 3-6: Human biases in estimating variance [87].

exist some sort of “self-correcting mechanism” or increasing viscosity that slows down such growth. For example, say the very nonlinear and rapid growth of a spreading forest fire in the first few hours of burn was plotted as the number of acres destroyed against time. The prediction of the future growth of the fire by fitting a curve to this data would attempt to follow the nonlinearity and growth of the data. Yet, a human would predict a much more linear and conservative fit, believing that the fire will be eventually extinguished due to natural means, such as rain or lack of fuel, or artificial means by human fire-fighting agents.

Cause and Effect

Another common failure in human judgment of probability is a natural tendency to try to influence the outcome of probabilistic circumstances [15]. Ellen Langer at Harvard University performed a series of experiments to determine examples of this behavior. In one experiment, Langer observed that gamblers rolled the dice harder if trying to roll higher numbers than lower. In another, subjects participated in a lottery based upon NFL trading cards. In this lottery, every subject had a different NFL trading card, all cards were thrown in a bag, and the subject whose card was drawn won. For the first variation of this experiment, subjects randomly selected a card. For the second variation, subjects could select a player. In both variations, subjects were approached by a disguised experimenter who offered to buy their card. Now whether or not the subject had come into possession of the card randomly or not, the probabilistic outcome of the lottery would not change. However, the experiment found that in the second variation, when subjects chose their own NFL trading card, they demanded on average more than four times as much money for it. Thus, they behaved as if their choice of card affected the lottery outcome. In another experiment, Langer and Roth asked Yale undergraduate students to predict the outcome of a series of coin tosses [15]. Langer and Roth rigged the setup so that they could provide false feedback to the students. At the end of thirty tosses, every student had only “predicted” fifteen successfully, the expected chance outcome. The twist was that Langer and Roth grouped the number of successes for each student either at the beginning of the trial or end of it. Those students who were “successful”

near the beginning of the sequence described themselves as “better than average” at predicting the outcome. Those students who were only successful near the end of the sequence felt they predicted worse than average. Furthermore, Langer and Roth report that, “over 25% of the subjects reported that performance would be hampered by distraction. In the same vein, 40% of all the subjects felt that performance would improve with practice.” Therefore, in three common probabilistic situations - rolling dice, lottery, and coin tossing - humans behaved as if there was some sort of influence they could affect over the outcome.

A final example of misinterpreting probability is the human tendency to try and find meaning in random sequences. One simple example is the notion of a “hot hand” in basketball. If a three-point shooter makes two shots in a row, the player all of a sudden has a “hot hand,” and the team will continue to feed him the ball even if there are better plays available. This is not to discount the amount of training and skill in developing shooting capabilities and degrade it simply to luck. This is to say that if a shooter just made two shots in a row and the opposing team proceeded to call a timeout, that shooter’s prediction of his ability to continue to make two more shots in a row is no longer based on skill, but that he’s “hot.” This is irrational, and researchers have proven the concept to be false [15]. Another example, is the notion that accidents occur in bunches, or more specifically in “threes.” If three aircraft accidents occurred in a period of time, say as small as two weeks, could there not be some connection? People at least seem to think so. As a test, Dawes and Vaught gathered from the Federal Aviation Administration (FAA) the dates of all aircraft accidents between 1950 and 1970. For every pair of crashes, they plotted the time elapsed between the two. Then, they modeled aircraft crashes by a Bernoulli process, where every trial has a binary outcome of zero or one corresponding to the arrival of some event. In this case, the trial interval is one day, and the event arrival is an airplane crash. (They do not consider the possibility of multiple crashes occurring per day.) For a Bernoulli process, the interarrival times are geometric variables with constant probability, p . Equation 3.1 is the probability mass function for a geometric variable with constant probability, p , and k^{th} trial.

$$p_X(k) = (1 - p)^{(k-1)}p \quad (3.1)$$

In words, $p_X(k)$ is the probability of the occurrence of x at the k^{th} trial, where p is the probability of the occurrence of x at each trial. Therefore, Equation 3.1 states that the longer the time has elapsed after the occurrence of a single arrival, the more likely the next arrival will occur. Dawes and Vaught found that the interarrival times predicted by this stochastic process fit the data almost perfectly. Figure 3-7(a) depicts the number of crashes predicted by a Bernoulli process over the period 7300 days, which is equivalent to the twenty-year period they analyzed. Figure 3-7(b) is a histogram that depicts the frequency of occurrence of interarrival times, the time elapsed between the crashes in Figure 3-7(a). Dawes and Vaught did not publish what constant probability p fit the data, therefore, for this figure, the Bernoulli process assumes that the probability of an aircraft crash every day is one percent (a little less than four crashes per year, which is hopefully too high). Notice from Figure

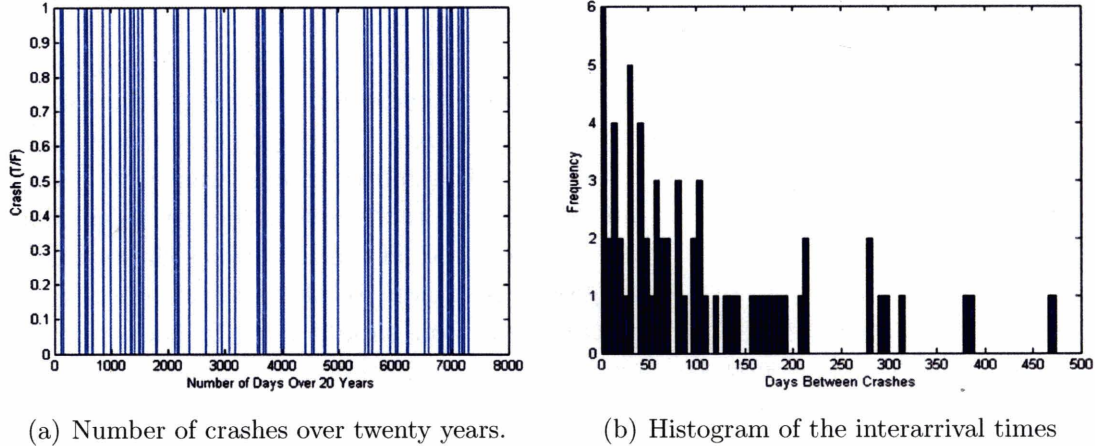


Figure 3-7: Hypothetical aircraft crashes modeled by a Bernoulli process with constant probability $p = 0.01$.

3-7(a) that random events do occur very close to each other, especially around days 700, 1500, and 7000. In fact, the conclusion from Figure 3-7(b) is that the highest frequency of occurrence of interarrival times, with a total of six, was the period from one to five days. Indeed, in continuous form, the probability mass function of the geometric random variable is the exponential probability density function of the form $p_X(x) = \lambda e^{-\lambda x}$ which equals λ when $x = 0$ and curves down to zero. Then is it true that accidents occur in threes? The conclusion from Figure 3-7(b) is that they at least occur frequently in pairs. Upon analyzing Figure 3-7(a) further, the solid band around 6800 days is actually four arrivals that occur over a time span of 43 days. Four crashes in a month and a half would be enough for humans to try and find meaning in this random sequence.

Statistical Models

One of the first steps necessary in forming a decision is to determine the information value of a cue (see Section 3.3.1). This information value consists of multiplication of the independent probabilities of cue diagnosticity and cue reliability. However, humans exhibit two fundamental errors in determining cue worth. First, humans tend to treat all cues as if they were of equal value. This is much easier than trying to differentially weight them. Second, humans present a salience bias, whereby they tend to select and give more attention to those cues which are more attention-getting. The combination of these biases with the mistakes in properly applying probability has led researchers to call for the end of statistical prediction by human decision makers [15, 16, 74, 87].

These researchers have conducted numerous experiments where expert decision makers predict an outcome of some case subject based upon a list of attributes and subject's scores in each of the categories. For example, Robert Libby had forty-three bank loan officers predict which thirty of sixty firms would go bankrupt within

three years of a financial report [15]. The loan officers requested information about each firm, primarily important financial ratios, such as the ratio of liquid assets to total assets, upon which they could make their prediction. Their predictions were 75% correct, but a linear regression was 82% correct. Furthermore, using just one attribute - the ratio of assets to liabilities - the linear model was 80% accurate. These researchers have concluded that “experts correctly select the variables that are important in making predictions, but . . . a linear model that combines these variables in an optimal way is superior to the global judgment of these very same experts.” Therefore, domain experts in statistical prediction should classify the appropriate predictor variables, show how to measure and encode them, and identify the correct direction of the variable’s weighting (i.e. - higher ratio of liquid assets to total assets equates to smaller chance of bankruptcy). At this point, statistical analysis through computation should make the appropriate prediction.

3.3.4 Heuristics and Biases

There are many decision heuristics and biases that humans employ because of the bounded rationality problem. These will be important in analyzing the data resulting from the knowledge-elicitation experiments so that we will have a complete picture of why the human subjects made particular decisions. These are presented in no particular order, but there definitely exists connections and common themes between many of them.

Representativeness Heuristic

The representativeness heuristic describes the tendency to first match existing cues with similar patterns represented in long-term memory and then to derive causality for the current situation based on the causality in the past. So far as a method, this is not a bad heuristic. How else do decision makers learn from past mistakes? However, the problem with the heuristic enters when cues are somewhat vague and uncertain and decision makers do not adequately take into account the base rate of occurrence. For example, consider the physician who must decide whether or not the presence of specific symptoms imply the existence of a disease. Say the patient has three out of five symptoms that represent and imply causality of disease X, and the patient also has four out of five symptoms that represent and imply causality of disease Y. The physician will most likely diagnosis the patient as having disease Y because of the representativeness heuristic irrespective of the probability of the actual occurrence of disease Y, which is its base rate. In Bayes Theorem, given by Equation 3.2, the event A represents the disease and the event B represents the collection of symptoms.

$$P(A|B) = \frac{P(A)P(B|A)}{P(B)} \quad (3.2)$$

The question the physician is trying to answer is what is the probability of A given the existence of B, that is $P(A|B)$. In the example scenario, the physician knows $P(B|A)$, which is the probability of the collection of symptoms given the existence of

the disease. He has already determined that the symptoms point to the disease. The physician also knows $P(B)$, the probability of the existence of the collection of symptoms. He has already determined the symptoms are present. What the physician fails is to take into account $P(A)$, the base rate of the disease. In fact, this failure to account properly for base rates and prior probabilities over and above the conservatism humans already exhibit in estimating probabilities (see Section 3.3.3) led Tversky and Kahneman to pronounce “in his evaluation of evidence, man is apparently not a conservative Bayesian: he is not a Bayesian at all.” The representativeness heuristic hinders the decision maker when causality is determined by representativeness of the cues without considering the probability of the cause itself.

Availability Heuristic

Tversky and Kahneman defined the availability heuristic as the “ease with which instances or occurrences can be brought to mind” [87]. The availability heuristic is a bias that decision makers use to estimate the probability of an event. Dawes uses the example of the homeless [15]. How many people are homeless due to long-term mental illness? Research shows that about one-third could be characterized as mentally ill due to either “current mental distress or a history of psychiatric hospitalization.” Most homeless people are just simply poor. The biased national opinion of correlating homelessness with mental illness can be explained by the availability heuristic. Dawes states:

Search your memory for the homeless people you saw most recently. What were they like? The unobtrusive homeless person is easily forgotten. We tend to remember the person who sings on the bus, who accosts strangers with stories of lost fortunes, who is drunk, or who is obviously high on some drug. Moreover, we may prepare ourselves to behave in certain ways if such a person approaches us . . . [which] enhances the recall of the event leading to it.

Alternatively, consider the two F-16 pilots (see Section 2.2.1) who stated that unmanned aircraft should never operate in the same airspace as manned. Why? An eyewitness account of a Predator drifting into the path of a group of fighters almost caused a mid-air collision. How many times has an unmanned aircraft operated in the same airspace as manned and not caused any problems? Those events are not remembered because they are not salient. Who has actually kept a statistic of safe, regulated operation of manned and unmanned aircraft in the airspace? Who has actually compared the number of times manned aircraft operating in the same airspace as other manned aircraft almost caused a mid-air collision? The pilots’ conclusions, based both on a biased recall of events and a biased experience (how much have they actually interacted with unmanned vehicles), are irrational. People who use the availability heuristic overestimate probabilities. The salience of the event, the “elaboration in memory” of the experience, the recency of the memory, and the simplicity of it (Wickens notes that a single failure is much easier to remember than a compounded double failure [87]), all lead to the availability heuristic.

Anchoring

The anchoring heuristic describes the tendency of humans to propose an initial solution to a problem, especially an ambiguous problem, and then simply adjust their answer from there. It is as if they have attached a “mental anchor” to it. For example, Tversky and Kahneman asked a group of students what percentage of African countries were members of the United Nations [15]. The correct answer is 35%. Before the students gave their answer, Tversky and Kahneman presented a wheel with numbers between one and one hundred. They would spin the wheel and then ask the students to simply judge whether the answer was higher or lower than the number on the wheel. Then the students would give their numerical answer. In actuality, the wheel was rigged to stop on either ten or sixty-five. When the wheel stopped on ten, students made an average estimate of 25%. When the wheel stopped on sixty-five, students made an average estimate of 45%. Therefore, the numbers ten and sixty-five became anchors for the students’ answers “even though the subjects were led to believe that these numbers were generated in a totally arbitrary manner.” In another example, students were asked to mentally estimate the value of $8!$. To describe the concept of a factorial, students were either presented with $1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$ or the reverse order of $8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$. The median answer given by students presented with the ascending sequence was 512. The median answer for those students presented with the descending sequence from 8 was 2,250. The correct answer is 40,320. The first number then in the multiplicative sequence served as an anchor. Note that once again this highlights the necessity to carefully design how an order of information is presented to the human. Finally, another variation of the anchor-and-adjust principle is when new information is given to the decision maker after an initially formed hypothesis. A study of Army intelligence analysts found that they gave considerably more weight to new pieces of information that supported their original hypothesis than those that did not [87]. Rather than forming a new hypothesis, humans tend to only adjust their answer based on the originally conceived answer.

Confirmation Bias

The confirmation bias describes decision makers who actively seek only that information that confirms their initial hypothesis. If in this active seeking, contradictory information is revealed and cannot be interpreted in the context of the initial hypothesis, it is discounted or completely discarded. Furthermore, completely undiagnostic or ambiguous cues will be interpreted in light of the conceived hypothesis and be found to support it. In the *USS Vincennes* incident of July 1988, the Ticonderoga class cruiser shot down an Iranian Airbus airliner. The context of the situation revealed tension in the Persian Gulf. Iranian forces had successfully exploded a mine into the *USS Samuel B. Roberts* earlier that April. Radar operators aboard the *Vincennes* initially believed the radar contact was an Iranian fighter because it appeared to follow a fighter aircraft’s profile out of Iran as depicted by intelligence reports. However, when the identification-friend-or-foe equipment returned a non-hostile con-

tact, the crew discounted it. The exact nature of the incident is still controversial, but it demonstrates the danger of the confirmation bias. The group of Army intelligence analysts described in the anchoring heuristic also exhibited the confirmation bias. Not only did they give undue weight to contradictory information that was given to them (passively received) and thus demonstrate a “mental anchor,” they were explicitly given the choice of what information they still desired. Consistently, the analysts chose to actively seek for information that confirmed the hypothesis rather than seeking for disconfirming evidence. However, this is not to say that maintaining a “working hypothesis” is the wrong strategy in decision making. Rather, the decision maker needs to consciously look for or at least rationally consider disconfirming evidence and ascribe to it the appropriate amount of weight that it is due.

Overconfidence Bias

The overconfidence bias describes decision makers who are overconfident in the accuracy of their predictions or solutions. As an example, Fischhoff and MacGregor asked human subjects to predict the outcomes of certain events, such as sports games and elections [87]. They then asked the subjects to rate their confidence in the predictions. The outcomes consistently revealed that confidence exceeded accuracy, sometimes as high as 20–30%. Consider the student who has spent a long time studying and really feels that he or she has a good grasp of the knowledge only to take the test and receive an average grade. Furthermore, J. Reason notes that when overconfidence and the confirmation bias are combined,

Humans will be loathe to change a completed plan set for action especially under the following conditions: [57]

1. when the plan is very elaborate, involving the detailed intermeshing of several different action sequences
2. when the plan was the product of considerable labor and emotional investment and when its completion was associated with a marked reduction in tension or anxiety
3. when the plan was the product of several people, especially when they comprise small, elite groups
4. when the plan has hidden objectives, that is, when it is conceived either consciously or unconsciously, to satisfy a number of different needs or motives

The implications of overconfidence are to disregard additional information by either ceasing to gather information or discounting its worth.

Elimination by Aspects

Elimination by aspects is a sequential search method for deciding between alternative choices. Elimination by aspects purposefully seeks to find a satisficing answer rather than an optimal one. The focus is on the aspects of the alternatives not necessarily

the alternatives themselves. The process is to choose the most desirable aspect, eliminate all other choices that do not exhibit it, choose the next most desirable aspect, eliminate all other choices that do not exhibit it, etc. At the end there should be either a single answer or only a small number that could then be evaluated more thoroughly. The accuracy of the method hinges on the order of elimination. If the aspects are ranked according to their desirability, the answer is typically satisfactory. If the aspects are probabilistically chosen according to a model of their importance (i.e. - the decision maker is more likely to consider a particular aspect first but not necessarily by explicit rank order), the results can be less than satisfying. If the aspects simply considered as they “come to mind,” the approach has no validity.

Sunk Costs

A rational decision, according to the criteria from Dawes, is based on the decision maker’s current assets, the consequences of the choice, and the evaluation of uncertainty in accordance with the laws of probability. A sunk cost is a past investment that has failed according to some standard of failure. To make a current decision based upon sunk cost is therefore irrational. Sunk costs deal with the past whereas rationality deals only with current assets and future consequences. Although honoring sunk costs is irrational, it occurs all the time:

“I am already here. I might as well go run that errand. I completely forgot about it. Oh well, it would be a waste of time and gasoline to have to come back.”

“We cannot cancel it now. We already paid the deposit. That would be a waste of money.”

“To terminate a project in which \$1.1 billion has been invested represents an unconscionable mishandling of taxpayers’ dollars.”

The last quote is from a senator in November 1981 [15]. He was responding to critics of a particular government program who had stated that the total value of the program, if carried on to completion, would be less than the amount of money yet to be spent completing it. The senator essentially responded by declaring that that to stop the program now would result in a deficit of \$1.1 billion with nothing to show for all the expense. However, the senator is justifying the expense of more money for the program to create something worth less than the money to be spent. This behavior is irrational. Note that if damage to reputation may occur due to canceling the program and it poses serious future problems, then it would be rational to honor the sunk cost. (Would any senator openly assert that the damage to his reputation outweighs the cost that taxpayers must shoulder to finish a doomed program?) Therefore, decisions should only be based on the future consequences of choices.

Framing

The framing effect found in decision making reveals a tendency of humans to make decisions based on their relative level of asset worth rather than on the total level of

asset worth. Consider the following scenario. There is a potential outbreak of bird flu in the Hawaiian Islands that is expected to kill 600 people in a local village. Scientists have proposed two alternative programs to combat the virus. They have estimated the consequences of the two programs as follows:

If program A is adopted, 200 people will be saved.

If program B is adopted, there is a $1/3$ probability that 600 people will be saved and a $2/3$ probability that no one will be saved.

Tversky and Kahneman posed this question to university students [15] and 72% chose program A. A second group of students was then given the same scenario, but with the two programs worded slightly different.

If program C is adopted, 400 people will die.

If program D is adopted, there is a $1/3$ probability that no one will die and a $2/3$ probability that 600 will die.

78% of this second group of students chose program D. However, program A and program C are identical because the concept of 200 people living is equivalent to 400 dying. Likewise, program B and program D are identical. The difference between the presentations of the programs resulted in an incredible 50% difference of acceptance (72% chose program A whereas only 22% chose the identical program C).

This behavior is irrational, and it has been termed the framing effect. The concept is simple. When a problem or choice between two options is presented in terms of gain (lives saved, money won, etc.), people tend to be risk-averse. They do not want to take chances if they know they are gaining something. On the other hand, if the problem is presented in terms of losses (deaths, debt, etc.), people tend to be risk-seeking. The irrationality is that the decision is not made upon the total outcome, but upon a relative outcome due to a presented frame of reference.

Tversky and Kahneman propose that the law of diminishing returns describes the human tendency to quantify the positive and negative consequences of decisions according to their own value function. These concepts are enumerated as follows: [15]

1. An individual views monetary consequences in terms of changes from a reference level, which is usually the individual's status quo. The value of the outcomes for both positive and negative consequences of the choice then has the diminishing-returns characteristic.
2. The resulting value function is steeper for losses than for gains.

Figure 3-8 depicts this value function and diminishing returns characteristic. The independent variable is money or the objective worth of the item, and the dependent variable is the subjective value as perceived by the decision maker. The diminishing returns characteristic is evidenced by the continually decreasing, positive slope away from zero (a negative second derivative) in both directions. Furthermore, note that the slope for 0^- is steeper than the slope for 0^+ , or in other words, the absolute magnitude of the slope near zero at 0^- is greater than the absolute magnitude of the positive slope near zero at 0^+ . Therefore, the interpretation of the value function

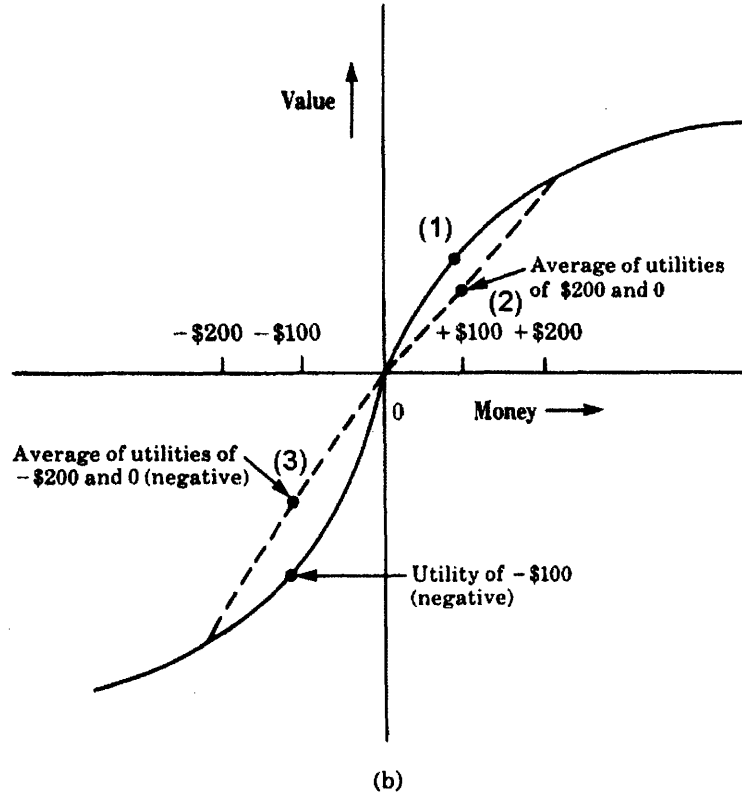


Figure 3-8: Displaying two aspects of prospect theory [15].

is that small positive increments in money away from zero (or from the frame of reference) maintain their total worth in terms of subjective value. However, the larger the positive increment of value away from zero, the less the perceived worth. The interpretation of small negative increments is the same, except that negative losses are perceived to carry greater negative worth due to the steeper slope. The earlier presentation of human failures in understanding probability presented a graph, Figure 3-5 that displayed subjective probability versus stated probability. In this figure, the curve was noticeably offset downwards, which means that the human's perception of probability is generally less than the truth. This offset depicts the framing effect. When a choice is presented in terms of gains, the probability of the riskier positive outcome will be underestimated due to the offset. Therefore, the gain associated with that outcome will be underestimated, and the decision maker will choose the sure option. When a choice is presented in terms of losses, the probability of the risky negative outcome will also be seen as less. The decision maker will then underestimate the expected loss, and therefore choose to be risk-seeking.

Figure 3-8 illustrates another framing effect where the upper right quadrant represents a choice given in terms of gains, and the lower left quadrant represents a choice given in terms of losses. For the positive choice, the scenario begins with the decision maker given \$100. This point is not shown on the graph because it is the

new reference level for the following decision. The decision maker must then choose between the option of receiving another \$100 or taking a 50-50 chance of receiving another \$200 or nothing. The sure gain of \$100 is point (1). Note that the expected outcome of the 50-50 chance option should be \$100. However, graphically it is the average of \$200 and 0 as found on the value function, given by point (2). Therefore, because of the diminishing returns characteristic, the average of the utilities (personal value) lies below the curve, and in terms of value, the decision maker opts for the risk-averse choice of a sure \$100. On the other hand, the same options are framed in terms of losses in the lower left quadrant, and it can be seen that the diminishing returns characteristic results in an average of utilities that lies above the curve, given by point (3). Therefore, the decision maker chooses the risk-seeking option of the 50-50 chance.

The choices to not wear a seatbelt, buy insurance, and play the lottery are three everyday examples of the application of Figure 3-8. For seatbelts, if the decision maker believes that the choice to not wear a seatbelt returns a small gain in comfort, that small gain is highly valued according to the value function. On the other hand, because of the diminishing returns characteristic, the large objective negative cost of an accident without wearing a seatbelt is not perceived as all that bad. For buying insurance and playing the lottery, the expectation is to lose money. Brenner notes that the reason for paying out insurance premiums and the cost of lottery tickets can be framed in terms of relative wealth distribution.

They [people] perform both acts for the same reasons: in both cases individuals expect to lose relatively small amounts, either the price of the lottery ticket or the insurance premium. But these small amounts are worth losing since these are the only ways by which people can either change or avoid changing their relative position in the distribution of wealth. Thus people gamble in order to try to become richer and change their relative position in the distribution of wealth, and they insure themselves in order to prevent becoming poorer, thus avoiding a change in their relative position. [15]

The relative frame of reference for the decision maker choosing to buy insurance is the net worth of positive assets. In terms of gain, the decision maker is risk-averse and would rather pay the sure loss than risk becoming poorer. For buying lottery tickets, the decision maker's relative frame of reference is negative, because he or she observes so many others as more wealthy. Therefore, in terms of this loss, the decision maker is risk-seeking and gambles to possibly become richer.

Finally, note that the earlier discussion of sunk costs can be understood in terms of framing. The decision maker clearly views the status quo as negative. Thus, the decision maker would rather keep spending the money in hopes that something good will come out of the already wasted effort rather than accept the loss and move on. The choice is between a sure loss of terminating the program and a gamble, and the gamble is perceived as having greater value. However, this commitment to honoring the past is precisely why companies bring in new heads of leadership. For the new leader whose reference level is not negative, the choice of terminating the

program or continuing to make a bad investment is easy. Considering only the future consequences of decisions, the new leader terminates the program for two reasons. First, the past history of the program holds zero value for the leader. Second, the future of the program is grim. The choice is between zero and a loss.

Overuse of Resources

In dynamic, uncertain environments, high level decision makers always want more information than can be had. There is never enough intelligence. Yet, is more necessarily better? To answer this question, Omodei et. al. used a high fidelity, interactive simulation that allows teams of humans to combat forest fires [51]. They focused particularly on the human subject chosen to be the firefighter commander and varied the amount of feedback information on the status of extinguishing the forest fires given to the firefighter commander. They found that when more resources were available to the commander, the commander felt obliged to take the time and effort to consider it, even when they were already cognitively overloaded. This applied to all kinds of resources, from information gathering, opportunities for action, and communication input. The perceived need to try and grapple with more information, opportunities, and communication in an already high pressure, high tempo, and uncertain environment significantly degraded performance. However, Omodei et. al. concluded that this bias to overutilize resources occurs outside conscious awareness and could most probably be the result of the activation of one or more schema in long-term memory (see Section 3.1.1). They proposed the following to explain the commanders' behavior:

1. A preference for errors of commission rather than omission. In most emergency situations, time is short, so there could very well be a general task bias for action over delay.
2. An illusory sense of greater control via activity...decision makers, to avoid subjective uncertainty, act in such a manner as to achieve an illusory sense of cognitive control over the environment. That is, activity regardless of its adequacy provides a sense that one is having some desirable effect in the problematic situation.
3. An illusory sense of greater self-competence via activity...to guard against a sense of personal incompetence.
4. An overestimation of personal ability... with respect to both speed of information processing and amount of information that can be concurrently managed in working memory.

Thus, due to a possible need for control and competence as well as an overconfidence bias (see Section 3.3.4), commanders would rather implement action, watch it unfold, and correct any mistakes rather than miss something. This suggests a failure in metacognitive activity, which is the awareness of one's own limitations. Humans do not have the ability to "optimally regulate" their cognitive activities because they are not consciously aware of their cognitive limitations.

Search Order

Nisbett and Wilson performed a series of experiments where they presented a rack of clothing to the subject [15]. All the clothes were in a single row facing the subject, and the subject was asked to state his or her preference for the article of clothing they desired. Nisbett and Wilson found that most subjects scanned the clothing from left to right, and then amazingly proceeded to choose the article of clothing on the far right. One explanation for this choice of the far right article is that in scanning from left to right, each new article of clothing might exhibit some desirable feature the previous one did not. Upon reaching the far right, there is no other article further to the right to be compared to, and thus it is chosen. This is not to say that subjects did not go back and look at other pieces of clothing, but that it is possible that the initial scan from left to right unconsciously produced a desire to choose the last article initially scanned. This experiment shows one main theme that the order in which items are presented is decidedly important. Humans are not parallel processors in comparing alternatives. The alternatives must be searched through serially. Furthermore, when combined with the limitations of working memory, not all the alternatives previously searched through can be held in working memory. Therefore, the order is of even greater importance. Moreover, the subjects in Nisbett and Wilson's experiments were presented with every possible decision choice. This is rarely encountered in everyday life. Thus, in making decisions, humans do not have all the alternatives, cannot hold all the alternatives in memory, and must search sequentially through all the possible alternatives.

Independence of Irrelevant Alternatives

When deciding between two options, A and B, the introduction of a third option C should not affect the decision between A and B. For example, if A has been chosen as more desirable than B and then another option C is considered along with them, A should still be more desirable than B. Option C has no relevance between preferring A over B. If option C is the most desirable in the set of A, B, and C, then ultimately C should be chosen. The question, then, is can the introduction of C be shown to reverse the preference between A and B even when C is not ultimately chosen? Such a reversal is contradictory and thus irrational. Yet, an experiment by Harrison and Pepitone displayed the contradictory contextual effects in decision making [15]. They asked students to train a rat through electric shocks. These shocks had different strengths. During one portion of the experiment, only two shock strengths were available. They were labeled "mild" and "slightly painful." In the second portion of the experiment, three shock strengths were available. They were "mild", "slightly painful", and either "moderately" or "extremely painful." The students were explicitly told not to use the moderately or extremely painful setting. Therefore, only mild and slightly painful shocks could be administered in both portions of the experiment. The results showed that the slightly painful shock was chosen 24% of the time when there was no third alternative, 30% of the time in the presence of irrelevant moderately painful level, and 36% of the time in the presence of the irrelevant extremely painful level. Therefore,

the contextual environment, though irrelevant to the decision, coerced the decision makers into making contradictory decisions.

3.3.5 Summary of Human Subjective Probability and Decision Heuristics and Biases

Humans exhibit an inability to properly apply the principles of probability theory. They would rather believe they could affect the outcome or derive some causality from a random sequence. Other failures in probabilistic reasoning have brought serious questioning as to whether the human expert should be relied upon to make intuitive predictions about the future. An analytical, quantitative, and deliberate analysis of a set of choice alternatives is slow and mentally taxing. Humans utilize heuristics to find quick, satisficing answers. For example, elimination by aspects quickly narrows down a set of large alternative choices to a few. However, because of certain biases, these solutions are not always appropriate. In the same example, if elimination by aspects proceeds according to the availability bias (the order of aspects is chosen as they come to mind), there is no validity to the solution. Deliberate analysis requires high cognitive effort. It is much easier and quicker to match important cues to representative and available information stored in long-term memory. The anchoring, overconfidence, and confirmation biases reveal a failure of meta-cognition (see Section 3.3), where the human expert is not fully aware of his or her own decision making limitations. How can the expert be trained to seek new hypotheses in ambiguous situations, to actively search for contradictory information, and to question the accuracy of his or her own situational awareness? Humans tend to get too attached to past mistakes and choose between risky options based on a present frame of reference rather than the final outcome. Due to the serial processing nature of human cognition and cue salience bias, the order of presentation and the attention-grabbing qualities of information is vital. Furthermore, human decision makers feel compelled to use and integrate all available information, and thus the amount of information provided to the decision maker also affects performance. Without proper filtering capabilities, irrelevant contextual influences bias the choice between relevant alternatives. In all, these misconceptions, heuristics, and biases appear to paint a bleak picture of human decision making. However, shortcut methods to quick decision making arise not from laziness but from necessity. In this research which attempts to capture tactical knowledge, such as reactive decisions, there is not time to find the best course of action. Thus, it is expected that some of these decision heuristics and biases will be found in these tactical knowledge elicitation experiments.

3.3.6 Naturalistic Decision Making

To summarize the chapter so far, we have presented the cognitive structures of long-term and working memories, an information-processing model of human problem-solving and decision-making, the concepts of normative decision making theory, and the human tendency to depart from these normative concepts by relying on heuristics and biases to make choices. One major conclusion thus far, is that in the real world of

uncertainty and time pressure, analytical and quantitative decision making methods appear infeasible. Furthermore, as the previous section described, humans are already prone to make decisions based on heuristics. The field of naturalistic decision making attempts to provide a stronger framework of real-world decision making rather than simply grouping departures from how decisions should be made under an umbrella of heuristics and biases.

Domain

Naturalistic decision making theory arose out of the seemingly impossible task of decision theorists seeking to reconcile normative decision making propositions with experimentally determined heuristics and biases all under the umbrella of a cognitive information-processing model. As always, when theories become so complex and so vast that individual parts now threaten its coherence, theorists return to the underlying foundations and seek simple answers. Klein did exactly this in 1998 when he observed the real time decision making of firefighter commanders. His goal was simple. How do firefighter commanders make decisions in the real world of high pressure, high tempo, and high risk? His work invigorated the decision making community to leave the laboratory and seek to understand decision making in real world situations. In doing so, the naturalistic decision making community has produced the following list of real-world characteristics that they seek to integrate into decision frameworks [45].

1. Ill-structured problems
2. Uncertain dynamic environments
3. Shifting, ill-defined, or competing goals
4. Action/feedback loops
5. Time stress
6. High stakes
7. Multiple players
8. Organizational goals and norms

One glance at this list confirms its accurate description of tactical decision making in battlefield environments. Indeed the military, especially the United States Marine Corps and the Swedish Army, has been seeking out naturalistic decision making theorists to train and re-teach decision making skills to its leaders because of its effectiveness in real world environments [25, 79, 59]. The military's traditional problem-solving models that rationally break out decision making into processes of problem definition, alternative generation, refinement, and selection are too time-consuming. Schmitt and Klein stated that, "these models are inconsistent with the actual strategies of skilled planners, and they slow down the decision cycle. As a result, the formal models are usually ignored in practice, in order to generate faster tempo" [79].

Recognition-Primed Decision Model

Figure 3-9 depicts Klein's proposed recognition-primed decision (RPD) model. RPD

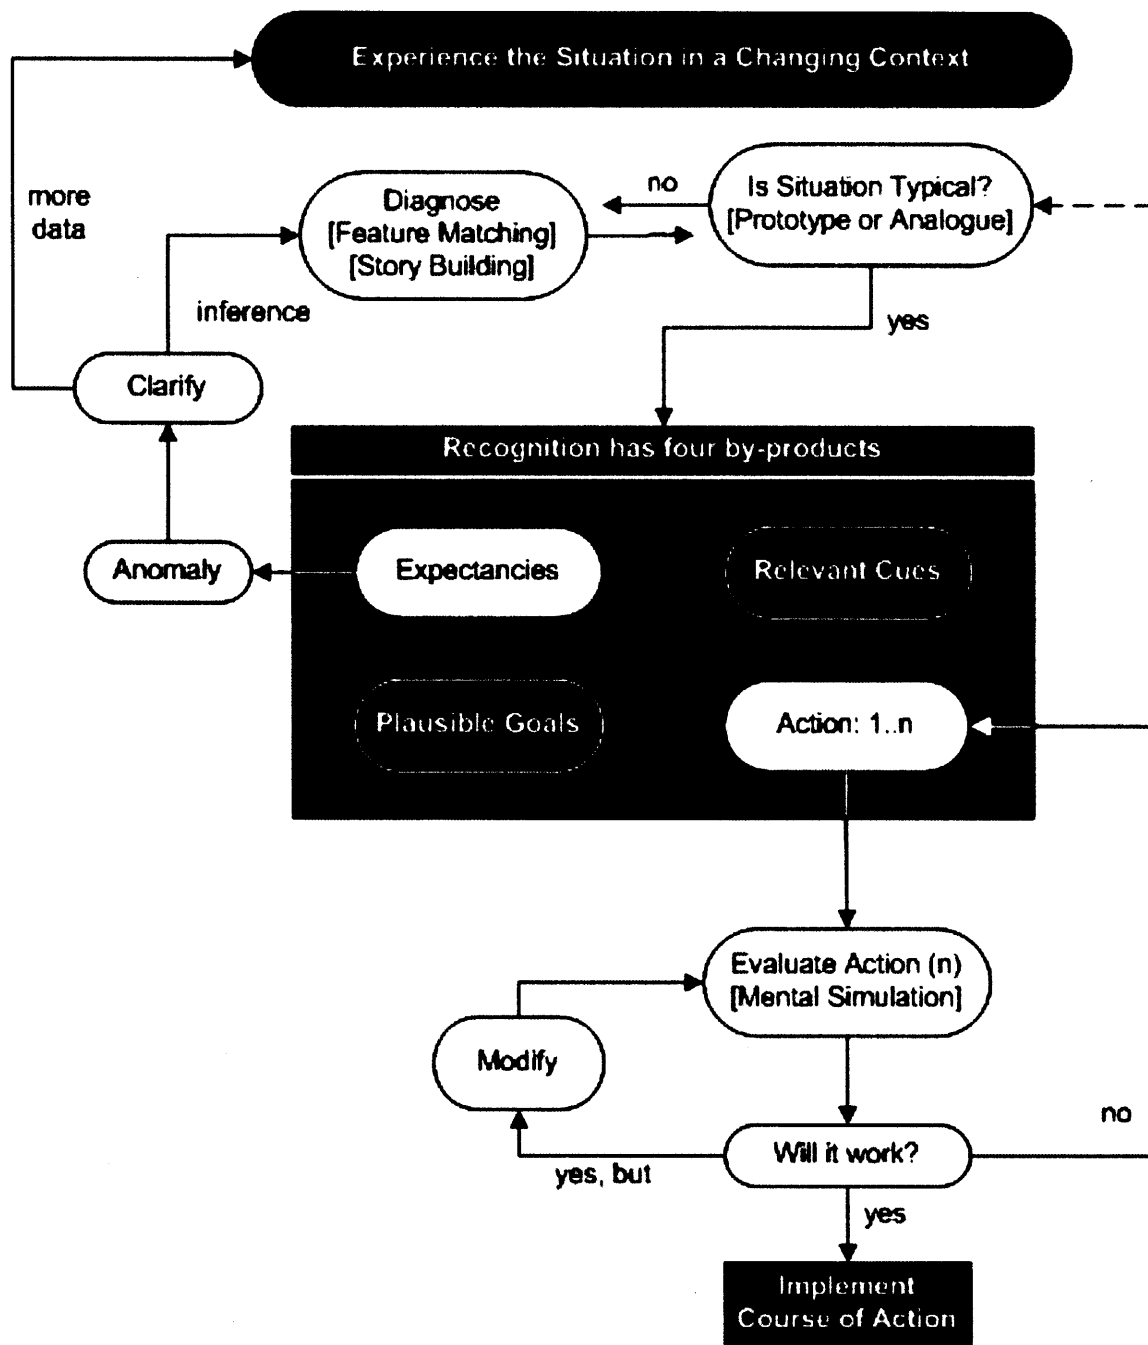


Figure 3-9: Recognition-primed decision model [38].

accounts for the two major aspects in decision making, situational awareness and choice of action. In experiencing the environment in a changing context, the decision maker begins with the following question: is the situation typical (prototype) or familiar (analogue)? If the answer is yes to either category, the decision maker categorizes recognition into four areas. The decision maker recognizes what goals are appropriate, which cues are important, what expectations to keep in mind, and what actions should be taken. By doing so, the decision maker sets priorities, filters the environmental information so there is not an overload, prepares for the next step and any surprises, and possesses a course of action. Situational awareness, then, is the combination of all four recognition by-products. This is markedly different from the decision maker beginning with a set of goals and expectations and then trying to interpret the situational context, which would quickly lead to misinterpretation and exacerbate heuristics and biases such as anchoring, confirmation bias, and overconfidence bias.

The upper left feedback loop in Figure 3-9 represents a situation in which the decision maker needs more information. In this circumstance, the decision maker may not recognize the situation as typical or familiar and therefore must spend time gathering more data in order to further diagnose the situation. At this point, the decision maker may attempt to match important features of the situation to alternative interpretations to find the best fit. This appears to be closely related to the elimination by aspects heuristic (see Section 3.3.4). Also, the decision maker may attempt to build a story that accounts for any missing features. However, this feedback loop does not just stem from an initial uncertainty in recognizing the situation. The decision maker may have misinterpreted the context and does not realize it until some expectancies have been violated. This anomaly also forces the decision maker to seek further clarification, more data, and a better diagnosis as described above.

The lower right feedback loop in Figure 3-9 represents a situation in which the decision maker has recognized the situation, chosen a course of action, but is not exactly sure of how that course of action will play out. For example, the decision maker may be trying to implement a tried-and-true course of action to an environment that possesses some novel characteristics. To understand if there may be any anticipated difficulties, the decision maker mentally simulates how the course of action will interact with the environment. In doing so, the decision maker could either modify the selected course of action or reject it and choose another. Furthermore, the mental simulation may highlight so many difficulties that the decision maker realizes more diagnosis is necessary.

There exists a simple rule-based analogy to RPD. For the simple scenario of recognizing the situation; forming a set of goals, cues, expectancies, and actions; and implementing a course of action, the model exhibits an “if ... then” response. For the upper left feedback loop where more information is needed, the model displays an “if (???) ... then” response. The decision maker no longer completely recognizes the situation and thus questions the initially formed antecedent. For the lower right feedback loop to evaluate a course of action through mental simulation, the model depicts an “if ... then (???)” response. The decision maker is unsure of whether the associated consequent is contextually appropriate and thus mentally simulates its

results.

Naturalistic vs. Normative

The following is a list of key features of the RPD model as presented by Klein (italics added for this thesis). The RPD model claims that with experienced decision makers:

1. The focus is on the way they assess the situation and judge it familiar, *not on comparing options*.
2. Courses of action can be *quickly evaluated* by imagining how they will be carried out, not by formal analysis and comparison.
3. Decision makers usually look for the *first workable option* they can find, not the best option.
4. Since the first option they consider is usually workable, they do not have to generate a large set of option to be sure they get a good one.
5. They generate and evaluate options *one at a time* and do not bother comparing the advantages and disadvantages of alternatives.
6. By imagining the option being carried out, they can spot weaknesses and find ways to avoid these, thereby making the option stronger. Conventional models just select the best, without seeing how it can be improved.
7. The emphasis is on being *poised to act* rather than being paralyzed until all the evaluations have been completed.

Notice the major departures here from normative decision making. Expert decision makers do not generate, evaluate, and compare alternatives. Rather, they find a workable option, imagine its implementation, and either look for improvements or new workable option. In the following quote, Klein compares RPD to normative, rational decision making theory. A “rational choice strategy” that seeks to overcome all of the human decision heuristics and biases and is mathematically sound requires the following:

[D]efine the evaluation dimension, weight each one, rate each option on each dimension, multiply the weightings, total up the scores, and determine the best option - that is, unless you do not have all the data you need, or are not sure how to do the ratings, or disagree with the weights, or run out of time before you have finished . . . The problem is that the assumptions of the rational choice strategy are usually too restrictive. Rarely is there the time or the information needed to make this type of strategy work.

Normative decision theorists, however, would proclaim that rational choice results in reliable, quantitative decision making. This is true, and Klein agrees. Furthermore, rational choice strategies, like subjective expected utility, may be helpful for beginners who are not experienced in a domain and have no long-term memories to draw upon. Yet, it simply hinders experts who must make real time decisions. Norma-

tive decision making calls RPD yet another heuristic. Naturalistic decision theorists dismiss normative decision making as neither natural nor practical.

RPD's Consistency and Applicability

RPD accounts for several major themes in this chapter. First, RPD accounts for the human tendency to match previous experiences from long-term memory with present scenarios. Second, it describes how the decision maker chooses within the constraints of working memory. The decision maker does not compare a list of alternatives. The decision maker does not even generate a list of alternatives, but devotes all of the power of working memory on one course of action at a time to mentally simulate its implementation. Third, RPD shows how the decision maker seeks a *satisficing* answer because time constraints do not allow further analysis. Therefore, not only does RPD account for the cognitive processes, limitations, and natural tendencies of pattern-matching and satisficing of a human decision maker, it also provides a framework drawn specifically from high pressure, high risk, high tempo operations. If tactical knowledge is to be learned from experts by observation in simulated environments, RPD should be the decision making model to answer the underlying questions of how the human expert is making the decision which will reveal the answer to the more important strategic question of why.

As a conclusion, the process of learning tactical knowledge can follow these steps, which correspond to the three levels proposed in the beginning of this chapter. First, we observe the actions and score the performance. Second, we use the think aloud reports and surveys to determine the human subjects' strategies. Third, to understand the process of the decisions, we search for the expectations, cues, goals, and actions involved in the human subjects' recognition of the tactical situations. We are most interested in the tactical strategies, because they provide the expertise to generalize tactical knowledge beyond the current simulated scenario. Both the actions and the decision making processes point to the strategies.

3.4 Other Cognitive Frameworks

Before finishing this chapter, we present two other cognitive frameworks and discuss the nature of both expertise itself and the ability to learn expertise. The two subsequent cognitive frameworks are first, the Belief, Desire, Intent behavioral framework, which has been given considerable attention in the AI community to model human agents (see Section 1.3.1). The second is the Generic Error Modeling System, which analyzes decision making from a human error standpoint. This is critical to learning tactical knowledge because we need to learn and understand both the good and bad tactical decisions, both successes and failures. They both provide equal importance to encoding the best tactical strategies in AVs.

3.4.1 Intentions, Plans, and Practical Reasons

Michael Bratman's book *Intentions, Plans, and Practical Reasons* proposed the Belief, Desire, Intent (BDI) model as a common sense, behavioral framework [7]. His in-depth treatment of intentional action, at first glance, appears to be more applicable to lawyers and fit for courtroom discussions not engineering theses. However, the AI community has embraced the BDI model as a framework for building human agents [3, 27, 48]. In fact, it has been so well accepted and applied over the past two decades that papers are now published asking if it has outlived its time and where to move on from here [27]. Furthermore, it appears that most designers of agent architectures can apply BDI to their architecture frameworks by simply casting the verbiage in the right form, although it was not conceived explicitly within the BDI framework. One example is a discussion of the SOAR architecture [27]. Yet, BDI is attractive for any form of modeling human behavior for three reasons. First, BDI specifies the use of planning to achieve some intended goal, which is equivalent to a means-end analysis put forth in Newell and Simon's general problem-solver. Second, these generated plans are both partial and hierarchical, and thus they branch off into subgoals as the agent moves towards its intended goal, again like the general problem-solver. Third, the agent's intention works as a filter to constrain the set of actions and subgoals it will pursue as it fills in the partial plans. These three attributes have allowed BDI-modeled agents to exhibit a coherent yet rich set of behaviors. As Pomerol notes, the beauty of BDI is that it presents a framework that seeks an equilibrium between reactivity and planning [55]. Therefore, the following sections present the BDI framework drawn exclusively from Bratman's book.

Intention

Intention is central to understanding humans and characterizing both people's actions and minds. As a general concept, intention frequently refers to the future. Humans intend to do something later. If Drew, a graduate student in Cambridge, MA, intends to eat lunch at the local thin-crust pizzeria Emma's, and is presently doing so, one would say he is no longer intending to eat but rather am now eating. This does not detract from the idea of present directed intentions, but rather emphasizes that intention typically refers to a future oriented commitment. Yet, what is the definition of intention? To refer to intentions as future-oriented commitments seems too vague to be applicable. Bratman argues that the nature of intention can only be determined in the context of planning. Before presenting that argument, there are three main objections to discussion of intentions as future oriented commitments. These are as follows. First, does intention, now, control actions, tomorrow, as if intention possessed some "ghostly hand over time?" Second, if another graduate student Kara intends, today, to take a flight to Boston, tomorrow, that intention should persist throughout tomorrow and help guide actions towards the desired end of flying to Boston. But is this intention irrevocable? Is there no other choice to take the flight even if unanticipated events arise between now and then that could possibly change the desire to go to Boston? That sort of total abandonment to prior intentions is

irrational. Third, if intentions, then, can be changed and are not permanent, “why should I bother deciding today what to do tomorrow?” Therefore, future directed intentions have been argued to be conceptually objectionable in the following three ways. They are “metaphysically objectionable” since they involve action at a distance. They are “rationally objectionable” since they are irrevocable. They are ultimately objectionable as just a waste of time. To counter these arguments, Bratman situates the concept of intentions in a planning framework.

Need for Planning

Humans are planning agents, and the need to plan for the future is rooted in two needs. First, because humans are rational, to some extent, they reflect and deliberate on what to do. If deliberation only occurred at the time of action, the ability of such deliberation to affect the action would be constrained. This is because deliberation requires time and resources and so would hinder the timeliness of the present action. Therefore, humans need to deliberate and rationally reflect on future actions due to resource limitations. Second, humans require both intrapersonal and interpersonal coordination. Intrapersonal coordination is prioritizing one’s own time and schedule to squeeze in all the demands of the day, including work, meetings, errands, etc. Interpersonal coordination is arriving and participating at a committee meeting with several other people. Furthermore, plans are never total and complete. The inability to predict the uncertainty and dynamic change of the future requires partial plans. Therefore, the incompleteness of plans requires humans to engage in some sort of reasoning that fills in partial plans with appropriate means, steps, and actions.

Plans \cap Intentions

What is the relationship between plans and intentions? “Plans . . . are intentions writ large,” and, “intentions are the building blocks of larger plans” [7]. Plans are partial and hierarchical. Because of the need for intrapersonal and interpersonal coordination, limitations of present deliberation, and ever-changing environment, partial and hierarchical plans are made. In order to be successful, plans must also be consistent and coherent. There are two consistency requirements. The plan, as a whole, should be internally consistent. If execution of one section of the plan hinders the execution of another section of the plan, it is not internally consistent and has no hopes of successful completion. The plan should also be consistent with the external world, given that the beliefs about the world are true. An internally consistent plan means very little if it cannot be successfully executed and completed within the world’s context. Finally, along with consistency, plans must be means-end coherent. The overarching intention guides the elaboration of plans with sub-plans of appropriate means, preliminary steps, and relatively specific courses of action.

The partiality and hierarchical structure of plans necessitate the ability to fill in high-level plans as needed. The requirement of means-end coherence forces the human to deliberate about which additional sub-plan or further intention among many should be chosen. Also, the requirement of consistency constrains which options

should be considered. Therefore, this framework of prior intentions and plans that must be consistent and means-end coherent helps define rational reasoning through two features. First, the establishment of “standards of relevance . . . provide a clear, concrete purpose for deliberation, rather than merely a general injunction to do the best.” Second, the use of a “filter of admissibility . . . narrows the scope of the deliberation to a limited set of options.” Therefore, Bratman argues that the only way a human can rationally choose low-level actions that contribute to high-level plans in a myriad - literally an infinite amount - of options is through a human’s prior plans and intentions, which are consistent and means-end coherent, because they determine which options are relevant and admissible. This has been termed *Bratman’s Claim* [27], and AI researchers have been particularly attracted to it precisely because it reduces the space in which agents have to choose actions. Intentions, in the AI community, then are a mechanism for constraining the set of options about which the agent must reason.

Filtering

Because this “filter of admissibility” is so important in implementation of rational agents, the following quote from Bratman elaborates on the concept.

This leads to the following understanding of the relation between the demand for strong consistency and the admissibility of new options. Consider a new option, *O*. Hold fixed the agent’s prior intentions, but add to the agent’s web of intentions and beliefs a new intention to *O*. Also add changes in belief that would be justified given that new intention, but without any other revision in the agent’s prior intentions. The option *O* is admissible if these changes in the web of intentions and beliefs would introduce no new inconsistency in that web. What matters for admissibility of a new option are one’s intentions *prior* to a decision concerning that option and the beliefs one would reasonably have *after* a decision in favor of that option. [7]

The graduate student, Kara, intends to fly to Boston tomorrow and has decided on when to leave, what route of transportation to take, has packed her luggage, and has arranged to be picked up from the arrival airport. Now Kara presently considers another intention to go to Loew’s Theatre tonight and watch a movie, do the beliefs along with that intention create inconsistency in my prior plans and intentions to fly to Boston tomorrow? If Kara believes that to watch the late-night movie might exacerbate her sleep-deprivation and potentially make her so tired as to miss her alarm in the morning, this would introduce inconsistency and be rejected. If, on the other hand, there is nothing in movie-watching that will hinder the intention to fly to Boston tomorrow, it is admissible. Therefore, the power in this framework is that it constrains actions through an admissibility filter, but not so strongly that the agent’s actions are mundane, simplistic, and limited. The framework allows very rich and complex agent behavior while maintaining consistency and coherence.

Conclusions from BDI

To summarize and conclude the discussion on BDI, consider the following quote by Bratman:

Practical reasoning, then, has two levels: prior intentions and plans pose problems and provide a filter on options that are potential solutions to those problems; desire-belief reasons enter as considerations to be weighed in deliberating between relevant and admissible options. This two-level structure is an essential part of the way in which intentions and plans play their coordination-facilitating role, and so part of the way in which intentions enable us to avoid being merely time-slice agents - agents who are constantly starting from scratch in their deliberations. So this two-level structure of practical reasoning has a pragmatic rationale, one grounded in its long-run contribution to our getting what we (rationally) want - given our limits and our complex needs for coordination. We need not leave a broadly instrumental conception of practical reason in order to allow intentions to have direct relevance to the rationality of action. [7]

Practical reasoning and rationality are related to intentions. However, note that never in this discussion has Bratman stated how to choose an initial intention. It is as if the problem has already begun, and the question is how to proceed rationally from here. Rational action is defined relative to an initial intention. Furthermore, Bratman does not propose how to weigh desires and beliefs when deliberating between admissible options, but that they simply must be weighed in accordance with consistency and coherence. Therefore, BDI is a framework that describes how humans with limited resources and needs for coordination must plan for the future through building blocks called intentions so that actions can be means-end coherent and consistent.

Return, now, to the beginning of this section on BDI and the statement of Pomeroy that the beauty of BDI is in its balance between reactivity and planning [55]. Does tactical knowledge reside in both reactivity and planning or just one or the other? We argue for both and have already identified reactivity as tactical knowledge we wish to learn (see Section 1.2). Planning is also a necessary element of tactical knowledge because the battlefield environment is characterized by dynamic change and uncertainty, which are the precise reasons Bratman identifies as to why humans must plan (see Section 3.4.1). The humans' intentions guide the balance between reactivity and planning, and thus discovering the human experts' intentions is crucial to learning and applying tactics.

3.4.2 Generic Error Modeling System

The last cognitive model to be discussed in this chapter is the Generic Error Modeling System (GEMS) by J. Reason [57]. The importance of this model is three-fold. First, its design is based on a cognitive approach to understanding human error. This is quite unique and helpful in seeking to understand both the right and wrong actions as displayed by human experts. Second, it classifies human performance

into three cognitive levels. These levels are consistent with the human's default desire to match present information with past experiences so that quick decisions can be made with minimal effort. Third, it comes directly from Rasmussen's skill-rule-knowledge framework which "has effectively become a market standard within the systems reliability community" [57]. Figure 3-10 depicts the flowchart of GEMS.

As seen by Figure 3-10, there are three levels as described from Rasmussen's framework. At the skill-based level, the decision maker is acting in a routine manner in familiar environments and only occasionally makes checks on progress. These are actions performed out of complete intuition. Conscious thought is not required at this level other than during the attentional checks. For example, the pilot of an aircraft does not have to consciously think during takeoff about how to pull back the stick at rotation speed, when to raise the flaps and gear, or how to maintain a constant airspeed, rate of climb, and direction heading. These are all skill-based actions performed so many times, that they are deeply embedded in long-term memory. As long as system monitoring indicates normal operation, all these actions are automatic. However, once an instrument, for example, indicates the gear did not completely raise, the decision maker moves into the rule-based level. Now the pilot must gather as much local state information as possible to diagnose the error. If the problem and local state information indicate a pattern previously encountered, the decision maker applies the stored rule that matches the pattern. Note that this does not only apply to trouble-shooting an error. All decision making can be cast in the form of solving a problem. Thus, this also applies to the chess player seeking after the right move. There are certain rules that a chess player may abide by, such as the familiar "maintain control of the center of the board" rule or knowing when to castle or move the queen out. Yet, if the context is completely novel, the decision maker moves out of the rule-based level to the knowledge-based level. Even after moving into the knowledge-based level, the decision maker may continue to seek a higher level analogy between some features of the current situation and past experiences to apply some stored rule. If this fails, the decision maker must expend a high level of cognitive effort to create a mental model of the problem space, analyze the relevant cues, diagnose the situation, and formulate and apply corrective actions.

GEMS depends entirely on the theme, discussed throughout this entire chapter, that when confronted with a problematic situation which requires a choice, humans first look for "prepackaged" solutions. As Rouse said, "humans, if given a choice, would prefer to act as context-specific pattern recognizers rather than attempting to calculate or optimize" [60]. Thus, GEMS asserts that the default level of problem solving by human preference is the rule-based level. That is why two feedback loops from the knowledge-based level lead back to the rule-based level. In fact, humans display a "rigidity" in trying to apply rule-based solutions even when the situation does not warrant it. For example, the famous jars test by Luchins and Luchins 1950, which involved over 9,000 adults and children, proved the "blinding effects of habit" [58]. In the experiment, the human subjects were given a series of problems based on three water jars of different sizes, as displayed in Figure 3-11. Those jars were to be used as measuring jugs. The problem began by specifying how much water could be held in each jar, and the goal was to end up with a specific amount of water in one

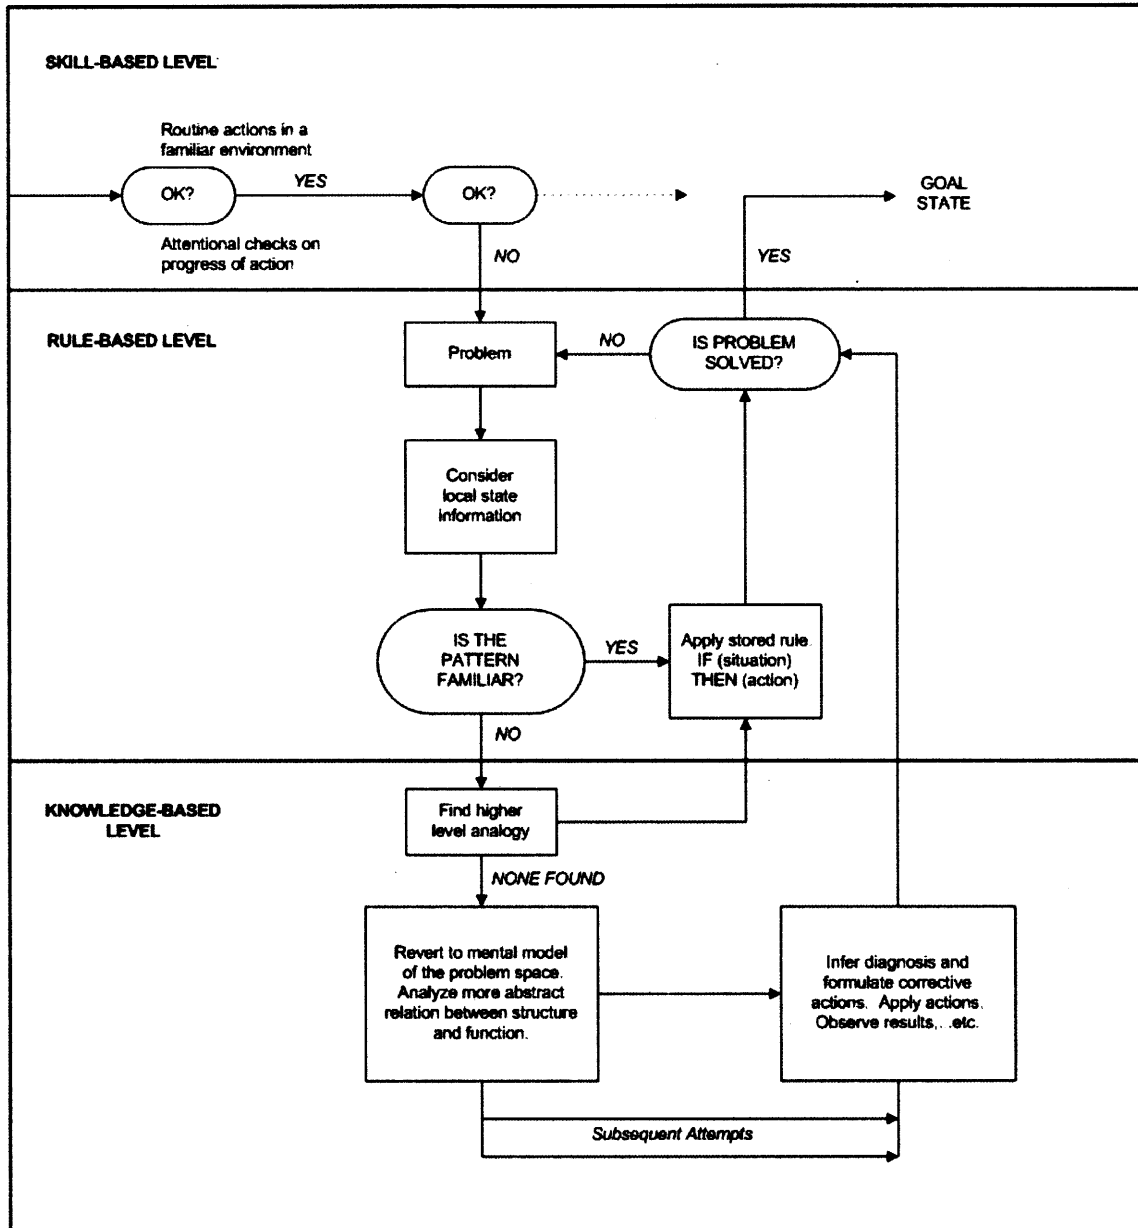


Figure 3-10: Rasmussen's performance level framework and the Generic Error-Modeling System [57].




Figure 3-11: Luchins' water jars problem [58].

jar. For example, if jar A can hold 18 liters, jar B can hold 43 liters, and jar C can hold 10 liters, how do you end up with 5 liters? The solution is to first, fill jar B and empty into jar A. This leaves 25 liters in jar B. Then, fill jar C by jar B, empty jar C, and repeat. This leaves 5 liters in jar B. Luchins and Luchins presented several problems that were solved by this rule $B - A - 2C$. In the last problem, however, the solution was much simpler. In this problem, if jar A can hold 28 liters, jar B can hold 76 liters, and jar C can hold 3 liters, how do you end up with 25 liters? The answer is to fill jar C with jar A, and jar A now contains 25 liters. However, subjects who had correctly induced the previous rule of $B - A - 2C$ had difficulty finding this much simpler rule of $A - C$ because they kept trying to incorrectly apply the more complex one to this simpler problem. In fact, it took these subjects significantly longer to solve the $A - C$ problem than another control group of subjects who were simply given the $A - C$ problem. The test proved that humans stubbornly activate familiar rules which are overly cumbersome and complex. In the words of the experimenters, past experience ceases to be a tool, "when, in a word, instead of the individual mastering the habit, the habit masters the individual" [57].

GEMS is a useful framework in evaluating tactical knowledge because it deals with human error, and we are interested in learning both why successful actions worked well and why failed actions did not. Thus, after differentiating between good and bad actions through scoring metrics, we cast the context of the bad actions in the form of problem solving and apply the GEMS framework. At what level did the human expert attempt to solve the problem? What were the stored rules that were applied? Did failures occur in forming the wrong mental model of the problem space? The answers to these questions will all help determine the tactical strategies.

3.5 Expertise in Cognitive Processing

To conclude this chapter, we present how expertise is found not only in superior actions but also in superior cognitive processing. Together, the actions and processes point to expert strategies. Yet, expertise in cognitive processes can still suffer from heuristics and biases. Thus, we finish the section with six problems associated with learning from expertise, how they emphasize observation of human performance

rather than simple verbalization of expertise, and their ability to affect the results of conducting human-in-the-loop experiments to determine human tactical knowledge.

3.5.1 Experts versus Novices

Experts possess a larger knowledge base in their area of expertise than a novice, by definition. Furthermore, experts can solve more difficult problems in their area of expertise than a novice, presumably because they possess a larger knowledge base. Long-term memory is assumed to have an infinite capacity for storage, and retrieval processes are very rapid, parallel, and independent of memory size. Experts, then, have more extensive knowledge structures in long-term memory that have also been more frequently accessed. Yet, any expert must cope with the same working memory limitations of all humans. Though an expert may be able to operate mostly in the rule-based level of performance when solving problems in the domain of expertise, if an expert must resort to a knowledge-based level of problem solving, the cognitive effort and limitations are no different than the novice who must begin at the knowledge-based level. Why does it seem, then, that experts, even when confronted with novel situations, solve problems quicker and more accurately than novices? There are two answers. First, the acquisition of expertise is linked directly with the ability to associate and package larger and larger amounts of information into recognizable patterns. Second, expertise is also directly related to cognitive mediation of performance, in that the expert exhibits greater cognitive control in solving new problems.

The most well known and often cited example of high level pattern matching is the analysis by Adriaan de Groot and later by Simon and Chase of expert chess players [15, 18, 47, 57, 87]. These researchers found that the expertise of chess grandmasters did not reside in analyzing every possible move out to five turns and weighing out the benefits of each approach. In fact, eye monitoring equipment found two surprising facts. Most of the time the grandmaster ultimately chose the move he or she first considered. Second, the grandmaster only looked, on average, two or three moves ahead. Therefore, it appeared as if the grandmaster simply knew the right move from the beginning and took a little bit of time to double-check that no other option appeared better. To validate the hypothesis that pattern recognition was the mechanism behind this intuition, expert chess players were presented with a hypothetical chess board taken from real games for a period of time and then asked them to recreate all the pieces on the board. Amazingly, after only five seconds of observation, chess grandmasters could accurately reconstruct the board up to 90%. Novices had no chance of such accurate and rapid observation and recreation. Even more intriguing was that if the pieces were placed randomly about the board, rather than as depicted from real games, the grandmasters' ability to recreate the board dropped quickly to the level of novices. This was not photographic memory but high level abstraction and association. Simon and Chase summarized by stating, "the most important processes underlying chess mastery are ... immediate visual-perceptive processes rather than the subsequent logical-deductive thinking processes" [68]. Thus, expertise resides in the ability to integrate larger and larger amounts of information and data into recognizable chunks stored in long-term memory for rapid retrieval. When solving a new

problem, experts must expend the same amount of cognitive effort in analysis and diagnosis as novices. However, in working out the problem, experts quickly identify previous patterns and can start employing long-term memory retrieval to speed up the search for the solution.

Expertise is not static. There might exist some standard whereby a novice can finally perform as well as an expert, but expertise does not plateau. If it did plateau, this would imply that everything that can be learned has been learned and expertise could then solely reside in automatic retrieval of correct actions. However, it is not, and the expert is still required to learn and to maintain problem solving abilities through slow and deliberate analysis and inference. In his analysis of chess players, De Groot also found that even chess grandmasters discovered better moves through mental simulation. Thus, “the performance of experts is mediated by increasingly complex control processes” [18]. Experts are better positioned by these increasingly complex control processes to react to changing environments and anticipate future consequences. It is all part of a continual process of learning, applying, and evaluating.

3.5.2 Problems with Learning from Experience

It seems intuitively obvious that the best way to learn the rules and skills involved in a new domain is to learn from a domain expert. Yet, there are precautions that must be taken into consideration due to the nature of human cognition (see discussion on heuristics and biases in Section 3.3.4). The following are six problems with learning from experience as given by Dawes [15]. First, decision makers will overestimate the probability of certain salient memories or salient features of past experiences due to the availability heuristic. Second, the set of experiences for the expert decision maker is always a subset in that domain. The F-16 pilots who would not fly near unmanned vehicles do not rationally relate their condemnation of dually-occupied manned, unmanned airspace to the set of all events of manned and unmanned aircraft operating in the same airspace. They just quote the one eyewitness account of a mid-air collision. Dawes calls this the “biased generation of experience.” Third, the dynamics of changing environmental conditions call into question the relevancy of past experiences where the conditions were different. This is the “superannuation problem,” and it is a main concern in engineering. How well do the results apply or generalize to the dynamics of the real world outside of the simulation environment? It is a question of robustness of the solution method. Fourth, the structure of the long-term memory, composed of schemata, tends to reconstruct experiences biased towards current attitudes and beliefs. It is what Bartlett described as “effort after meaning” [2]. George Valliant writes, “It is all too common for caterpillars to become butterflies and then to maintain that in their youth they had been little butterflies. Maturation makes liars of us all” [15]. Present attitudes and expectations bias the reconstruction of past experiences. Fifth, expert decision makers sometimes fail to appreciate the randomness in their experiences. They would rather infer some causality due to their choices rather than accept chance in the details. Sixth, due to the hindsight bias, expert decision makers who can review the details of a previous event believe that they can accurately predict its occurrence again. However, they are now too quick to

associate current details with that past event, and, closely related to the availability heuristic, they overestimate the probability of occurrence.

All of these six warnings in learning from experience underscore the need for observing the expert's performance within his or her domain. Simply interviewing the expert and asking them to verbalize their expertise through questions pertaining to specific past events, questions that seek broad interpretation, or mentally stepping through a simulated task allow these cognitive biases to taint the results. At least in a real time, simulated environment, the subject matter expert must demonstrate expertise by performance rather than simply by memory recall. There is still the question of the subject matter expert misapplying past experiences to present actions in the simulator due to the above problems. That is why the performance should be scored to help filter out good from bad actions and combined with the cognitive processes to find the best tactical strategies.

3.6 Conclusions from Cognition

The hope of this chapter is that, though long, it would provide a brief, yet thorough overview of the engineering psychology efforts to understand cognitive processes in decision making. The notion of learning tactical knowledge is an ill-defined problem, yet human experts find solutions. A qualitative analysis of experimental results requires interpretation, and this chapter has sought to illumine some of the objective theoretical decision making frameworks. The consistent application of this chapter to the research are the three levels in understanding tactical knowledge of actions, strategies, and processes. It is tempting to conduct experiments and move straight from actions to strategies. Not considering the cognitive processes, however, has several consequences. First, the experimental setup could unknowingly bias the human subject towards certain information or actions, such as with the framing effect. Second, too much interpretation may be used and too much read into the humans' performance to move from actions to strategies, when a simple decision heuristic or bias provides sufficient explanation. Third, the human subject may verbalize the same strategy in two cases, yet perform terrible in one and very well in the other. This could occur if the human subject, for example, selectively attended to a different cue in the second scenario. Therefore, utilizing cognitive frameworks and being aware of decision heuristics and biases bring objectivity both to the experimental setup as well as the data analysis in learning tactical knowledge.

Chapter 4

Procedures and Methodology

The emphasis in Chapter 2 was to answer the question of what are the right kinds of expertise to learn from experts and how they should be carefully applied in automation design to achieve required levels of reliability. The focus of Chapter 3 was to answer the question of whether it is possible to learn human expert knowledge, how to model it, and the limitations and mechanisms involved in human decision making. In Chapter 4, we present the procedures and methodology for this research's effort at learning and applying human tactics.

In this chapter, we begin with discussion of the experimental framework. Next, we describe the human-in-the-loop experimental setup, including the chosen scenarios, simulation environment, experimental visual display, human-simulation interface, software, and a brief overview of the human subjects employed in the experimentation. Then, we present human factors considerations that must be kept in mind for all human-in-the-loop experiments by discussing training effects, think aloud reports, use of surveys, and the limited number of cases allowed in the experiments. Next, we describe the timeline and the presentation of the experimentation to the human subjects. After that, we outline in detail the governing equations and behavior of the vehicle and enemy simulated entities and the interactions between them. We then present the baseline autonomous vehicle (AV) statechart derived without human subject inspiration. Finally, we conclude by discussing the limitations in the experimental procedures and methodology.

4.1 Experimental Framework

In order to test the concept of learning tactical knowledge from human subject matter experts and applying that knowledge to AVs, the proposed experimental framework is as follows. First, we identify scenarios wherein the AV's ability to complement manned missions is considered beneficial. Next, we develop a human-in-the-loop simulation environment to model these scenarios whereby we can design and change a set of independent variables to force various decision points. Then, we take several human subjects through a set of experiments and surveys to collect tactical, decision-making data. After collecting the data, we filter the good from bad decisions by performance

metrics. In parallel with these human experiments, we design an autonomous vehicle whose behavior is controlled through a statechart diagram. Initially, this AV's behavior is strictly baseline, that is without any human inspiration other than the common sense of the human designer. The baseline AV flies through the exact same scenarios as the human subjects and scored with the same performance metrics. Then, we compare both the performance between human subjects as well as between the baseline AV and human subject performance. Next, we find candidate, human-inspired tactics and encode them into the baseline statechart diagrams. In order to ensure the candidate tactic performs well above and beyond the small number of training data it was taken from, we employ Monte Carlo simulation and test the baseline and improved AV against a large number of scenarios with randomly varying parameters. Finally, we compare the baseline and improved AV performance and add those tactics to the statechart playbook that exhibit performance improvements.

4.2 Human-In-The-Loop Experimental Setup

The human-in-the-loop experiments took place in February 2006. Between the two rounds of experiments, five human subjects were given fifteen total scenarios and two surveys to complete over the course of two hours. The details of the experimental setup are given in the following sections.

4.2.1 The Scenarios

In choosing a scenario, we focused on identifying those tactical situations where the inclusion of an AV appeared beneficial. After discussing the research's focus with a U.S. Army officer and UH-60 Blackhawk helicopter pilot (who was later included as a human subject) we derived the following list of candidate scenarios [69]:

1. air corridor reconnaissance
2. target lazing
3. suppression of enemy air defense
4. battle damage assessment
5. finding and assessing landing zones
6. communications relay
7. forward-looking urban terrain scout

First, air corridor reconnaissance is needed when a company of troops are being moved to another location by utility helicopters. The specific route to the drop-off point, the air corridor, must be searched and cleared of enemy contacts to ensure the safe passage of the troops. Second, in target lazing, the AV can take the place of the special operations ground forces who must manually laze a target so that another platform may lock on and fire a precisely guided missile to the target. Third, suppression of enemy air defense, for example seeking out and destroying surface-to-air missile (SAM) sites, is crucial to many modern engagements but extremely dangerous. Fourth, after a conflict, someone must perform battle damage assessment to get an

accurate picture of the casualties for future intelligence. Fifth, just as the air corridor has to be searched and cleared for the passage of troops, helicopter landing zones must also be searched and cleared. The helicopter is most vulnerable during landing. Sixth, for long convoys and supply lines, ground vehicles in the rear of the line are not able to communicate to the vehicles in front with line-of-sight radio equipment. Instead of consuming bandwidth for satellite communications, an AV that maintains flight above the middle of the supply line can act as a communications relay platform. This would be especially helpful if sections of the convoy or supply line run into obstacles that must be circumvented. Seventh and finally, modern control system design for small autonomous helicopters have taken advantage of the large thrust-to-weight ratio and agility of these platforms [26]. A small, agile helicopter that can maneuver within tight spaces would be invaluable to troops conducting urban warfare.

Of all these missions, the air corridor reconnaissance mission was chosen for two main reasons. First, the mission allows for the inclusion of all three desired elements of designing scenarios to test and learn human tactics, as discussed in Chapter 2. These elements are uncertainty, a hierarchy of objectives, and flexibility in mission completion so that the human expert can experiment with and discover novel problem-solving techniques. They are accomplished by forcing the human to find a balance between planning and reaction, between the nominal mission of searching over terrain for enemy contacts and reacting to the presence of an enemy. Second, the reconnaissance mission has been and continues to be the primary role of AVs. Table 4.1 displays the historical reconnaissance roles AV's have fulfilled since their inclusion to battlefield operations [50]. Note that the Pentagon has changed its terminology to Unmanned Aerial Systems (UAS), as seen in this table. AVs have been tasked with reconnaissance roles for all echelons of the military and all branches. In the Unmanned Systems Roadmap 2005 from the Office of the Secretary of Defense (OSD), the military's combatant commanders ranked the importance of eighteen missions across the following four classes of AVs: small, tactical, theater, and combat AVs. For every class of AV, the reconnaissance mission was ranked as the most important mission. In fact, the reconnaissance mission for the combat AV ranked higher than the strike mission itself [50]. Furthermore, the Unmanned Undersea Vehicle (UUV) 2004 Master Plan of the Office of the Secretary of the Navy identified the intelligence, reconnaissance, and surveillance (ISR) mission as the highest priority for UUVs [77]. Thus, the air corridor reconnaissance scenario was chosen not just because AVs only perform reconnaissance, but because smarter, armed reconnaissance missions, such as ISR tasks, are the next vital step in the continuing evolution of desired missions for AVs.

4.2.2 The simulation environment

An initial effort in designing and building the simulation environment focused on developing a six degree-of-freedom (6-DOF) helicopter model in Simulink [42]. The 6-DOF model implemented the oblate, rotating earth equations for any rigid aircraft [72]. The helicopter forces and moments were modeled by thrust vectoring of the main propeller, where the joystick inputs controlled the blade flapping angles, β_{lat} and β_{lon} ,

Brigade/Division Asset for Reconnaissance, Surveillance, and Target Acquisition	
<i>Proponent</i>	Army, Marine Corps
<i>Heritage</i>	Falconer (1950-60s) - Aquila (1970-80s) - Pioneer (1980-2000s) - Dragon Drone (1990s) - Outrider (1990s) - Shadow 200 (2000s)
Shipborne Asset for Reconnaissance and Weapons Support	
<i>Proponent</i>	Navy
<i>Heritage</i>	DASH (1960s) - Project Blackfly (1970s) - Pioneer (1980-2000s) - Fire Scout (2000s)
Small Unit Asset for Over-the-Hill Reconnaissance	
<i>Proponent</i>	Marine Corps
<i>Heritage</i>	Bikini (1960s) - Pointer (1980-90s) - Dragon Eye (2000s)
Survivable Asset for Strategic Penetrating Reconnaissance	
<i>Proponent</i>	Army / Air Force / Navy
<i>Heritage</i>	Osprey (1960s) - D-21 (1960s) - Classified Program (1980s) - Dark Star (1990s) - J-UCAS (2000s)
High Altitude Endurance Asset for Standoff Reconnaissance	
<i>Proponent</i>	Air Force
<i>Heritage</i>	Compass Arrow (1960s) - Compass Dwell (1970s) - Compass Cope (1970s) - Condor (1980s) - Global Hawk (1990-2000s)

Table 4.1: Historically Validated Unmanned Aerial System (UAS) Roles [50]

and the magnitude of the main propeller thrust. This model was slightly unstable and some work was begun to add low-pass, notch, and high-pass filtering to make it controllable. The next version of the 6-DOF model scaled back the inputs from blade flapping angles to kinematic inputs only. Now the joystick controlled pitch rate q and roll rate p directly, and the magnitude of the thrust only acted in the $-z$ direction in body-axis coordinates. This model was controllable, and saturation blocks provided realistic dynamic constraints.

The original hope was that this helicopter model could interact with prepackaged enemy models, terrain databases, and visualization software through a distributed simulation network. This effort was progressing forward, but the time required to research, test, and acquire the best licenses for these databases and software tools took too long and appeared to leave minimal time for actual experimentation. A tradeoff was necessary between fidelity and timeliness [64]. Therefore, this three-dimensional distributed simulation environment was put on hold. Instead, the simulation environment changed from three-dimensional to two-dimensional and became self-contained in Simulink rather than distributed over a network. Enemies had to be designed and added to the simulation model, and predetermined interactions between enemies and the vehicle had to be encoded. Furthermore, the display for the human-in-the-loop had to be designed to exist within Simulink. Simulink does have a virtual-reality toolbox, but the learning curve to implement it would have negated any time-savings that this two-dimensional effort sought after. Thus, the best solution was to include a block in Simulink that sent the necessary variables into the MATLAB workspace

to be plotted by a MATLAB script at each time step.

The simulation ground terrain environment comes from the loading of a Level 1 Digital Terrain Elevation Data (DTED) file. The entire file covers a 1° of latitude by 1° of longitude portion of Kosovo. The southwest corner of the terrain is the intersection of latitude 42°N and longitude 021°E . Mountainous regions of this terrain were pulled out and trimmed down into several different smaller terrain patches that could be used for scenarios. This database was originally procured to be used in the three-dimensional simulation experiments. However, when the decision was made to conduct experiments in a two-dimensional world, we did not wish to abandon this terrain database. Thus, to operate in two dimensions, all platforms are restricted to operating within one altitude plane of the three-dimensional terrain. The environment is displayed from a “Global Hawk’s-eye-view,” where the camera is above the terrain matrix and looking straight down onto it.

4.2.3 Display, Interface, Software

Figure 4-1 depicts the visual display of the simulation environment as seen by the human subject. There are several features in this display to be discussed. First, there are two windows, or subplots to the display. The right window depicts the entire map, the vehicle, the air corridor, waypoints, and the location of any enemies encountered during the scenario. The left window is a zoomed-in picture, centered on the vehicle. In this window, the vehicle’s two radars are displayed (see Section 4.3.2). The first is the primary sensor for searching over terrain and detecting any enemy contacts, which is circular. The second is the weapons radar for firing upon enemy contacts, which is a cone extending to either side of the vehicle’s heading marker. Also, in this left window, enemy contacts are displayed with their primary, circular radars. Note that an enemy contact is only displayed with its primary circular radar. The enemy’s weapons cone and heading marker are never displayed but must be learned and inferred by the human subject. Finally, underneath the right map window is a simple counter that displays the number of weapons rounds remaining that can be used by the human subject.

There are several visual cues to help the human subjects differentiate between enemy contacts, between progressing stages of engagement, and the health of the AV and the enemies. Figure 4-2 displays these visual cues. Section 4.3.2 discusses the logic to progress from full health and no detection, through the steps of engagement, to firing upon and hitting an enemy.

Figure 4-3 displays the experimental setup with the visual display, joystick for vehicle control inputs, and microphone for recording verbal data. The human subjects interacted with the Simulink model through a Microsoft Sidewinder joystick. Moving the joystick side-to-side changed the vehicle’s heading angle θ , and moving the joystick forward-and-back in combination with the joystick’s lever changed the vehicle’s velocity V along its heading. Also, the TELEX PC microphone was used for voice recordings of the human subjects [76]. Finally, the software Camtasia Studio 3 recorded an integrated audio/video file for each experiment [73]. This software captured the human subjects’ spoken “think aloud” reports through the microphone

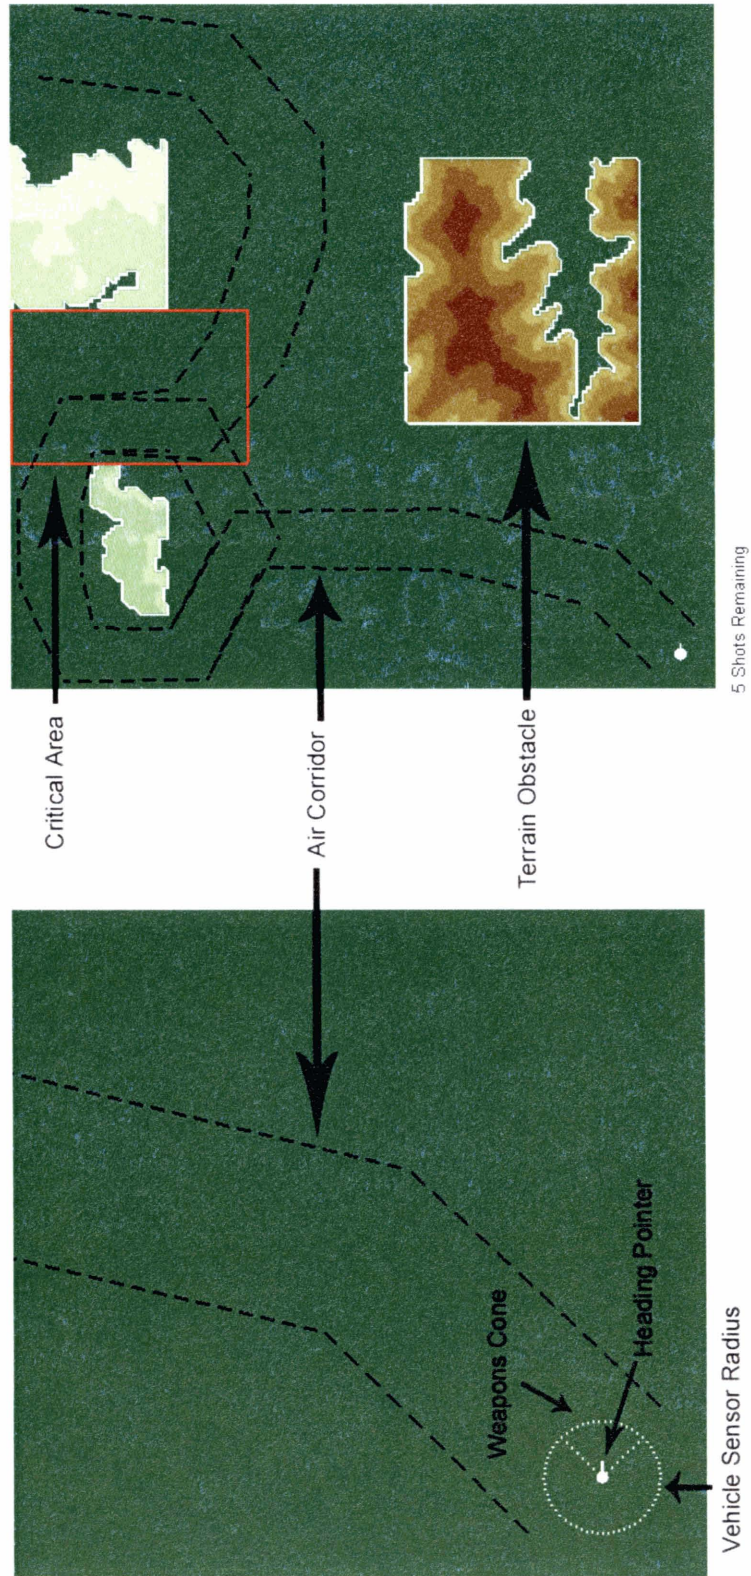


Figure 4-1: Two-window visual display of simulation.

	Normal	Detected	Locked	Damaged	Destroyed
Your Vehicle	○	○	●	●	●
Tank	□	□	■	■	■
S.A.M.	△	△	▲	▲	▲
U.A.V.	◇	◇	◆	◆	◆

Figure 4-2: Visual cues.




Figure 4-3: Human-in-the-loop experimental setup.

located near the joystick. It also recorded the simulation display, as shown in Figure 4-3 in the center screen. Most of the variables during each experiment were recorded by MATLAB and stored in data files, including the joystick inputs.

4.2.4 The Human Subjects

The original aim of the experiments was to first develop a simple simulation environment and test out the concept of learning tactics from humans. Then, the process would be tested on human experts in a higher fidelity simulation environment. However, due to the high demand of both military pilots and dedicated simulation facilities for training and testing, we could not finish the second step in the desired experimentation process. Therefore, the human subjects in the knowledge elicitation experiments were five graduate students at the Massachusetts Institute of Technology (MIT). Four out of the five are military officers, one U.S. Army, two U.S. Air Force, and one U.S. Navy. The fifth is a civilian whose father is a retired Colonel in the U.S. Air Force. All except the U.S. Army officer had come to MIT straight from their undergraduate education. The U.S. Army officer is a UH-60 Blackhawk helicopter pilot who came to MIT after spending a tour of duty in Iraq during the initial phases of Operation Iraqi Freedom.

4.2.5 Human factors considerations

Training Effect

In conducting human-in-the-loop experiments, there are human factors considerations and some specific guidelines to ensure sound methodology in data collection. First, the ordering of the scenarios presented to the human subjects should be randomized to avoid the training effect. The training effect occurs when the subject is exposed to more than one variation of the same scenario. Although some variables have changed, the human subject now has a larger knowledge base to accomplish the new missions because of the previous simulation runs. There is no perfect avoidance of the training effect. The purpose of randomization of scenarios between subjects (subject 1 is presented the cases ABC, subject 2 is given BCA, and subject 3 is given CAB) is to average out any training effects which occur from person to person [64]. Section 4.4.7 discusses the balance between the training effect and desiring human subjects to learn expertise.

Think Aloud

The purpose of having a human subject “think aloud” while participating in an experiment is to understand the thought processes occurring in working memory. These thought processes are not stored in long-term memory, and thus, after the the experiment has finished, the human subject cannot remember or access this crucial data that occurred in the past. The human subject is asked to access his short-term memory and simply verbalize those thoughts. Think aloud is a continuous verbalization of what the human subject is thinking at that particular moment [4, 29].

It has been utilized in many human experiments, such as Newell and Simon's work, where the verbal reports provided the foundation for their general problem-solver theory [47]. Ericsson and Simon were the first to provide a rigorous theory of using think aloud reports in data collection in their book titled *Protocol Analysis: Verbal Reports as Data* [20]. This book lays the complete theoretical foundation for using think aloud techniques as concrete and justified data collection. It also provides techniques for organizing and classifying the think aloud reports. However, they freely admit that it is an extremely time-intensive process. Unfortunately, this book was not discovered until the time of this writing. For future work in learning human tactical knowledge, this book should be a primary reference. Appendix C gives the verbal reports for all human subjects and cases in the second round of experiments.

Surveys

Another method for collecting data in human-in-the-loop experiments is through written surveys. These surveys can contain such high-level questions as, "Describe any strategies you used during each scenario," to questions asking participants to rank the priority given to some list of items. Essentially, these written surveys after the simulation runs give some time to the human subjects to thoughtfully evaluate their actions. Surveys and think aloud reports help reveal the human subjects' intentions that otherwise would not be clearly understood by observation of the human subjects' actions and scores alone. Appendix B presents the surveys obtained after the experiments.

Finite Number of Cases

A major limitation of human-in-the-loop experimentation is the attention span of the human subjects. After a couple hours of experimentation, it will be difficult for most subjects to continue to perform at their best due to weariness. Also, another problem is subject boredom. Therefore, the experiment designer must be careful to avoid building in too much dimensionality in the independent variables. If a full factorial search of the design space is required, there will be tradeoffs between the number of independent variables, time line of experimentation, and required time for each experiment. There had to exist some limit to the experiment design because the interactions and dimensionality of tactical scenarios are by no means small or simple. Thus, each scenario was constrained to last a total of four or five minutes, which was short enough to both avoid subject boredom as well as make it difficult to complete the entire mission. Each round of experiments contained seven or eight scenarios. Furthermore, interactions between the human subjects and enemies were discretized into four levels of engagement - detection, weapon has line-of-sight, weapon has radar lock, and firing. Finally, independent variables such as the importance of pieces of terrain and the probability of enemy contact were discretized into two and three levels, respectively.

4.2.6 Time Line

On February 1, 2006, the five human subjects went through the first round of experiments. During this first round, each session consisted of two practice scenarios and five scored scenarios. Human-in-the-loop experiments typically begin with a practice session to make sure the human subject is familiarized with the simulation environment. From the human subjects' perspectives, they only had the two practice scenarios in the first round to become to become acquainted with the simulation environment, when in actuality the entire first round was treated as practice (see Section 4.4.1). After all seven scenarios had finished, each human subject filled out a four-question survey. Each scenario lasted five minutes, and the total time for each human subject was about one hour. On February 16, 2006, the five human subjects went through the second round of experiments. In this second round, each session consisted of one practice scenario and seven scored scenarios. After all eight scenarios had finished, each human subject answered a multi-part two-question survey. In this round, each scenario lasted four minutes, and the total time for each human subject was about fifty minutes.

4.2.7 Presentation of Experiments to Human Subjects

At the beginning of the first round of experiments, each subject signed a "consent to participate in non-biomedical research" form [52]. They were then given a two-page description of the experiment, scenario context, their specific mission objectives, joystick and simulation setup, scoring logic, thinking aloud procedures, and visual cues and symbology. Before each scenario began and the four or five minute time limit began counting down, the human subjects were given time to look at the map display and read an intelligence report that defined the probability of running into enemy contacts. At this point, the human subjects were asked if there were any questions. Then, the Camtasia Studio 3 software was activated to begin recording both audio and video. Once that began, the Simulink model and the timer were started. For the first round of experiments, the human subjects had five minutes to complete each scenario, and for the second round, they only had four minutes. Once the time finished, the human subjects took their hands off the joystick, the Simulink model was stopped, and the Camtasia recording ended. At the end of each round, every human subject filled out a survey. Appendix A contains the two-page experiment description and the intelligence reports.

The role of the human subjects was that of an Army helicopter pilot. Their task was to search through both the air corridor and any patches of terrain designated as "critical" to ensure the terrain was safe for the passage of troops. Both the air corridor and critical area were visually designated on the map display (see Figure 4-1). A "critical area" was defined as an area believed to be either an ambush site by the enemy or a possible landing zone for the follow-on troop-carrying helicopters. Therefore, the human subjects were told that it was very important to search carefully through the critical area. They only had four or five minutes to search through as much of the terrain as possible, and they only had five shots to use against enemy

Enemy Types	Tank, SAM, UAV
Terrain	Samples 1 – 6
Ratio of Vehicle Sensor Radius to Air Corridor Width	5:16, 5:22
Stated Probability of Enemy Contact	Slim, Possible, Very Good

Table 4.2: Independent variables.

contacts. The human subjects were told to assume all contacts were hostile, and it was up to their discretion whether to engage or avoid the enemy.

Table 4.2 depicts the independent variables used in the experimentation. In the first round of experiments, there were a maximum of two enemies present in each scenario. In order to increase the difficulty in the second round, there were a maximum of three. Terrains 1–3 were used in the first round, and Terrains 4–6 were used in the second round. The ratio of vehicle sensor radius to the half-width of the air corridor indicates how much the vehicle had to maneuver to cover the entire air corridor. The area of the vehicle’s sensor coverage is that of a circle, $2\pi R^2$, and with $R = 5$ equal to 50π (units²). Note that the units are generic and are defined relative to visual perception of the vehicle on the display (see Section 4.4.3). If the half-width of the air corridor is 8 units, then a square with sides equal to the width of the corridor would have an area equal to 256 units². Thus, the vehicle’s sensor only covers $\frac{50\pi}{256} \approx 61.4\%$ of the area of a square of air corridor with sides equal to its width. For a 5:22 ratio, the vehicle only covers 32.5% of a square of air corridor with sides equal to its width. Finally, if the scenario included both an air corridor and critical area (which was true in almost all cases), each of these regions had a separate stated probability of enemy contact, where slim < possible < very good, notionally speaking. These ranges of probability were chosen because they are more intuitively acceptable to a human vice an exact percentage, such as trying to understand the difference between a 37% versus 42% of enemy contacts [17]. The three probability categories can be viewed as covering equal thirds of 100%.

4.3 Autonomous Comparison

The development of a baseline AV behavior had two purposes. First, the baseline behavior was designed without any inspiration from the human-in-the-loop experiments. The baseline behavior was simple but logical. Therefore, the baseline behavior should exhibit reasonable performance, but one of its purposes was to highlight where the human subjects were stronger. Second, after deriving human tactics by filtering good from bad decisions, the experiment would come full cycle, and the baseline behavior augmented with the new tactics. Thus, human tactics could be applied to the baseline behavior to validate the research’s goal of improving AV performance through human expertise.

4.3.1 Vehicle Governing Behavior and Equations

The vehicle’s governing behavior, apart from interactions with enemy contacts, is defined by two main categories of search and obstacle avoidance. After presenting these categories, we show the vehicle’s dynamics of 2D motion.

Search

Stored in the vehicle’s navigation log are a terrain database and list of waypoints. These waypoints mark the centerline of the air corridor. These waypoints also tend to mark heading changes in the corridor’s path. The vehicle’s search behavior is to fly diagonally across the air corridor, zig-zagging from waypoint to waypoint. Once the next waypoint falls within the vehicle’s sensor radius, the vehicle alters its course to fly directly to the waypoint. Once the vehicle reaches the waypoint, the waypoint index increases by one, and the vehicle begins searching along the corridor to the next waypoint. If the vehicle misses the waypoint, which means the vehicle passed by without the waypoint ever falling within its sensor radius, the vehicle takes a direct path toward the waypoint’s stored location until it acquires it. At this point, the waypoint index increases by one, and the search process continues. Thus, the search process can be divided into a series of steps between $waypoint_i$ and $waypoint_{i+1}$.

Once the vehicle begins a new search segment from $waypoint_i$ to $waypoint_{i+1}$, two things occur. First, the vehicle calculates the number of diagonal segments required to reach the next waypoint, as a function of its search angle off the air corridor’s centerline and the total distance to cover to the next waypoint. This number of diagonal segments, called *segmentRatio*, is given by the following equation:

$$segmentRatio = \frac{\|\overrightarrow{waypoint_i} - \overrightarrow{waypoint_{i+1}}\|}{halfWidth} \tan(\theta) \quad (4.1)$$

where $\overrightarrow{waypoint_i}$ and $\overrightarrow{waypoint_{i+1}}$ are the current and next waypoint (x, y) ordered pairs, respectively, along the AV’s path; *halfWidth* is half of the width of the air corridor; and θ is the search angle measured from the centerline. Equation 4.1 assumes the vehicle is beginning from the current waypoint. If this is not the case, the vehicle’s distance along the corridor centerline away from the current waypoint can either be added or subtracted as necessary. As will be discussed later, *segmentRatio* serves as a logical comparison to know if the AV has missed its next waypoint or not. Second, to begin searching the vehicle turns by default to its left, as facing the next waypoint along the centerline, by the angle θ .

While the AV is diagonally searching along the corridor, it calculates the perpendicular distance from its position to the air corridor centerline. Once this perpendicular distance almost equals the half-width of the corridor minus two-thirds of the vehicle’s sensor radius, the AV has essentially reached the corridor boundary and makes a 90° turn to cross back through the air corridor. The reason that two-thirds of the vehicle’s sensor radius is subtracted out of the half-width is due to efficiency. If the vehicle continued its diagonal search all the way to the edge of the air corridor, its

sensor radius would fall on terrain outside the air corridor, which, as to be discussed, would contribute nothing to its performance score. Thus, this is just a simple tradeoff that was qualitatively determined.

In order to search through a critical area, a function calculates the number and location of waypoints needed inside the critical area to guarantee an adequate amount of coverage. This function is dependent on the dimensions of the critical area and the sensor radius of the vehicle. The first critical area waypoint, by default, is located in the upper left corner of the critical area box. From there, the vehicle diagonally searches clockwise around the perimeter of the box, and then the vehicle proceeds vertically up and down through any inside columns of waypoints.

Obstacle Avoidance

At every time step, the AV projects its current path forward a few steps into the future and compares the final point of this path to the terrain database. If the final point of this path matches an obstacle cell (i.e., if the vehicle will hit an obstacle when continuing on the same path), the vehicle enters into a **while** loop and iteratively modifies its current heading by $\pm 10^\circ$. At every iteration, the AV projects the temporary heading forward and tests seven points along the path against the terrain database. Whichever direction of $\pm 10^\circ$ finds a new path where all seven points are clear of obstacles, the AV replaces its current, desired heading by this new, temporary one. A default parameter called *terrainAvoidanceCounterVeh* defines how long the AV must hold this modified heading. If the AV was searching along the corridor at the time of terrain detection, the AV transitions to flying directly to the next waypoint, once it has cleared the terrain obstacle. If the AV was engaging an enemy at the time of obstacle detection, the obstacle avoidance takes precedence. While it is avoiding the obstacle, the AV is also stepping through its engagement logic, but any desired maneuvering outputs specified by the engagement logic are not passed through until the AV has cleared the obstacle. Therefore, the obstacle avoidance logic overrides all other outputs determined by the vehicle's search and engagement logic.

Early tests of this logic found that the vehicle sometimes chose paths that barely cleared the obstacle. If the vehicle maintained a path close to the obstacle and the counter designated by *terrainAvoidanceCounterVeh* finished before it had cleared the terrain obstacle, the vehicle would try to find a new heading. However, that new heading would still be hindered by the obstacle, and the vehicle would iterate for another temporary obstacle avoidance heading. During that process of trying to change its heading and iterating again due to the obstacle, the vehicle would take a one time step movement towards the desired heading into the obstacle. If this occurred more than three times, the vehicle tended to run into the obstacle. Thus, a 30° buffer angle was added to the temporary obstacle avoidance heading to keep the vehicle safely away from the obstacle.

Vehicle Equations

The vehicle two-dimensional dynamics are defined only from its velocity V and heading angle θ .

$$\dot{x} = V \cos(\theta) \quad (4.2a)$$

$$\dot{y} = V \sin(\theta) \quad (4.2b)$$

where \dot{x} and \dot{y} are the first derivatives of the x and y positions of the vehicle with respect to time. For the human-in-the-loop experiments, V and θ are inputted by the humans' movements of the joystick (see Section 4.2.3). For the AV, V and θ are outputted by the statechart, which is essentially the controller in the plant's feedback loop and will be discussed in detail in Section 4.3.4. The vehicle's speed in rectangular coordinates, \dot{x} and \dot{y} , are calculated and then numerically integrated to update the vehicle's x and y position. To approximate heading rate so that it takes time for the vehicle to turn, $\dot{\theta}$ is approximated by a finite difference equation, given by:

$$\dot{\theta} \approx \frac{\theta_{i+1} - \theta_i}{\Delta t} \quad (4.3)$$

where θ_i is the current heading. By setting a default heading rate, Equation (4.3) can be solved for θ_{i+1} . Now, when the statechart outputs a new commanded heading, the vehicle turns by the amount $\dot{\theta}\Delta t$ every time step until it reaches the commanded heading.

4.3.2 Enemy Interactions

The two-dimensional dynamics of the enemy platforms are also defined by Equations (4.2) and (4.3), and the enemy platforms also follow the same obstacle avoidance logic as the vehicle. The difference is the modes that govern how the enemy platforms move within the simulation as compared to the vehicle's zig-zagging, waypoint-following, search mode. After presenting these enemy behaviors, we then discuss the process and logic of interactions between the AV and the enemy platforms.

Enemy Behaviors

The three enemies called Tank, SAM, and UAV correspond to general enemy types of ground moving, ground static, and airborne, respectively. Note, of course, that operating in a two-dimensional environment negates the technical terming of air and ground vehicles. There are three distinct modes that define the enemy's movement and/or tracking behavior during the simulation. Note that these modes only apply to the Tank and UAV. The SAM is completely static and does not need to track the vehicle with its radar because its weapons cone extends around a full 360° azimuth. The first mode exists from the start of the simulation until vehicle detection. In this mode, the moving enemies can either be static or following a predefined "patrol route." The second mode exists during detection and engagement. In this mode,

the moving enemies are reactively following the vehicle. Note that the enemy does not project the vehicle's motion into the future to try and find intercept paths. At every time step, the enemy determines the range and heading to the vehicle and simply tries to follow the vehicle. The third mode is post-engagement. If the vehicle successfully out-maneuvers the enemy and breaks away from an engagement, the moving enemies randomly choose one of five locations on the map and move to that location. The enemy then enters into a "holding pattern" where it simply waits until another detection occurs. These five locations roughly corresponding to the four corners and middle of the map.

RADAR

The AV has two simple radars. They both have the same range, which is a fixed radius termed *vehicle.range*, but different field-of-views, and they serve different purposes. The first is the vehicle's primary radar which covers the entire 2π azimuth around the vehicle. This sensor is the primary search-and-detect sensor. At every time step during the mission, the vehicle is said to have "seen" that terrain which lies inside the primary sensor's radius range. Thus, at every time step, the vehicle covers a circular footprint of the terrain. Also, the vehicle has no knowledge of enemy presence until the enemy's position falls within this radar circle. The second radar is the vehicle's weapons radar which only extends 45° from the nose to both sides of the aircraft. In a three-dimensional world, the appropriate concept is of a radar cone extending from the aircraft. Once the enemy's position is inside this weapons cone, the weapon begins acquiring the target and is said to have "line-of-sight" (LOS) to the target. The test to determine if the enemy is within the vehicle's weapons cone is two-fold. First, the vector defined from the enemy's position to the vehicle's position, $\overrightarrow{enemy2vehicle}$ is projected onto the vehicle's current heading vector, $\overrightarrow{headingVeh}$. Using the dot product, the included angle, α_{IN} is found by:

$$\alpha_{IN} = \cos^{-1}\left(\frac{\overrightarrow{enemy2vehicle} \cdot \overrightarrow{headingVeh}}{\|\overrightarrow{enemy2vehicle}\| \|\overrightarrow{headingVeh}\|}\right) \quad (4.4)$$

If $\alpha_{IN} < \text{vehicle.coneANG}$, where *vehicle.coneANG* is the 45° half-angle of the cone, then the enemy lies within the vehicle's field-of-view. The second test is to ensure that the magnitude of the distance between the vehicle and the enemy is less than the sensor radius range, that is $\|\overrightarrow{vehicle.position} - \overrightarrow{enemy.position}\| \leq \text{vehicle.range}$. If the vehicle maneuvers to keep the enemy inside this weapons cone for a fixed time parameter called *time2lock*, the weapon is said to have "radar lock" on the enemy. At this point, the vehicle can now fire upon the enemy. The probability of hitting the enemy is a function of the distance between the vehicle and enemy at time of firing. Thus, if the vehicle achieves radar lock on the enemy and chooses to fire, it inputs $\|\overrightarrow{vehicle.position} - \overrightarrow{enemy.position}\|$ to a look-up table, linearly interpolates, and outputs a probability of hitting the enemy. This probability is compared to a uniformly generated random number. If the random number is less than the probability, the vehicle is said to have hit the enemy. There is one final

time count once radar lock is achieved and the probability of hitting the enemy is true. The motivation of this counter is to account for the time required between shots. A SAM site shooting missiles takes longer to fire due to the wait required to start the missile's engine than a tank that is firing shells. Once this counter runs out, the enemy loses one health point, and the shot is displayed on the screen. This sequence from detection to LOS to radar lock to firing upon the enemy is summarized by Algorithm 1.

```

Given:  $\overrightarrow{vehicle.range}, \overrightarrow{vehicle.coneANG}, \overrightarrow{time2lock}, \overrightarrow{tableVeh}$ 
Input:  $\overrightarrow{enemy2veh} = \|\overrightarrow{vehicle.position} - \overrightarrow{enemy.position}\|, \alpha_{IN}$ 
if  $\|\overrightarrow{enemy2veh}\| \leq \overrightarrow{vehicle.range}$  then
    Detection;
    if  $\alpha_{IN} \leq \overrightarrow{vehicle.coneANG}$  then
        Line-of-Sight;
         $tempCounter = 0;$ 
        while  $tempCounter < \overrightarrow{time2lock}$  do
             $tempCounter++;$ 
        end
        Radar Lock;
        if  $flag2FIRE = 1$  and  $\overrightarrow{tableVeh}(\|\overrightarrow{enemy2veh}\|) \geq \overrightarrow{uniformRandNum}$ 
then
            Enemy to be Fired Upon;
             $tempCounter2 = 0;$ 
            while  $tempCounter2 < \overrightarrow{time2shoot}$  do
                 $tempCounter2++;$ 
            end
            Enemy Hit;
        end
    end
end

```

Algorithm 1: Radar logic from detection of enemy to firing upon the enemy.

4.3.3 Parameters of Simulated Entities

Table 4.3 displays the parameters of each simulation platform. Two important considerations must be noted with this table. First, the vehicle's units of position are generic. This was noted before and the reason why will be detailed in Section 4.4.3. Second, the maximum velocity, heading rate, and time to shoot parameters are given as variables and not scalars because they changed after the first round of experimentation. The reason for this change is tied directly to the first note of generic units and discussed as well in Section 4.4.3. As can be seen in this table, the vehicle and UAV are almost identical in their defining parameters, except for the heading rate.

	VEHICLE	TANK	SAM	UAV
sensor radius (units)	5	3	8	5
weapons cone angle (deg)	45°	45°	180°	45°
max velocity (units/sec)	V_{max}	$(\frac{1}{2})V_{max}$	—	V_{max}
heading rate (rad/sec)	$\dot{\theta}$	$(\frac{1}{4})\dot{\theta}$	—	$(\frac{1}{2})\dot{\theta}$
<div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 5px auto;"> probability table $(\ \overrightarrow{enemy2veh}\ P(hit))$ </div>	0 1.0	0 0.75	0 1.0	0 1.0
	3.00 1.0	1/3 0.75	3.00 1.0	3.00 1.0
	3.01 0.75	1/3 0.5	3.01 1.0	3.01 0.75
	4.00 0.75	1/2 0.5	4.00 1.0	4.00 0.75
	4.01 0.5	1/2 0.3	4.01 0.9	4.01 0.5
	5 0.5	3 0.3	8 0.9	5 0.5
	> 5 0.0	> 3 0.0	> 8 0.0	> 5 0.0
	time to shoot (time steps)	3(plotCounter)	2(plotCounter)	4(plotCounter)

Table 4.3: Defining parameters of simulation platforms.

It takes the UAV twice as long to turn to a heading as the vehicle. The tank has a small sensor radius, maximum velocity, heading rate, and probability of hit than the vehicle. The tank, however, does have a lower required time to shoot counter and can thus fire quicker than the vehicle. The SAM has a larger sensor radius, weapons cone angle, and greater probability of hit. Yet, the SAM is static and takes the longest to shoot.

4.3.4 Statecharts

Hickie used statecharts in an initial attempt to capture tactical knowledge [31] from interviewing human experts. However, Hickie had to translate the statecharts into software code to test them against enemies in the U.S. Army’s force-on-force simulation tool, One Semi-Automated Forces (OneSAF) Testbed Baseline 2.0 (OTB 2.0). This was a time-consuming process. MATLAB offers a stateflow toolbox that allows the programmer to design statecharts through a graphical user interface (GUI) [43]. This toolbox was used to design the baseline AV behavior as well as to capture human-inspired tactics.

Background

Flight environments are highly transitory, reactive, dynamic, and very difficult to model. Harel proposed statechart diagrams as a way to characterize large, complex reactive systems [30]. The innovation behind statechart diagrams is the ability to include depth, orthogonality, and broadcast communication. Traditional state diagrams or flow charts are insufficient for representing large, complex systems because of the “unmanageable, exponentially growing multitude of states, all of which have to be arranged in a ‘flat’ unstratified fashion” [30]. For example, the well known human

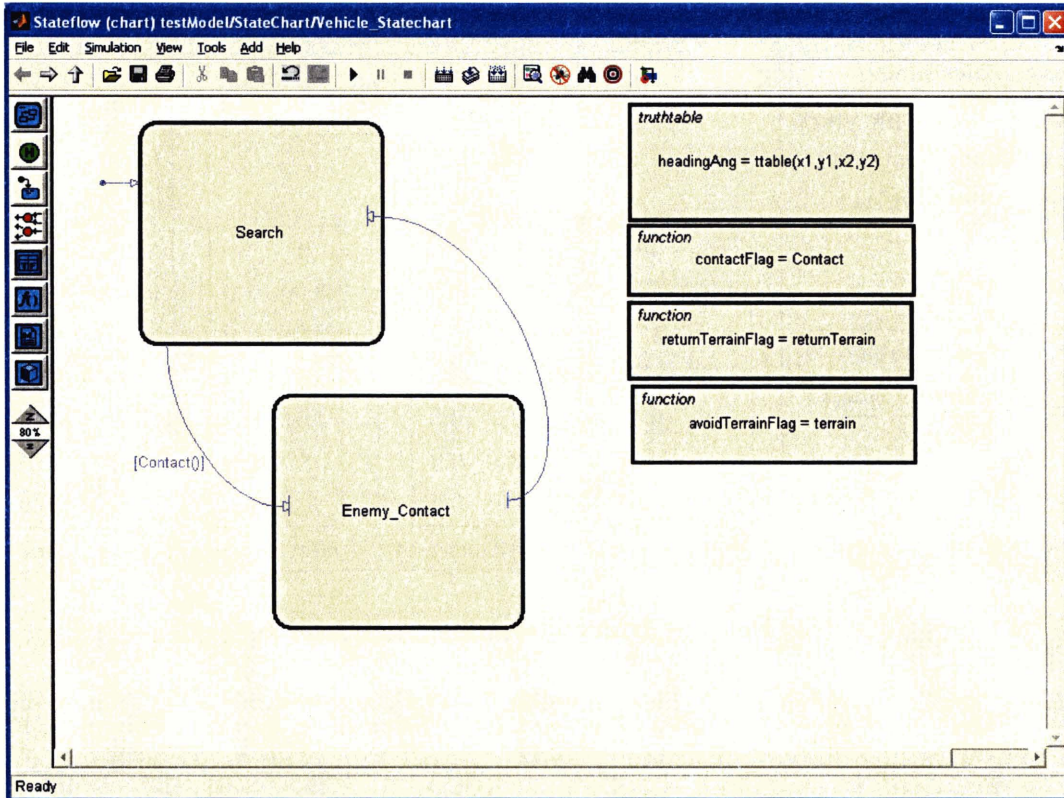


Figure 4-4: The AV's top level nominal and reactive superstates.

behavior model (HBM) in military simulation exercises, TacAir-SOAR, is a rules-based system where tasks are hierarchically decomposed from high-level actions, such as moving to a target, to low-level actions of changing heading to fire a missile. The designers acknowledge that the novelty behind TacAir-SOAR “is primarily a matter of scale and integration.” In 1991, there already existed over 5,200 rules encoded in TacAir-SOAR [35]. As the design space for AVs continues to expand in response to fidelity and domain size requirements, statecharts offer the advantage of moving beyond conventional two-dimensional representations. They possess the ability to add multi-dimensional depth by lumping states together into superstates, to execute parallel transitions between separate events in orthogonal states, and to broadcast a transition across multiple states.

Baseline Statechart

The baseline statechart begins at the very top level with two exclusive superstates, Search and Enemy_Contact, as shown in Figure 4-4. The AV is either operating in the Search superstate or the Enemy_Contact superstate, but not both. The statechart defaults to the Search superstate, as depicted by the arrow in the top-left corner of Figure 4-4. This default arrow can be visually recognized by the dot at the tail. This default arrow designates the first state or junction the vehicle enters when activating

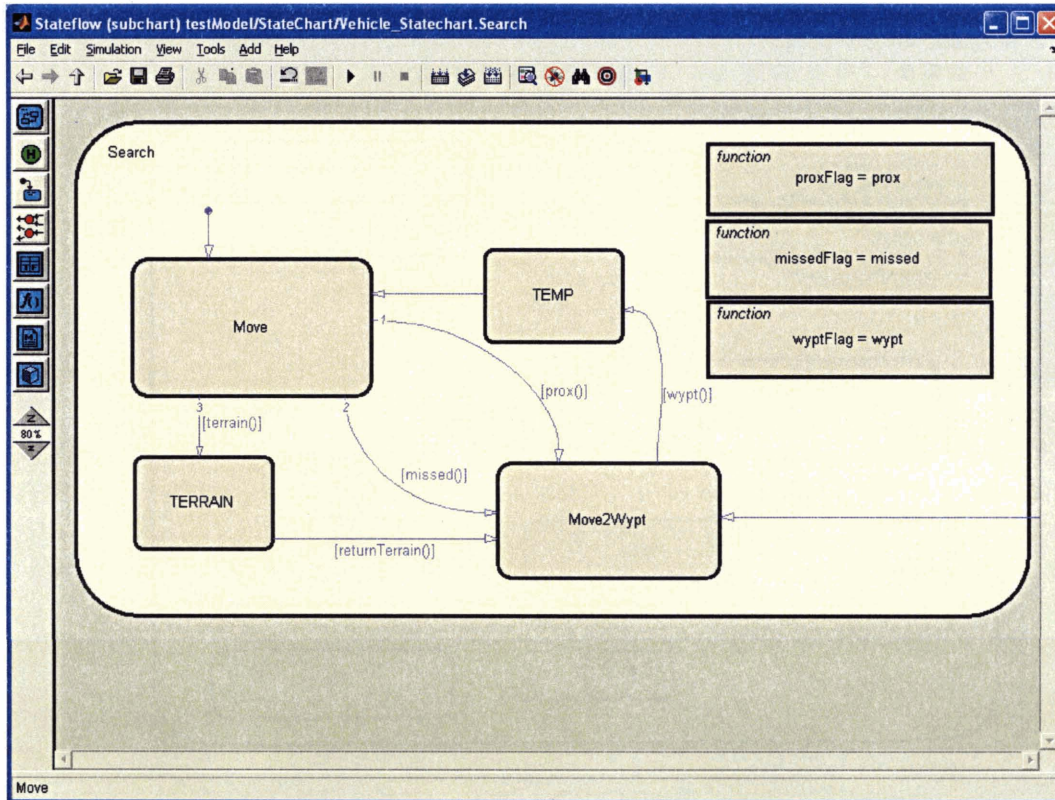


Figure 4-5: Search superstate - the search behavior of the AV.

that level of the statechart. The Search superstate is the nominal behavior for the AV. The AV's mission is to search through the corridor and react to enemy presence if required. The Enemy_Contact superstate is the reactive component of the AV's behavior that responds appropriately to pop-up threats. Thus, the AV is either following along the nominal search path or reacting to pop-up threats.

Search superstate There are four states in the Search superstate: Move, TERRAIN, Move2Wypt, and TEMP, as shown in Figure 4-5. As described earlier (see Section 4.3.1), the vehicle moves diagonally through the air corridor from waypoint to waypoint. Once the next waypoint lies within the vehicle's sensor radius, it triggers a transition out of the Move state to the Move2Wypt state. If it misses the waypoint, it transitions to the Move2Wypt state, and thus, it turns around to head back towards the missed waypoint. This way the vehicle is forced to always begin the next search segment at a known location or reference point. Once the vehicle reaches the waypoint, it transitions back to the Move state and begins searching until the next waypoint is reached. If at any time a terrain obstacle is present, the vehicle transitions to and remains in the TERRAIN state until it is cleared. From the TERRAIN state, it always transitions to Move2Wypt. The TEMP state is necessary to allow one time step to pass as the vehicle increments the waypoint index to begin searching again.

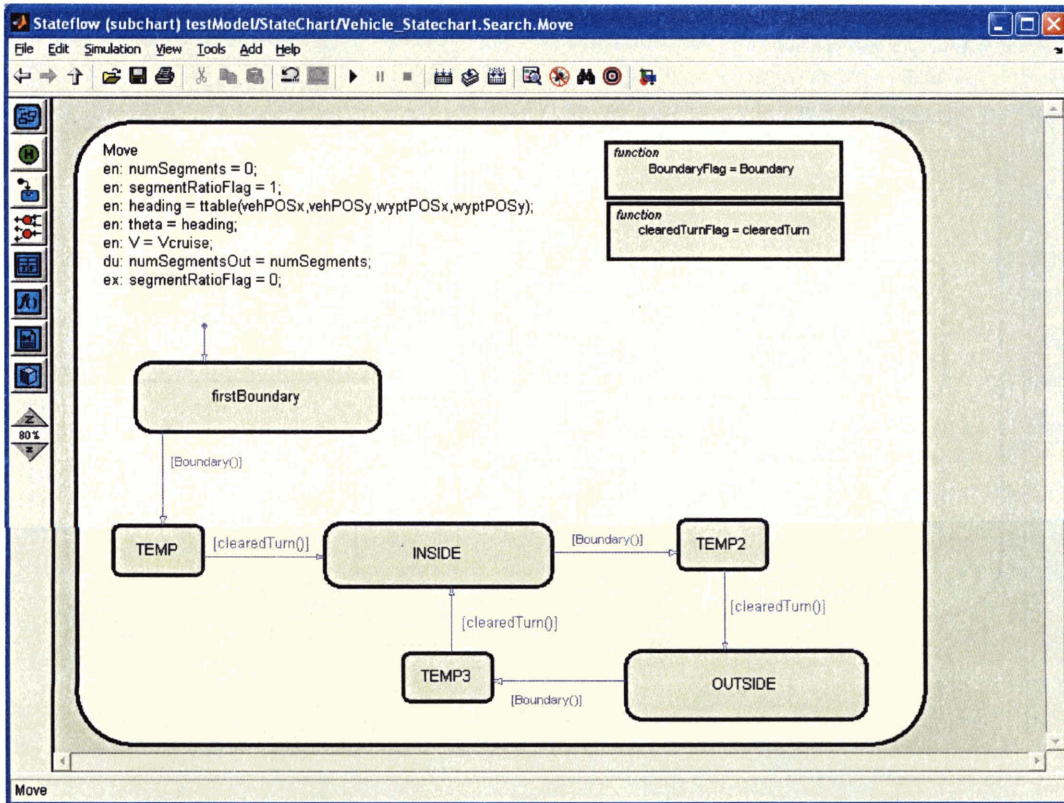


Figure 4-6: Move state - diagonal searching between waypoints.

Search.Move The Search superstate defaults to the Move state. Inside the Move state are three substates: `firstBoundary`, `INSIDE`, and `OUTSIDE`, as displayed by Figure 4-6. Upon entering the Move state, two important variables, in regards to determining if the vehicle has missed the next waypoint, are initialized. The first variable, `numSegments` is set to zero, and the total number of half-width corridor segments, `segmentRatio` is calculated (see Section 4.3.1). Every time the vehicle crosses the air corridor centerline, `numSegments` increases by one. However, because the vehicle always begins the next search portion at a waypoint which lies in the air corridor centerline, the first diagonal segment only covers half of the width of the air corridor whereas all other diagonal segments cover the entire width. Thus, the vehicle always defaults to turning to its left. Upon reaching the first boundary, the vehicle turns 90° back across the corridor, and `numSegments` increases by one. Upon reaching the next boundary, the vehicle has now crossed the entire length of the corridor. When the vehicle turns for the next diagonal segment, `numSegments` increases by two. If $numSegments > \text{floor}(segmentRatio)$, then the vehicle has missed the waypoint. The temporary states exist because it takes time for the vehicle to turn upon reaching the air corridor boundaries. The test for the `boundary` function is $d \geq (\text{halfWidth} - \frac{2}{3}\text{vehicle.range})$, where d is the perpendicular distance from the vehicle to the air corridor centerline (see Section 4.3.1). Because it takes time for the vehicle to turn, the above condition for the `boundary` function does not evaluate to be

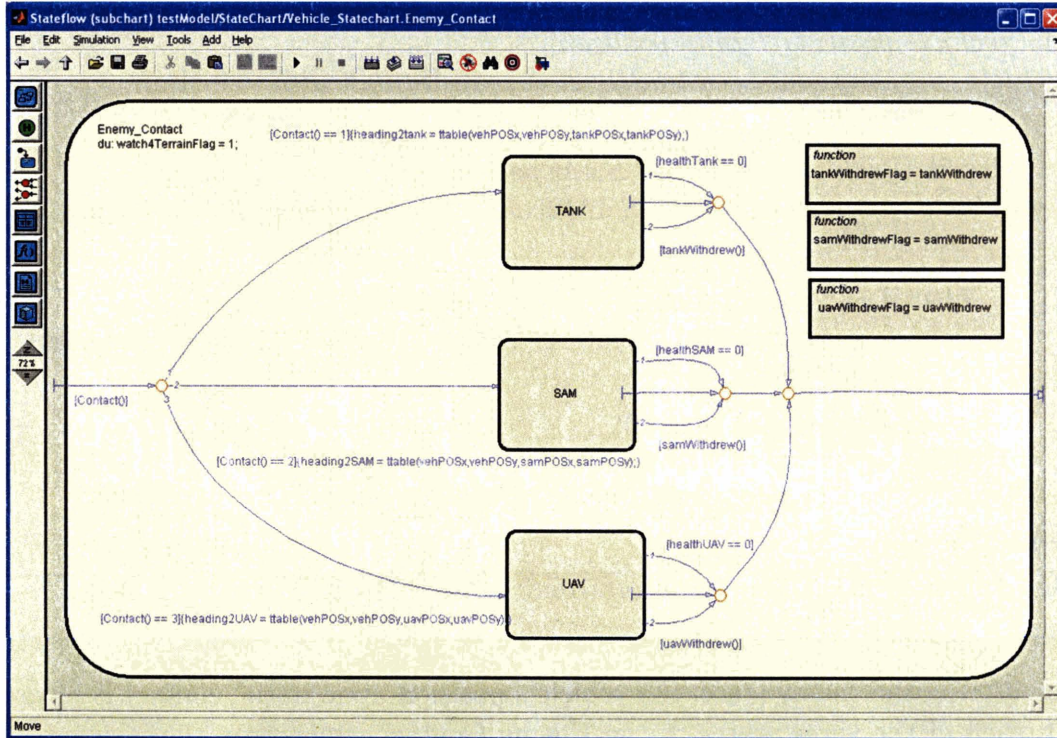


Figure 4-7: Enemy_Contact superstate - detection of enemy contact.

false at the next time step. Therefore, a second function called *clearedTurn* activates the transition to the next side of the corridor once the above condition evaluates to be false.

Enemy_Contact superstate The transition from the Search to the Enemy_Contact superstate ends inside the Enemy_Contact superstate on a connective junction, as depicted by Figure 4-7. This connective junction is the red circle at the left of the figure. The vehicle moves from the connective junction along one of its branching paths depending upon which of the branch's transition logic is true. It's a graphical form of IF-THEN logic. Proceeding from this connective junction are three paths that transition to one of three new states - TANK, SAM, and UAV. Therefore, the transition out of the Search superstate assumes away any identification phase and proceeds directly to the enemy platform in contact.

The transition from the Search to the Enemy_Contact superstate is given by the function, *contact*. Its logic is given by Algorithm 2. Once the *contact* function evaluates to true, the Search superstate is exited and the transition ends on the connective junction. The transition out of the connective junction can follow one of three paths as described above, depending on the integer value of *contactFlag*. During the transition from the connective junction to enemy state (Tank, SAM, UAV), a truth table called *ttable* is consulted wherein the heading from the vehicle to the enemy is calculated. This truth table is located at the statechart's top level, Figure 4-4. Also,


```

if ( $\|\overrightarrow{vehicle.position} - \overrightarrow{tank.position}\| \leq$ 
 $vehicle.rangeorradarLock.tank2veh = 1$ )and $tank.health \neq 0$  then
    contactFlag = 1;

    else if ( $\|\overrightarrow{vehicle.position} - \overrightarrow{SAM.position}\| \leq$ 
 $vehicle.rangeorradarLock.SAM2veh = 1$ )and $SAM.health \neq 0$  then
        contactFlag = 2;

        else if ( $\|\overrightarrow{vehicle.position} - \overrightarrow{UAV.position}\| \leq$ 
 $vehicle.rangeorradarLock.UAV2veh = 1$ )and $UAV.health \neq 0$  then
            contactFlag = 3;

            else contactFlag = 0;
        end
    end
end

```

Algorithm 2: Logic for transition from SEARCH to ENEMY.CONTACT.

upon entering the enemy state, the vehicle's position upon initial contact is recorded.

In each of the three states in the Enemy_Contact superstate, the new substates and transitions that define the engage and avoid criteria are the same, except for one. If the vehicle has contacted a SAM, and as long as the SAM parameters have not been changed, it will always default along a transition path to an AVOID state. It will never attempt to engage.

Enemy_Contact.ENEMY In each of the enemy states (TANK, SAM, UAV), there are three substates called ENGAGE_TA, AVOID, and ENGAGE_TD, as shown in Figure 4-8. For the engagement states, TA and TD stand for "tactical advantage" and "tactical disadvantage," respectively. In this statechart, tactical advantage or disadvantage is a single-parameter function, based on the platform's sensor radius range. Note that there are other parameters which define a platform's behavior and do contribute to its tactical advantage or disadvantage (see Table 4.3). For example, the larger heading rate of the vehicle over the UAV is a tactical advantage. However, the sensor radius range parameter defines tactical advantage or disadvantage in this statechart because the vehicle will never be able to shoot the enemy if the enemy is not first within its sensor radius range. Thus, if the vehicle has a smaller sensor radius range than the enemy, the enemy has the advantage of using its superior sensor radius range as a stand-off capability to freely fire upon the vehicle. Algorithm 3, then, defines which of the three substates - ENGAGE_TA, AVOID, ENGAGE_TD - the vehicle enters upon contacting an enemy.

A transition exists from each of the engagement states to the AVOID state if the vehicle, while engaging, runs out of ammo. In that case, the condition $numAmmo = 0$ evaluates to true, and the vehicle moves to AVOID. Conversely, there are transitions out of the AVOID state to a connective junction, which then ends at one of the engagement states. The first transition is defined as follows: if at any time, when

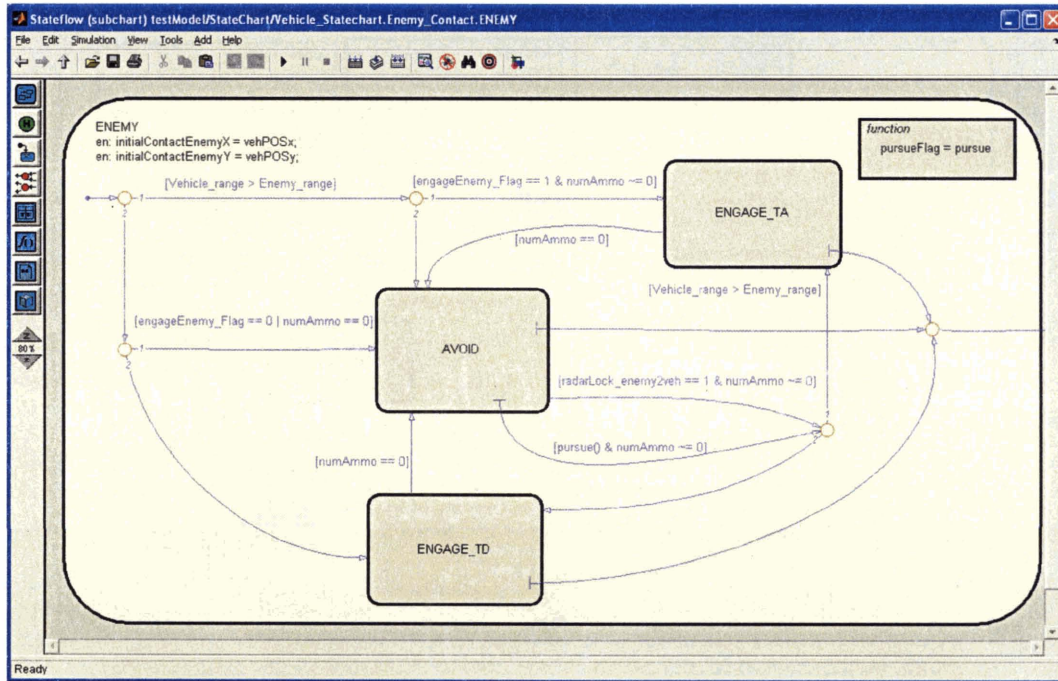


Figure 4-8: Enemy_Contact.ENEMY - three engagement modes.

```

if vehicle.range > enemy.range then
  if engageEnemy.flag = 1 and numAmmo ≠ 0 then
    move to ENGAGE_TA;
  else move to AVOID;
  end
else
  if engageEnemy.flag = 0 or numAmmo ≠ 0 then
    move to AVOID;
  else move to ENGAGE_TD;
  end
end

```

Algorithm 3: Logic for engaging or avoiding enemy contact.

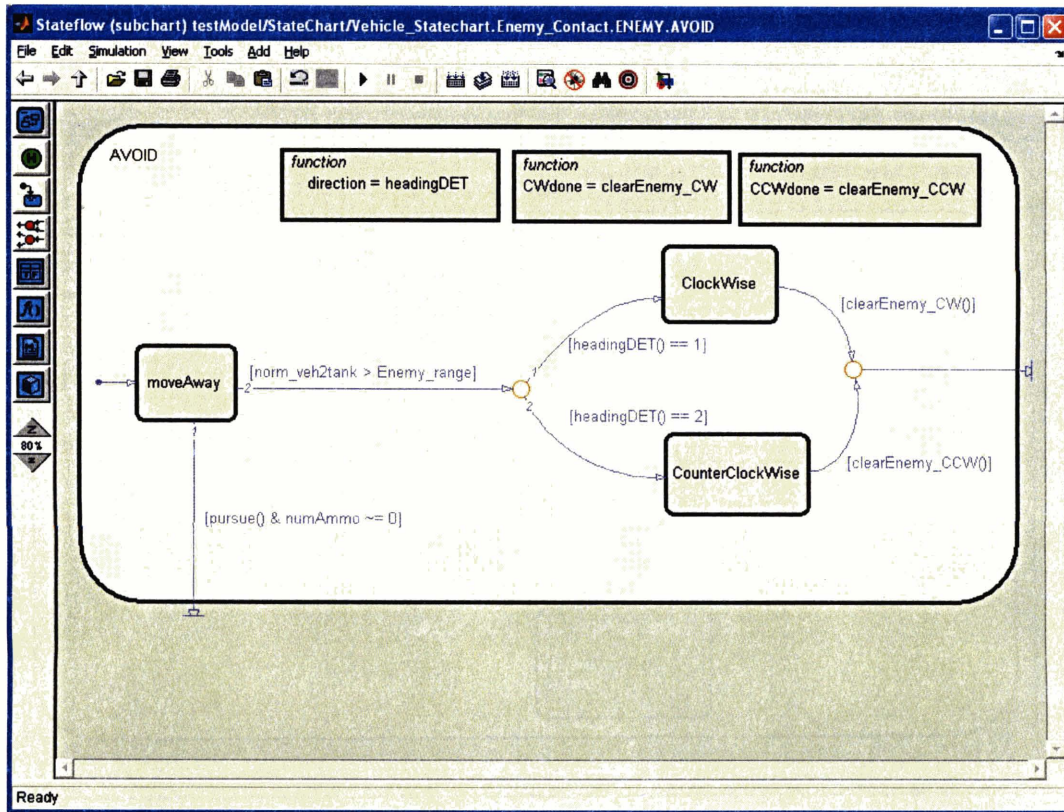


Figure 4-9: Enemy_Contact.ENEMY.AVOID - maneuvering to avoid enemy.

the vehicle is executing the contents of the AVOID state, the enemy gets radar lock on the vehicle, the vehicle transitions out of AVOID to ENGAGE_TA/TD as long as $numAmmo \neq 0$ is true. The second transition is a super-transition, a transition out of both the substate and the parent state. In this case, the default substate in the AVOID parent state is `moveAway`, where the vehicle seeks to move out of the enemy's sensor radius range, as displayed by Figure 4-9. Remember, that the position of the vehicle upon initial contact with the enemy was recorded. A function called *pursue* evaluates the condition $\|vehicle.position - vehicle.positionInitialEnemyContact\| > 4 * vehicle.range$. This *pursue* function is used as a transition for a few different states, but here it serves as a transition out of the `moveAway` substate, out of the AVOID parent state, and to a connective junction that splits to either one of the ENGAGE_TA/TD states. The reasoning is as follows: if the vehicle is trying to avoid the enemy, but while executing the contents of the substate `moveAway`, the vehicle is never able to move completely out of the enemy's sensor radius range, it can only be because the enemy is pursuing the vehicle. Therefore, if the vehicle has flown a distance greater than four times its sensor radius range and has not already transitioned out of the `moveAway` substate, the vehicle is being pursued and turns to engage, as long as there is ammunition. Of course, this assumes that the enemy's sensor radius range is significantly less than four times the vehicle's range.

As a summary up to this point, we have discussed the logical conditions and parameters which take the vehicle from its default Search superstate into the Enemy_Contact superstate and ending in one of the TANK, SAM, UAV states, the ENEMY states. Next, we discussed the logical conditions that define the path from entering one of the ENEMY states to one of its three substates, ENGAGE_TA, AVOID, or ENGAGE_TD. Finally, we have discussed the inter-state transitions between the engage and avoid states. Therefore, we are left with describing the contents of each of the engage and avoid states, any other transitions out of these states, and finally the transitions out of the ENEMY states and ultimately out of the Enemy_Contact superstate.

Enemy_Contact.ENEMY.AVOID As already mentioned, the default substate in the AVOID state is called *moveAway*, as displayed by Figure 4-9. Upon entering this substate, the vehicle turns 180° from its heading to the enemy to exit most quickly out of the enemy’s sensor radius range. There are two transitions out of the *moveAway* substate. The first one, mentioned above, used the *pursue* function to transition from avoiding to engaging. The second transition, evaluates the condition $\|\overrightarrow{vehicle.position} - \overrightarrow{enemy.position}\| > enemy.range$. It is assumed, then, that the AV has been given an estimate of the parameter *enemy.range*. If the condition evaluates to true, then the vehicle has moved outside of the enemy’s sensor radius range and transitions to a connective junction and calls the function *headingDET*. In this function, the vehicle evaluates the cross product of the its orientation vector with $\overrightarrow{waypoint_{i+1}}$. By the right-hand rule, the sign of the third element in the cross product determines whether the waypoint lies to the left or right of the vehicle’s orientation vector. After evaluating the function *headingDET*, the vehicle then transitions from the connective junction to either the ClockWise or CounterClockWise substate. In the ClockWise substate, the vehicle determines, at each time step, its heading to the enemy and then adds 90° to calculate its own heading. On the other hand, in the CounterClockWise substate, the vehicle subtracts 90° from the calculated heading to the enemy. In this way, the vehicle, which has moved out of the enemy’s sensor radius range, moves in increments around the enemy’s sensor radius range towards its next desired waypoint. The transitions out of the ClockWise and CounterClockWise substates, determined by the functions *clearEnemy_CW* and *clearEnemy_CCW*, end on a common connective junction and transition out of the Avoid state. These functions evaluate as true when the third element of the cross product of the vehicle’s orientation vector with the next waypoint flips sign. Thus, the waypoint which was originally on the vehicle’s right, for example, and which stayed on the vehicle’s right as the vehicle made small heading changes to move counter clockwise around the enemy’s sensor radius range, is now on the vehicle’s left. The track from the vehicle to the waypoint is no longer obscured by the enemy’s presence, and the vehicle can transition from avoiding to moving to the waypoint and continue its searching.

Enemy_Contact.ENEMY.ENGAGE_TA The strategy in engaging with a tactical advantage is using the larger sensor radius range as stand-off capability. The

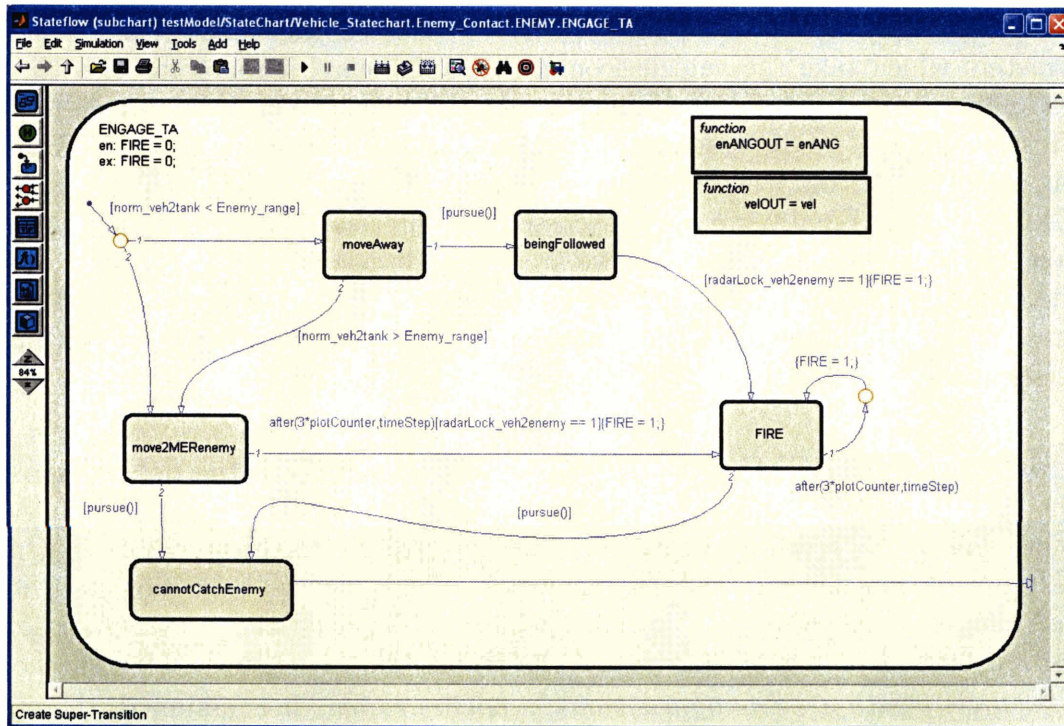


Figure 4-10: Enemy_Contact.ENEMY.ENGAGE_TA - maneuvering to engage the enemy with a tactical advantage.

vehicle moves just outside of the enemy’s sensor radius range, denoted here as MER which stands for "maximum effective range," waits for weapons lock, and fires upon the enemy.

In the `ENGAGE_TA` state, as displayed by Figure 4-10, there are five substates: `moveAway`, `move2MERenemy`, `beingFollowed`, `FIRE`, and `cannotCatchEnemy`. In entering `ENGAGE_TA`, the default path is to a connective junction with the following logic: if $\| \overrightarrow{vehicle.position} - \overrightarrow{enemy.position} \| < enemy.range$, then transition to `moveAway` state; else transition to `move2MERenemy` state. There are two transitions out of the `moveAway` state. First, if $\| \overrightarrow{vehicle.position} - \overrightarrow{enemy.position} \| > enemy.range$, then transition to `move2MERenemy` state. Second, if `pursue` (same function as in `AVOID`) evaluates to true, it means the enemy is pursuing the vehicle, and the vehicle transitions to the `beingFollowed` state. In the `beingFollowed` state, the vehicle stops, and turns to face the enemy, waits for radar lock, and transitions to the `FIRE` state. In the `move2MERenemy` state, the vehicle maintains its stand-off distance while waiting for weapons lock to transition to the `FIRE` state. Upon entering the `FIRE` state, the `FIRE` flag triggers true, and the vehicle fires upon the enemy. The `FIRE` state has two transitions. The first is a self-transition. It waits for a predetermined number of time steps, then it self transitions back to the `FIRE` state, firing on the enemy again. The second transition uses the function `pursue`. While the vehicle is in the `move2MERenemy` or `FIRE` state and `pursue` evaluates

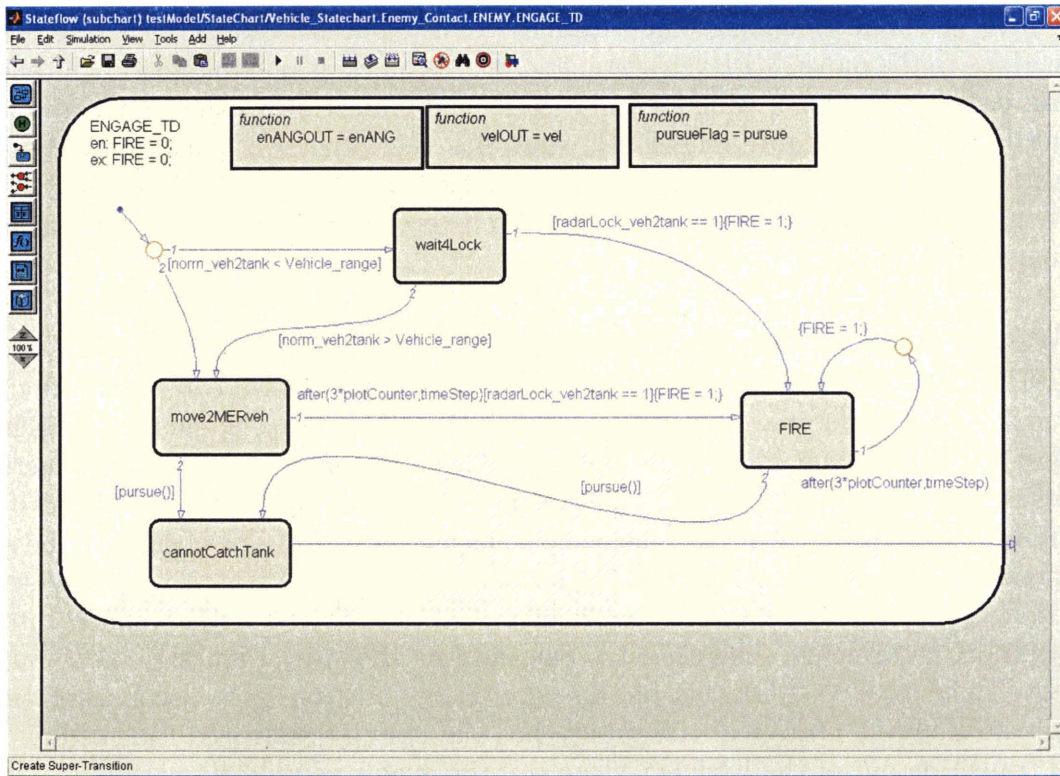


Figure 4-11: Enemy_Contact.ENEMY.ENGAGE_TD - maneuvering to engage the enemy with a tactical disadvantage.

to true, it means the enemy ran away before the vehicle could fire either initially or again. Thus, the vehicle transitions out of the `move2MERenemy` or `FIRE` state into the `cannotCatchEnemy` state, where the vehicle stops and transitions outside the `ENGAGE_TA` state to a connective junction.

Enemy_Contact.ENEMY.ENGAGE_TD In engaging with a tactical disadvantage, the vehicle's one hope is to move as fast as possible to its own maximum effective range, get the enemy in weapons lock, and fire upon the enemy quicker than the enemy can fire upon the vehicle. During this whole time, the vehicle will be inside the enemy's sensor radius range, and therefore, it becomes a question of who can get weapons lock the quickest and ultimately fire the quickest on the other.

In the `ENGAGE_TD` state, as depicted by Figure 4-11 there are four substates: `wait4Lock`, `move2MERveh`, `FIRE`, and `cannotCatchEnemy`. As in `Engage.TA`, the default path is to a connective junction with a slightly different logical condition: if $\|vehicle.position - enemy.position\| < vehicle.range$, then transition to `wait4Lock` state; else transition to `move2MERveh` state. Here, the vehicle, at the disadvantage, is already close enough to the enemy to put the enemy in the vehicle's weapons cone. Thus, the vehicle transitions to the `wait4Lock` state, stops, and turns to the enemy, hoping to get radar lock quicker to transition to the `FIRE` state. Else, if upon contact,

the enemy is not close enough to lie within the vehicle’s sensor radius range, the vehicle transitions to the `move2MERveh` state, and flies as quick as possible to put the enemy within its weapons cone, wait for radar lock, and transition to the `FIRE` state. As in the `ENGAGE_TA` state, the `FIRE` state has the same two transitions, the self-transition and the transition to `cannotCatchEnemy`. Also, as in the `ENGAGE_TA` state, a transition occurs from `move2MERveh` to `cannotCatchEnemy` when *pursue* evaluates true.

Transition back to Nominal Mission There are three transitions out of the `ENEMY` state. The first transition is a super-transition from a connective junction within the `ENEMY` state. This connective junction combines a single transition from each of the engage and avoid substates. From the `AVOID` state, once the functions *clearEnemy_CW* or *clearEnemy_CCW* evaluate to be true, the vehicle has successfully circumnavigated the enemy’s sensor radius range and has clear line of sight to the next waypoint in its navigation catalog. Thus, it exits the `AVOID` state to the connective junction. From the `ENGAGE_TA` and `ENGAGE_TD` states, if the vehicle entered the `cannotCatchEnemy` substate, it stops, and at the next time step exits out of `ENGAGE_TA` or `ENGAGE_TD` to this connective junction. The second transition out of the `ENEMY` state is if the vehicle has successfully destroyed the enemy. The third transition out of the `ENEMY` state is if the enemy completely withdrew, which is defined as the condition $\|\overrightarrow{vehicle.position} - \overrightarrow{enemy.position}\| > 4 * enemy.range$.

In conclusion, each `ENEMY` state has the same three transitions to a connective junction. These three connective junctions combine into a single connective junction, which transitions from the `Enemy_Contact` to `Search` superstate. Specifically, it transitions inside `Search` to the `Move2Wypt` state.

4.4 Limitations of Experiments

Now that the experimental framework, methodology, and the dynamics of the simulation have been discussed, we present the following limitations for these experiments. Because the whole purpose behind the human-in-the-loop experiments and comparison to baseline AV behavior is to test an untried concept of learning and applying human tactics, it is important to honestly assess the experiment limitations.

4.4.1 Experts?

First, the human subjects chosen for the experiments are by no means experts in armed reconnaissance missions. Therefore, the following question must be asked: how do you learn expertise from a human subject that is not an expert? This is a fundamental concern. We propose to ask the following re-phrasing of the question: can a human learn the right tactics in order to win the game? The picture is that of the ten year-old boy who learns all the right moves and shortcuts to beat the video game. Can humans learn expertise? Absolutely [18], and human learning is far superior to machine learning. Yet, the question is how much training is necessary

to learn that expertise? In answer to this question, we performed two rounds of experiments. The first round of experiments began with two practice scenarios. The second round began with one practice scenario. However, we treated the entirety of the first round of experiments as pure training. The results were merely used as comparison in determining tactics from the second round. Therefore, a first limitation of the experiments was the use of human subjects who could not be considered as true experts. However, this limitation was offset as much as possible by allowing an extended session of training. Furthermore, there is the question of how much fidelity is then necessary to ensure the expert is accommodated enough to the environment to display expert decision making? This question will not be discussed here, but it poses thought for future work. Note the answer to the question of fidelity versus expertise highlights the fundamental concept of what expertise are we trying to learn.

4.4.2 Real-time Interaction

A major issue in developing the simulation was trying to interact with Simulink in real time. The Simulink software is designed to simulate a model as quickly as possible as a function of the computer's performance capabilities. After talking to many Simulink users at Draper Laboratory, it was clear that no one had used Simulink in a real time process. An initial search found a real-time blockset offered publicly by Leonardo Daga, which succeeds in slowing down Simulink to mimic real time by changing priorities in the operating system's processes [14]. During the simulation development phase, this blockset was used. However, another major complicating issue surfaced. In order to allow human-in-the-loop interaction with the simulation, all the necessary Simulink variables have to be transferred to a MATLAB file, which then plots all of them. The plotting scheme only works by assigning handles to all of the plotting objects. Then, at every time step, each handle is deleted and then reassigned and thus re-plotted at the new locations. Unfortunately, this process of moving many variables from Simulink to a MATLAB script and then plotting of all them tremendously slows down the simulation. As the simulation continued to develop and more and more variables were being transferred due to increasing complexity and robustness in engagement scenarios, the simulation began to slow down more and more until it began running slower than real time. Now the original problem of slowing down the simulation to run at real time speed was irrelevant, and Daga's blockset no longer helped.

4.4.3 Arbitrary Scaling

The result of not being able to regulate the speed of the simulation made realistic scaling for this simulation impossible. The speed of the vehicle over the terrain and the heading rate of the vehicle were arbitrarily chosen numbers. They were determined solely by how smoothly the vehicle appeared to maneuver in response to human inputs of the joystick. This, then, affected how much time was allotted to each scenario and how much terrain could be covered and vice versa. For example, we found that the time constraint of five minutes for the first round of experiments

was too long in relation to the vehicle's speed and the amount of terrain that could be covered. In effect, the time constraint did not constrain the problem, and the human subjects could easily search through all of the area with time remaining. Therefore, the amount of terrain to cover was increased and the total amount of time for the mission decreased to four minutes for the second set of experiments. However, now the vehicle's speed and heading rate had to be redefined by trial and error until both its maneuvering appeared smooth and it gave the human subjects a decent chance at making it through the entire search area. Furthermore, the speed and heading rate of all three enemies in these scenarios had to be made relative to the vehicle's speed and heading rate. For example, the tank must move slower than the AV, but the question was how much slower? If the vehicle's speed was arbitrary, then the tank's speed had to be strictly relative.

4.4.4 Wall-Clock Time versus Simulink Time

Another difficulty related to the processing speed of the simulation concerned how to compare human runs against baseline vehicle runs. Because the speed and heading rate of the vehicle were determined by the perception of their smoothness as seen in real time, it seemed logical that the time constraint of each scenario for the human subjects should exist in real time. Therefore, a stopwatch marked the time. The human subjects were only given four or five minutes for each scenario as marked in real time or what can be called "wall-clock time." However, it was found that for a fixed amount of wall-clock time, the number of time steps Simulink processed varied significantly. Now remember, the processing rate of Simulink for these scenarios, as perceived by the observer in real time, was slower than wall-clock time. Therefore, for every one second of Simulink time passed, wall-clock time had passed through multiple seconds. Yet, the issue was not that Simulink was running at a fundamentally slower rate than wall-clock time, but that Simulink was running at a fundamentally slower rate that also varied as perceived in wall-clock time. Sometimes the simulation would run quicker as perceived by the human, and sometimes it would run slower. Thus, a fixed amount of wall-clock time, say four minutes, corresponded to twenty-five seconds of Simulink time in some runs, thirty seconds in some runs, and even up to forty seconds in others. It was assumed that the explanation behind this varying Simulink rate was the difference in processing from scenario to scenario which was occurring behind the scenes in the host computer's operating system. However, now the issue was how to constrain the time allotted for the autonomous baseline vehicle runs which existed solely in the simulation. The best answer was for each scenario, to average over all the human subjects the total number of time steps in Simulink at scenario termination. The autonomous baseline vehicle runs were then constrained by the total number of time steps allowed in the simulation. Therefore, where the time constraint for the human runs was marked off in real time, the time constraint for the autonomous baseline vehicle runs was marked off in Simulink time. After the fact, it was determined that the best solution was to place the time constraint in the simulation because it was the processing rate of Simulink that could not be regulated. The procedure would be to run through a number of trial simulations. Each scenario

would be stopped after a fixed amount of wall-clock time and the total number of time steps in Simulink recorded. After several runs for each scenario, the recorded time steps would be averaged. This value would then be entered into the simulation as a time constraint which would exist in Simulink for both the human and autonomous vehicle runs.

4.4.5 Varying Processing Rates

One final word on the above discussion is a limitation of the experiments due to the above problem with Simulink processing rates versus human interaction in real time. For every scenario, the human subjects had a fixed amount of time to accomplish the mission objectives. This fixed amount of time was five minutes for the first round of experiments and four minutes for the second round of experiments. As previously discussed, the number of time steps in Simulink at the termination of each run varied significantly. In post-experiment data analysis, we realized that this variation in Simulink time gave an unfair advantage to some human subjects and an unfair disadvantage to others. The reason is that the vehicle dynamics are calculated within the Simulink world. Thus, irrelevant of wall-clock time, if the vehicle is programmed to fly at a fixed speed at every time step, a larger number of time steps in Simulink correlates to a farther distance flown by the vehicle. If five minutes of wall-clock time for human subject *A* resulted in thirty-five time steps in Simulink time and five minutes for human subject *B* resulted in thirty time steps in Simulink, human subject *A* had five extra time steps in Simulink to search through the terrain than human subject *B*. Therefore, human subject *A* has the unfair potential for a greater score. In data analysis, this limitation of the experiments could not be removed by simply finding the scenario with the shortest number of time steps and chopping off all the other runs at that time. The reason is that the human subject would not have behaved the same way if time had been more limited. Thus, this limitation could only be kept in mind while comparing human subject against human subject and interpreting strengths and weaknesses.

4.4.6 Human and Automation Differences

In comparing the human subject runs against the autonomous runs, both the baseline and improved AV runs, there are a few major discrepancies that need to be highlighted. First, the human subject responds to visual stimuli. For example, the only way that the human subject knows he has achieved radar lock on the enemy is when the enemy icon turns completely red. This, however, takes time to visually process and generate a conscious response, a well-known delay in human processing [23]. On the other hand, because the AV only deals with data, radar lock on the enemy is known and responded to instantaneously. This gives the AV a realistic advantage over the human. Second, there are no noise sources in the simulation; all calculations are based on truth data. This combined with another well-known fact that humans introduce noise into the system because they cannot fine tune the control inputs as well as a computer gives the AV another advantage over the human. Third, the

human subject has no *time2shoot* counter after acquiring radar lock on the enemy. Once the human subject consciously processes the visual stimulus of achieving radar lock on the enemy, he can depress the trigger button and fire on the enemy. On the other hand, the AV must wait for the *time2shoot* counter to run out before it can fire on the enemy after achieving radar lock. This gives the human subject a slight advantage over the AV. Fourth, in the MATLAB plotting script, there is a logical expression that states if the vehicle has already detected the enemy (it is assumed that the engagement, if one occurred, was broken off and the vehicles separated) and the distance between the enemy and the vehicle closes to within twice the distance of the vehicle's sensor radius range, the enemy is plotted on the display. The rationale behind this expression was that if the computer on board the human subject's vehicle has already detected the enemy, it will retain some knowledge or make some prediction about the enemy's path. The rationale holds if the enemy is static, continues to move in the same path as its motion at separation, or the time between separation and re-acquiring is small. Regardless, this logical expression gives the human subject an advantage of visually reacquiring and responding to an enemy presence before the AV can. Fifth, the AV's behavior is purely reactive. There is no planning component to its behavior, whereas the human subject can visually scan the terrain map and plan out the best path to take. The AV must follow the waypoint list given to it.

4.4.7 Learning, Training Effect, and Assumption Violation

The training effect was described earlier (see section 4.2.5) as the unalterable learning side effect that presenting a number of cases to a human subject. It is important to consider, however, the relationship between the training effect and the limitation mentioned above about using human subjects who must spend time learning the simulation environment and tasks. Is there a difference between the training effect and allowing human subjects enough practice to become "experts" in the simulation? In the first round of experiments, there was no difference. We wanted the human subjects to come away from the first round as from a training session. Therefore, there was no randomization of the cases as presented from subject to subject during the first round. They saw the cases in the same order so that they all had an equal learning curve. In the second round of experiments, though, the cases were randomized between subjects to cancel out any learning effects. However, the limitation is that there was still plenty of learning occurring during this second round. This was most clear in enemy engagements. The necessity to discretize engagements into levels of interaction (see 4.2.5) violated the assumptions of the human subjects of how engagements should occur, even after the first round of experiments. In essence, the human subjects displayed an overwhelming tendency to try and "strafe" the enemy. In this maneuver, they flew high speed at the enemy and pulled the trigger as they flew over the enemy. However, this strafing method never allowed enough time for the enemy to be put in radar lock, and thus, they were never able to hit the enemy by strafing. Thus, during the second round, there was still a lot of learning occurring as to how to beat the enemies.

Chapter 5

Results of Human-in-the-Loop Experiments

To analyze the results of the human-in-the-loop experiments, we begin by scoring the performance over all cases. With five human subjects, a first round of seven cases, and a second round of eight cases, there were a total of seventy-five cases, including practice scenarios, to be scored in these experiments. Then, we compare the human performance against the baseline AV to ensure that there are definite areas where the humans outperform the baseline AV. Once we identify these areas of human superiority, we begin looking for human strategies by first comparing all human subjects' performance against each other. We are seeking the best tactics from the humans, and thus, we use the performance metrics to filter out those cases which resulted in the strongest performance. To identify strategies, we analyze the action sequences, think aloud reports, and surveys together, for they all make significant contributions to learning strategies. We then analyze the cognitive processes behind the decisions made to help further clarify the learned strategies. Finally, we create statechart representations of the tactics, augment the AV with the improved behavior, and test both the baseline and improved AV behavior against a large sample of cases with randomized parameters.

5.1 Performance Metrics

The total score for each case is composed of two main parts, a search and engage score. Both the search and engage scores are composed of several elements, and most of these elements are unitless ratios so that they may be linearly combined into a total score without mixing up units.

5.1.1 Search Score

The search score is obtained by the summation of the following elements: *areaMetric.AirCorr*, *areaMetric.CA*, *timeMetric.CA2AirCorr*, *timeMetric.exposure2total*, and *VelPenalty*. First, the *areaMetric.AirCorr* score is a reward for the percentage

of area covered out of the entire air corridor. Second, the *areaMetric.CA* score is a penalty for the percentage of area not covered out of the entire critical area. Because the critical area is, by definition, a more important piece of terrain - a fact emphasized in the presentation of the scenarios to the human subjects (see Section 4.2.7) - it is assumed the human subjects will cover the entire region. If they do not, they receive a penalty equal to the percentage not seen. Third, the *timeMetric.CA2AirCorr* score is a measure of how much time was spent searching inside the critical area against how much time was spent searching inside the air corridor. The numerator of this ratio is the time inside the air corridor subtracted from the time inside the critical area. The denominator is the sum of the time inside the air corridor and the time inside the critical area. The human subjects were told that they should spend a greater amount of time in the critical region than the air corridor. Fourth, the *timeMetric.exposure2total* score is a penalty for the percentage of time the human subjects were exposed to the enemy out of the total scenario time. To “be exposed” is to be within the enemy’s weapons cone radar. Finally, the *VelPenalty* score is a small penalty aggregated during the scenario if the human subject moved at full speed. The purpose of this score was to simulate a limited fuel supply. The joystick was setup so that there was a nominal search speed by pushing the lever forward (see Figure 4-3). If the human subject needed to go faster, he could push the joystick forward. If the human subject needed to slow down, he could pull the joystick backward. However, the human subject was told in the introductory two-page description of the scenarios (see Appendix A) that the longer the human subject moved at full speed (both lever and joystick forward), the more penalty he would incur. The penalty is equal to a constant α multiplied by the position measurement of how far forward the joystick is pressed.

5.1.2 Engage Score

The engage score is determined by the summation of the following elements: *exposureMetric.enemy*, *weaponsLockMetric.enemy*, *firingEfficiency*, *healthVEH*, *healthEnemy*, and *VelPenalty*. First, the *exposureMetric.enemy* is a measure of how much time the human subject had the enemy within his weapons cone against how much time the enemy had the human subject within its weapons cone. Second, the *weaponsLockMetric.enemy* is a measure of how much time the human subject had the enemy in radar lock against how much time the enemy had the human subject in radar lock. Both the *exposureMetric.enemy* and *weaponsLockMetric.enemy* can be either a reward or a penalty. Third, *firingEfficiency* is a reward for the ratio of the total number of hits over the total number of shots taken over all enemies. Fourth, if the health of the vehicle is less than two, where a health of one is damaged and a health of zero is destroyed, *healthVEH* is a penalty of -1 if damaged and -2 if destroyed. All simulated entities begin with a health of two. If the simulated entity takes a hit from any other entity, it loses one health point. This represents a damaged state. A second hit effectively kills the entity. Fifth, *healthEnemy* is a reward of $+1$ for a damaged enemy and $+2$ for a destroyed enemy, and it is summed over all enemies damaged or destroyed. Sixth, *VelPenalty* also applies to the engage score, because

SEARCH	
Score Element	Score Range
$areaMetric.AirCorr = \frac{areaCovered.AirCorr}{totalArea.AirCorr}$	[0, 1]
$areaMetric.CA = \frac{areaCovered.CA}{totalArea.CA} - 1$	[-1, 0]
$timeMetric.CA2AirCorr = \frac{time.CA - time.AirCorr}{time.CA + time.AirCorr}$	[-1, 1]
$timeMetric.exposure2total = \frac{timeExposed2Enemy}{totalTime}$	[-1, 0]
$VelPenalty = (\alpha)(joystickPos)$	$[-\alpha, 0]$
ENGAGE	
Score Element	Score Range
$exposureMetric.enemy = \dots$ $\dots \sum_{i=1}^3 \left(\frac{enemyExposedTime - vehicleExposedTime2Enemy}{enemyExposedTime + vehicleExposedTime2Enemy} \right)$	[-3, 3]
$weaponsLockMetric.enemy = \dots$ $\dots \sum_{i=1}^3 \left(\frac{TimeVehicleHasLockOnEnemy - TimeEnemyHasLockOnVehicle}{TimeVehicleHasLockOnEnemy + TimeEnemyHasLockOnVehicle} \right)$	[-3, 3]
$firingEfficiency = \sum_{i=1}^3 \left(\frac{numHitsEnemy}{numShotsEnemy} \right)$	[0, 3]
$healthVEH = vehicle.health - 2$	[-2, 0]
$healthEnemy = \sum_{i=1}^3 (2 - enemy.health)$	[0, 6]
$VelPenalty = (\alpha)(joystickPos)$	$[-\alpha, 0]$

Table 5.1: Performance Metrics.

there was no way in the simulation to define a clear line and extract those times when the human subject used full speed in searching only or in reacting to an enemy only. Furthermore, because *VelPenalty* is supposed to simulate limited fuel, it is counted in both the search and engage score.

5.1.3 Summary and Bounds

Table 5.1 summarizes the scoring elements and each score range. It is hard to predict, based on these metrics, the upper bound for scoring. For the search score, if the human subject covered the entire critical area and air corridor, split his time evenly between the two, was never exposed to the enemy, and never used full throttle, the human subject would receive a search score of 1.0. On the other hand, if the human subject covered the entire critical area, half of the air corridor, spent twice as much in the critical area than the air corridor, was never exposed to the enemy, and never

used full throttle, the human subject would receive a search score of 0.833. However, the relationship between percentage terrain covered and the time it takes to do so depends on the vehicle's speed, the total amount of area to search, and the total amount of time allotted. Thus, it does not seem realistic to obtain a search score greater than 1.0.

Likewise, the upper bound on the engage score is hard to predict. If the human subject could engage and destroy all three enemies without ever being exposed to the enemy (inside the enemy's weapons cone) and did so with perfect firing efficiency and without using full velocity, the human subject could receive a score of 15.0. However, this is not possible, because two out of the three enemies have a weapons radar range equal to or greater than that of the vehicle. To engage these enemies requires exposing the vehicle to their weapons radar. Furthermore, one of the enemies is faster than the vehicle, and to engage it requires moving at full speed. If the human subject could perfectly engage one enemy, say the easiest to kill, then the human subject could obtain a score of 5.0 from that one engagement. Because there is not much time in the scenarios to engage all three enemies, it appears that very good engagement scores would be close to 5.0 with some variation.

5.2 Human and Baseline AV Scores

Figure 5-1 depicts the total scores for each case in round 1. Figure 5-2 depicts the total scores for each case in round 2. Note that for human scores, each figure depicts the average, variance, absolute maximum, and absolute minimum of human scores. The average human score over all humans is shown by the height of the corresponding "Search Human" and "Engage Human" bar. The variance across all humans is given by the extended vertical line, where the height of the line from the bar to its horizontal cap represents one standard deviation. The absolute maximum and absolute minimum scores are displayed by the blue and red horizontal lines, respectively. Also, note that the baseline AV performance constitutes the untrained set of tactics. They were derived without any human subject inspiration, other than the engineering sense of the designer (see Section 4.3.4). Before discussing the results, it is important to highlight the differences between the two rounds. In round 1, the human subjects were first exposed to the simulation environment, and they were given two practice rounds before the five cases. In round 1, each case lasted five minutes, there were only a maximum of two enemies in each case, and the number of air corridor matrix elements (i.e. - how many elements of the entire terrain database did the air corridor cover) ranged from 3700 to 5000 out of the total 14,400 to 19,600 elements in the entire terrain matrix. In round 2, each case lasted only four minutes, there were a maximum of three enemies in each case, and the number of air corridor matrix elements ranged from 6000 to 8500 out of a total 32,400 elements in the entire terrain matrix. The larger terrain to cover was slightly offset by an increase in the vehicle's velocity gain for round 2 to keep the simulation appearance smooth (see Section 4.4.2). However, the human subjects consistently ran out of time before they could cover all of the air corridor and critical area in round 2, whereas in round 1, there




Figure 5-1: Human and baseline AV performance for the first round of experiments.




Figure 5-2: Human and baseline AV performance for the second round of experiments.

was plenty of time to backtrack and overlap. The time limit, then, actually became a constraint in round 2, which was the desired goal. A time limit is desirable in these experiments because it forces the human subject to seek creative solutions because there is not enough time for the straight-forward ones. Therefore, the problem was overall significantly harder in round 2 than in round 1.

To begin analyzing these two figures, start with the total scores for each case. The first question to ask is if there is something to be learned from the human subjects. The answer is undoubtedly *yes*. For the total scores, the best human performance beat the baseline AV in every case. Remember that we are interested in discovering the best decisions and strategies to apply to the AV, not just the average. In round 1, the total baseline score beat the human average in cases 1, 4, and 5. However, note that except for case 5, the maximum human total score significantly beat the AV baseline. In round 2, the total baseline score was higher than the human average in case 2 only. In fact, the human average total score was significantly higher than the baseline's total score in cases 3–7. Though the problem was harder in round 2, it therefore appears as if the human subjects performed better. The second question, then, is in what ways did the human subjects outperform the baseline AV.

To answer the question of this increase in human performance, consider first the search scores. In round 1, for positive search scores, the baseline beat the human average and nearly equalled the best human performance in searching. In round 2, for three out of seven cases, the baseline search score was greater than or equal to the best human search score. It is important to note, however, that in case 6, there were no enemies, and therefore, the humans exhibited better pure search performance on average than the AV. All told, though, it does not appear that search scores account for the higher human total scores in round 2.

Next, consider the engage scores. In round 1, the baseline engage score beat the average human engage score in cases 1,4, and 5. Only in case 5, however, was the baseline engage score higher than the best human engage score. In case 2, the baseline AV did not run across any enemies, and in case 3 the baseline AV was killed rather lopsidedly. In round 2, the baseline engage score only beat the human average engage score in case 2, and even then, the maximum human engage score was significantly higher than the baseline AV engage score. More importantly, the baseline AV did not engage any enemies in cases 3,4, and 7, and on average, the human engaged enemies in these cases very successfully. Therefore, the significantly higher average human total score noted in cases 3–7 above, are mainly due to baseline AV non-engagements and successful human engagements. However, note in case 5, the AV had a very low engage score, and in case 6, there were no enemies. When the baseline AV did engage enemies in round 2, the best human score either equalled or significantly surpassed the baseline AV engage score. This was not true for round 1, where the baseline AV beat the average human engage score three times and the best human engage score once as opposed to only beating the average human engage score twice in round 2 and never beating the best.

A high-level interpretation of the results, then, reveals that the human subjects engaged enemies more successfully than the baseline AV in round 2 versus round 1. This is encouraging because it means that the human subjects learned good reactive

tactics to pop-up enemy threats. One criticism before moving on, though, is the large variance in human engage scores. The large size of the one standard deviation away from the mean is not in and of itself the issue. Note that this variance comes from the set of five observations, which is the set of five human subjects performing each case. Large differences in scores between subjects is not altogether bad. It means different subjects have differing levels of success in engaging enemies. This is to be expected, and especially when the number of observations is so low. It is not the large standard deviation, then, which is necessarily the problem. Rather, the criticism is that the large standard deviation in round 1 did not diminish at all in round 2. The human subjects did not learn consistently better tactics, as a whole, between engagements in round 1 and round 2. Again, this does not mean that individual human subjects obtained consistently better scores through improved tactics in round 2. In fact, one interpretation of the large variance in round 2 could be that all human subjects learned better tactics to the same degree, and thus the scores showed stronger improvement but the variance did not change. The continued presence of large variances in engage scores in round 2 is disappointing because it shows that the human subjects, as a group, did not converge on better tactics. We would have liked to see that as the human subjects' experiences increased with each case, the variance decreased. Therefore, we must be selective in what cases and which subjects we choose to learn tactics from. This discussion highlights the learning limitation, mentioned in section 4.4.7.

5.3 Finding Human Expert Performance

To help determine exactly how the humans performed better, Figures 5.3.1, 5.3.2, and 5.3.3 break down the human performance across elements of the engage scores, subjects, and enemies. Note that in the legend of each figure, all five subjects are identified by "S1, S2, S3, S4, and S5," a color bar, and the number of observations in this statistical set given in the parentheses. Also, note that the *VelPenalty* scoring element is not present in these figures for three main reasons. First, in post-processing there was no clear line in how to distinguish between the use of maximum velocities for searching purposes and for engagement purposes. Second, because much larger engage scores are obtainable than search scores, *VelPenalty* does not noticeably affect the human subjects' engage scores like it does in the search scores. Third, the purpose of *VelPenalty* was to simulate limited fuel, but it can be argued that when the human expert must respond to a pop-up threat in a life-or-death engagement situation, limited fuel is not too important. Standard operating procedures protocol mandates the use of fuel reserves in all flights [80]. In a life-or-death engagement then, a small *VelPenalty* is not important. Finally, the process is to use the performance scores to identify those cases and scenarios where the best tactical decision making occurred. To that end, the verbal reports, surveys, and video recordings of the human subjects' actions will all be combined to form an interpretation of the best tactical decision making. Note, that just one of these interpretative elements cannot, by itself, reveal the human strategies. They must be integrated together to learn tactics.

5.3.1 Fighting the Tank

The enemy tank is slower, less maneuverable, has a smaller sensor radius, and has less probability of hitting another vehicle than the human subjects' vehicle (see Table 4.3). The only advantage the tank has is that it has the smallest *time2shoot* constraint. After the tank acquires weapons lock, it can fire the quickest on the vehicle, even faster than the AV can fire at the tank. However, as noted in Section 4.4.6, the humans did not have a *time2shoot* constraint other than the inherent human delay in processing cue information. Figure 5.3.1 depicts the human performance in engaging the tank. It is easy to notice that S4, the fourth human subject in the experiments, greatly outperformed everybody else. First, none of the subjects were ever damaged or destroyed by a tank, and thus no penalty was added due to *healthVEH*. Next, S4 killed the tank every time he engaged one. Thus, there is no variation in the full two points reward of *healthTank*, depicted by the lack of a one standard deviation vertical line extending from the orange bar. For exposure, S4 was rarely ever exposed to the tank, which indicates he was able to use his larger sensor radius to his advantage. Even if S4 was slightly exposed by the tank, he always corrected the situation so that he was never put into weapons lock by the tank. S4 also had the greatest firing efficiency, and for all four tanks engaged, he only missed once achieving a total efficiency of $\frac{8}{9} \approx 0.889$. Finally, S4 had a nearly perfect score of 5 points against every tank encountered. Note, that S4 had the greatest score and the least variance in every element composing the total engage score. Therefore, the focus in learning the best reactive tactics in engaging the tank is completely on S4's decisions. We then turn to the verbal reports, surveys, and video recordings of S4's encounters with tanks.

Verbal Reports (transcribed from simulation recordings)

Tank Engagements, S4 Verbal Reports
Case 7
Looks like that's a stationary object there ... [kills tank] That was easy, ok
Case 5
Ok, that's not good Alright, looks like I ... [kills tank]
Case 3
Stay centered Ok, nice I want to stay out of the range the whole time ... [kills tank]
Case 4
Yeah, well, let's try and go for him this time ... [kills tank] ... ok

As can be seen, S4 did not verbalize too much during engagements with the tank. In fact, S4 verbalized the least out of all five subjects during all experimentation. There are two main points to draw from his comments. First, in case 7, he called the tank a stationary object. In that case, the tank did sit and wait during the scenario until




Figure 5-3: Human performance against enemy tanks.

contact was made, rather than following a patrol route until contact (see Section 4.3.2). At the point of contact, the tank began to move, but did so slowly. S4's comment highlights his ability to use superior speed and maneuvering to kill the tank ... "That was easy." Second, in case 3, he reveals his strategy of consciously staying outside the enemy tank's weapons range. S4 employs his superior sensor radius to standoff from the enemy tank and destroy it.

Surveys

At the end of the second round of experiments, the human subjects responded to the following question regarding all three enemy types:

To the best of your knowledge, are you prone to engage or avoid the following three threats? What factors are involved in either engaging or avoiding? If you had to plan your strategy ahead of time, what actions would you take if you unexpectedly ran into each of these three threats in a typical scenario? (see Appendix A)

For enemy tanks, S4 wrote, "These are easy [to kill] in that they are slow and have small targeting areas. I would typically engage them because they didn't involve much of a threat or repositioning." Most interesting in this answer is S4's concept of "repositioning." Presumably, the concept of not having to reposition is equivalent to not giving chase to the enemy tank. Thus, if S4 encountered a tank, he could quickly kill the tank without diverging far from his originally intended search path.

Actions

In analyzing all four engagements with the tank, there was no real difference in how S4 proceeded from detection of the tank to destroying the tank. The process was almost identical in every case. First, upon detection, S4 either stopped and came to a hover right away next to the tank, or S4 moved a little past the tank on the original course and then stopped and came to a hover. Most importantly, every time S4 quickly came to a hover, he remained outside the tank's weapons range at a safe standoff distance. Next, S4 continued to hover and turn until the tank was within his weapons cone. Then, S4 tracked the enemy tank by simple heading changes as necessary until achieving weapons lock. Finally, S4 shot two rounds and quickly killed the tank.

Statechart Representation

Figure 5-4 depicts the human-inspired engagement tactics against the enemy tank. The first major difference in this statechart logic and the baseline AV Engage.TA logic (see Section 4.3.4 and Figure 4-10) is the Fly.By state. If the vehicle detects the vehicle and is inside the tank's weapons range, the vehicle transitions to flying past the enemy tank. In the baseline statechart, the vehicle would have transitioned to a moveAway state. In the moveAway state, the vehicle would always change its

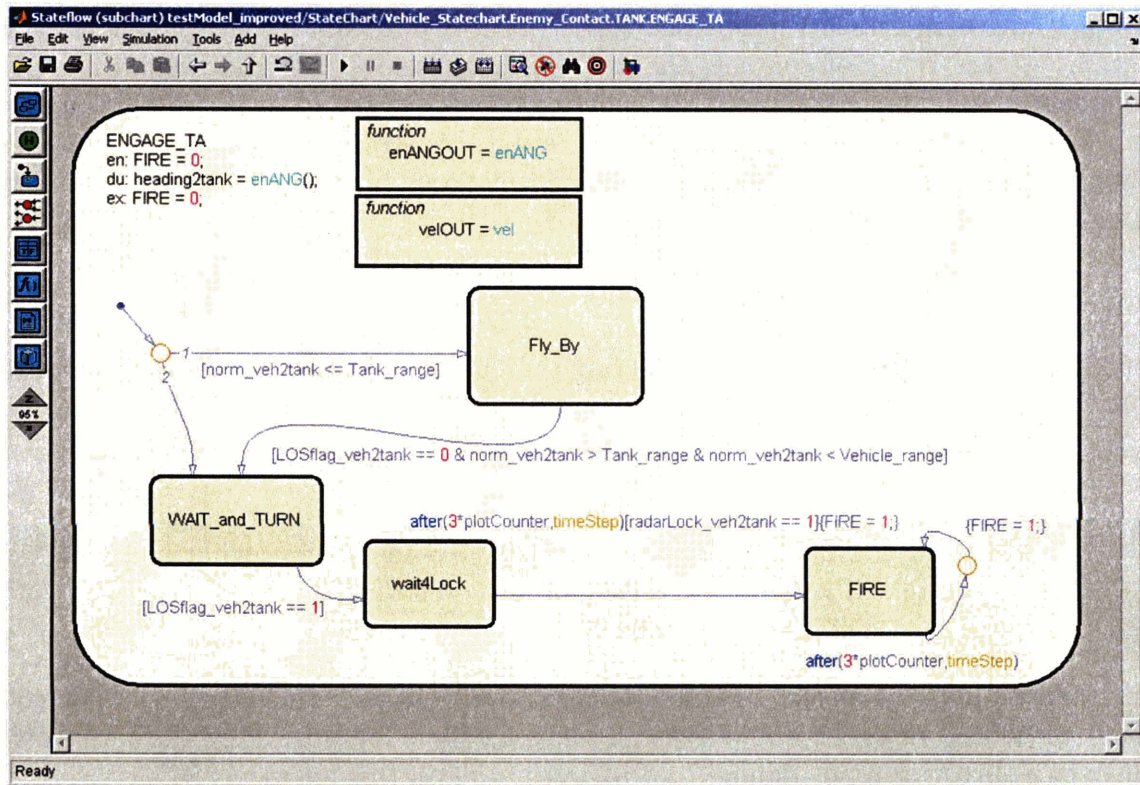


Figure 5-4: Human-inspired statechart to engage enemy tank.

heading to move away from the tank. In the Fly.By state, the vehicle intentionally moves toward the tank in the beginning to take advantage of the vehicle’s superior speed so that it can set up a standoff distance to fire upon the tank behind it. The logic in flying past the enemy tank is given by the true/false flag of the vehicle’s line-of-sight to the tank, *LOSflag.veh2tank*. Simply put, if *LOSflag.veh2tank* = 0 then fly away from the tank, else if *LOSflag.veh2tank* = 1, then fly towards the tank. At the beginning of detection, the tank lies within the vehicle’s weapons cone, and *LOSflag.veh2tank* = 1. The vehicle knows it has then flown past the enemy tank if *LOSflag.veh2tank* = 0. Once the vehicle has flown past the tank and is outside the tank’s weapons range but the tank is inside the vehicle’s weapons range, the vehicle transitions to the Wait.and.Turn state. In this state, the vehicle comes to a hover and turns and constantly tracks the tank with its weapons cone. Once the vehicle achieves weapons lock, it fires.

The speed logic of the Wait.and.Turn state is the second major difference in the human-inspired statechart. In the baseline statechart, the vehicle transitioned to a move2MERtank state. The speed logic of this state, given by the *vel* function, is depicted in Algorithm 4: Note that the units do not match on either side of the “velocity = ” equations. Velocity is being set equal to units of position. This again is a limitation in the simulation of arbitrary scaling, as discussed in Section 4.4.3. In order to keep excessively large velocities being passed to the vehicle dynamics when $\|vehicle.position - tank.position\|$ is large, velocity is first passed through a saturator

```

if  $\|vehicle.position - tank.position\| > vehicle.range$  then
  velocity =  $\|vehicle.position - tank.position\|$  ;
  else if ( $\|vehicle.position - tank.position\| \leq vehicle.range$ ) and
  ( $\|vehicle.position - tank.position\| > tank.range$ ) then
    velocity =  $(\frac{1}{2})\|vehicle.position - tank.position\|$ ;
    else if ( $\|vehicle.position - tank.position\| \leq vehicle.range$ ) then
      velocity = -1;
    end
  end
end
end

```

Algorithm 4: Speed logic for baseline AV behavior.

block whose upper limit is given by the defined maximum velocity, given in Table 4.3. Thus, in the baseline statechart, the vehicle is always moving, even when it achieves the proper standoff distance. What was learned from S4, however, is to simply stop. The tank is not fast enough to move sufficiently far away to break line-of-sight in the time it takes the vehicle to get weapons lock. Thus, once the vehicle achieves the proper standoff distance, the human-inspired tactic is to stop, hover, and turn to fire upon the tank.

5.3.2 Fighting the UAV

The parameters of the enemy UAV are almost all identical to the human subjects' vehicle. It has the same maximum velocity, equal sensor radius, and equal probability distribution of hitting another vehicle. There are two main differences. First, the enemy UAV operates at maximum speed at all times outside of engagements, and thus the humans must maintain full speed to simply keep up with it. Once the UAV would change modes from pursuing the vehicle to moving to a holding pattern, it would move away very quickly. The human vehicle could not catch it, but could only merely maintain a relative distance while pursuing it. Second, the enemy UAV was less maneuverable with a slightly slower heading rate, which is its main tactical disadvantage. Figure 5.3.2 depicts the human performance in engaging the UAV. Unlike the tank, it is much harder here to distinguish the best human performance. Note that the bars will not be displayed if the score for the specific performance metric equals zero. Only S5 had a positive average total engage score, and S4 had the least negative total engage score. Furthermore, only S4 and S5 took damage from a UAV. S2 and S4 damaged a UAV, while S5 destroyed a UAV. Thus, S1 and S3 never committed to engaging a UAV. Interestingly, only S4 had a positive average exposure score, while S2 was slightly negative. All subjects were put in weapons lock by the UAV longer than they had weapons lock on the UAV. On average, S2, S4, and S5 only hit a UAV once for every three shots taken. In summary, out of the twenty-two encounters with enemy UAVs, only two resulted in damaging a UAV and only one in destroying a UAV. There are only three cases, then, of out twenty-two that can be examined with the goal of learning successful engagement tactics. Also, S2 had the




Figure 5-5: Human performance against enemy UAVs.

most encounters with enemy UAVs, never sustained damage, damaged a UAV once, and stayed fairly neutral in exposure, i.e. - not being exposed by a UAV longer than he could expose it. Thus, it appears that S2 knows how to successfully avoid an UAV, and the next step is to understand these scores in light of the verbal reports.

Verbal Reports

UAV Engagements
Case 2, S4 Verbal Reports
Yeah, well, I'm not going to get to the rest of this, which is fine I definitely want to hit the CA Get ready for some bogeys ... Whoa, get out of there, whoa buddy, whoa buddy, he's gunning for me Yeah, he's definitely targeting me ... [they damage each other] Whoa, what just happened Well I got hit first of all Ok, well, not successful Well, that's ok
Case 2, S5 Verbal Reports
Approaching the terrain features, so I'm more cognizant of my maneuver desires I know that I can go to the left in the event of contact in the area but not to the right Corner here, whoo, ok, have enemy contact Looks like the same enemy UAV Approaching now back up into here ... [kills UAV] Plucked him, plucked him dead
Case 1, S2 Verbal Reports
Now that UAV I don't know how cautious I need to be around him It's going very slow ... [damages UAV] Is he running away now? ah coward

First, in case 2, S4 had decided that it was time to move from the air corridor to the critical area. The intelligence report had stated that there was a very good chance of enemy contact within the critical area. Therefore, he wanted to “get ready for some bogeys.” However, he was not completely prepared for what would happen next. When he made contact with the UAV, the UAV began to chase him: “Whoa, get out of there.” S4 had never been continually targeted and fired upon by an enemy UAV. S4 was a little confused on the exact nature of the sequence of events, but he broke away with both he and the UAV being damaged. S4 summarized by stating the engagement was “not successful” but “ok.” Second, also in case 2, S5 verbalized his maneuvering restrictions while flying within the critical area due to the presence of a terrain obstacle to his right. He runs into the UAV, and after only a little time, kills the UAV pretty easily, “plucked him, plucked him dead.” Third, S2 confessed that he did not “know how cautious” he should be around the UAV. S2’s knowledge of successfully engaging UAV’s was very limited, and in fact, this was the only time S2 was able to damage an enemy UAV. Also, during S2’s engagement with the UAV the simulation slowed down significantly, which was a limitation discussed before (see

Section 4.4.5). This may have actually been an advantage for him because he was able to process the consequences of his inputs more deliberately than before.

Surveys

In response to the survey question about UAV engagements, S4 stated that the UAVs, “were fast and often tried to track you. Initially I engaged them, but this proved to be a waste of time, generally speaking.” S4 had tried to engage UAVs, but the UAVs were so quick that S4 was never close enough for long enough to fire upon them. The one time S4 did damage a UAV, was when the enemy UAV tracked him, and S4 finally turned around and shot back at the UAV. S4 does not conclusively state how or if he would engage a UAV in the future, but he mostly likely would avoid them because it took too long to be beneficial to the overall mission of air corridor reconnaissance. To engage “proved to be a waste of time, generally speaking.” S5 answered that he would, “engage if given opportunity, but not a high probability engagement (i.e. - I will have to chase him) and not my [primary] mission.” S5 does not qualify what an “opportunity” for engagement looked like, but he does state that the probability of successfully engaging the enemy is not high. It is interesting that S5 qualifies a high probability engagement by stating “I will have to chase him.” This most likely reflects the fact that every time S5 chased a UAV, he was not able to successfully damage it. A successful engagement is only somewhat probable, not because the enemy UAV is so lethal and S5 expects to be fired upon, but because the enemy UAV is so fast that S5 does not expect to catch it. Furthermore, S5 agrees with S4 in that spending time to chase the enemy UAV detracts from the “primary mission.” Lastly, S2 wrote that he was, “prone to ignore [them]. [They were] too fast for me to shoot. [Therefore, you should] ignore, but don’t get shot.” S2 advises to ignore the enemy UAVs. However, it must be an active mode of ignorance to keep from being shot. S2 had more contacts with enemy UAVs than anyone else with a total of seven, and S2 learned how to quickly maneuver to avoid and ignore them more than learning how to damage or destroy them.

Actions

In S4’s engagement with the enemy UAV in case 2, S4 performed the following sequence of actions. Upon detection, S4 first turned away from the UAV a full 180°. As S4 tried to move away, the enemy UAV gave chase and tracked S4. S4 then quickly maneuvered again, but he turned towards a terrain obstacle with the enemy UAV still giving chase. S4 was now stuck between the terrain obstacle and the enemy UAV. Thus, he stopped and turned to face the UAV. By the time S4 had turned completely around to put the UAV in his weapons cone, the UAV had maintained line-of-sight and achieved weapons lock on S4. They then proceeded to shoot and damage each other. As soon as S4 shot the UAV once, he moved past the UAV. While quickly flying past the UAV so that he would not get shot again, S4 tried to shoot the UAV. Yet, he had already passed by the UAV and missed.

S5 was the only subject to kill an enemy UAV. The sequence of events in S5’s

engagement with the enemy UAV in case 2 is as follows. At initial contact, S5 continued to move past the enemy UAV not deviating from his southward course. S5 took time along this path to observe the enemy UAV's actions. After breaking contact, the enemy UAV turned away from S5 and moved northwest. However, this northwest track took the enemy UAV into a terrain obstacle. Because the enemy UAV had the same obstacle avoidance logic as the baseline AV (see Section 4.3.1), it iterated its heading until it found a new path away from the obstacle. Unbeknownst to S5, this obstacle avoidance path moved the enemy UAV back towards S5. Thus, when S5 turned north back towards the UAV, the enemy UAV had now begun moving towards him. Flying directly towards each other, S5 stopped his vehicle and came to a hover inside the enemy UAV's sensor range. S5 was able to maintain line-of-sight, achieve weapons lock, and fire upon the enemy UAV quicker than the enemy UAV could fire upon him. Thus, S5 killed the enemy UAV pretty easily, but for two main reasons. First, the enemy UAV was in obstacle avoidance mode, and its obstacle avoidance path allowed S5 to directly intercept the UAV and kill it. However, even if S5 intercepted the enemy UAV, why would the UAV slow down at contact so that S5 could easily maintain line-of-sight and fire upon it? Why did the enemy UAV not keep on moving at its quick speed away from S5 but on the obstacle avoidance path?

When searching for an answer to this question, we found that the solution revealed a flaw in the original simulation design. The original concept for the enemy UAV's speed logic called for the enemy UAV to move at its maximum speed unless approaching the human vehicle or a holding pattern. If the enemy UAV was chasing a human subject, caught up to the human, but maintained its maximum speed, it would fly right past the human vehicle. Thus, the speed logic was for the enemy UAV to move at maximum speed until the human subject was within its sensor range. At that point, the enemy UAV slowed down to move at a rate proportional to $\|vehicle.position - UAV.position\|$. If the human vehicle did not move, the UAV would move at maximum rate until the human vehicle was within its sensor range, slow down at a rate proportional to the closing distance between the enemy UAV and human vehicle, and finally come to a hover directly over the human vehicle. The flaw in this design is that it did not account for those times when the enemy UAV was trying to move away from the human subject after breaking off an engagement. If the human subject had the enemy UAV in weapons lock or the enemy UAV had line-of-sight on the human subject, the enemy UAV would track and chase after the human. However, once both of those became false, the enemy UAV changed modes from pursuit to moving to a randomly designated holding pattern. Now say, for instance, that the human subject had flown into the enemy UAV's circle to fire upon it, and they both had line-of-sight on the other. Next, say that the human subject maneuvered to fly around the UAV and come around behind it, and that this maneuver caused both the human subject and enemy UAV to temporarily lose line-of-sight on the other. According to the enemy UAV's logic, it would now transition from pursuing the human vehicle to moving to a holding pattern. This should occur at maximum speed. The flaw is that it does not because the human subject and enemy UAV are so close together that $\|vehicle.position - UAV.position\|$ is still less than the enemy UAV's sensor range. Therefore, until the human vehicle is outside the enemy UAV's circle,

the UAV will still move at a rate proportional to $\|vehicle.position - UAV.position\|$ even though it is no longer pursuing the human, but rather fleeing the human. This then, is the answer as to why the enemy UAV slowed down enough upon contact with S5 so that S5 could kill it, even though the UAV was following an obstacle avoidance path. S5 intercepted the UAV and stopped right inside the UAV's circle. Thus, the UAV now moves along its obstacle avoidance path at a rate proportional to $\|vehicle.position - UAV.position\|$, which is slow enough for S5 to kill it.

The speed logic for the enemy UAV also allowed S2 to damage the UAV in case 1. In this scenario, the UAV was chasing S2. S2 began to turn and held the turn long enough to come in behind the UAV. He had the UAV in his weapons cone, but fired before achieving weapons lock. By the time he would have achieved weapons lock, S2 had flown past the enemy UAV. As soon as S2 realized he had missed and had flown past the enemy, he initiated another turn until he could come back around and fire again. The same sequence of turning, firing, missing, and overflying occurred. Because of the human vehicle's superior maneuvering rate, S2 turned quickly enough to come in behind the UAV, put the UAV within his weapons cone, fired, and missed because S2 did not wait long enough for weapons lock. Furthermore, because S2 did not slow down, he flew past the UAV before achieving weapons lock and began another turn. This happened a total of three times. Finally, on the fourth attempt, he actually stopped the vehicle, turned, achieved weapons lock, and damaged the UAV. However, at this point, S2 did something different. Rather than turning right away back around the UAV, he flew straight for a while. By the time he turned back around, the UAV had fled. "Is he running away now? Ah, coward!" Why did the enemy UAV flee now after having been damaged? There was no logic that told the UAV to stay in the fight until damaged. Instead, the difference was the distance between S2 and the enemy UAV. The four times S2 missed the UAV, he initiated a rapid turn while *inside* the enemy UAV's sensor radius. Furthermore, during the turn, S2 stayed *inside* the enemy UAV's sensor radius. Because S2 was inside the UAV's sensor radius, $\|vehicle.position - UAV.position\| < UAV.range$ evaluated to true, and the UAV did not move at maximum speed but at a rate proportional $\|vehicle.position - UAV.position\|$. Once S2 flew past the UAV and *outside* the UAV's sensor radius, $\|vehicle.position - UAV.position\| < UAV.range$ evaluated to false, and the UAV fled at maximum speed.

These three scenarios form the building blocks for an improved tactic of engaging enemy UAVs. First, S4's engagement highlights what not to do. If the human subject, while being chased by an enemy UAV, stops, comes to a hover, and turns to shoot, the human's vehicle will be damaged. That maneuver gives the UAV enough time to fire upon the human vehicle. Now the human subject might be able to rapidly fire two shots to kill the enemy UAV in exchange for being damaged. However, given the *time2shoot* counter built into the AV behavior (see Algorithm 1), the best case for the AV would be exchanging hits and then moving on, exactly like S4's actions. Second, S5's engagement brought out the speed logic flaw in the enemy UAV's design. While that is important, the other important decision was S5 deliberately coming to a stop within the enemy UAV's circle. This action of coming to a hover *once the enemy is in the vehicle's weapons cone* and not before (like S4) would have greatly

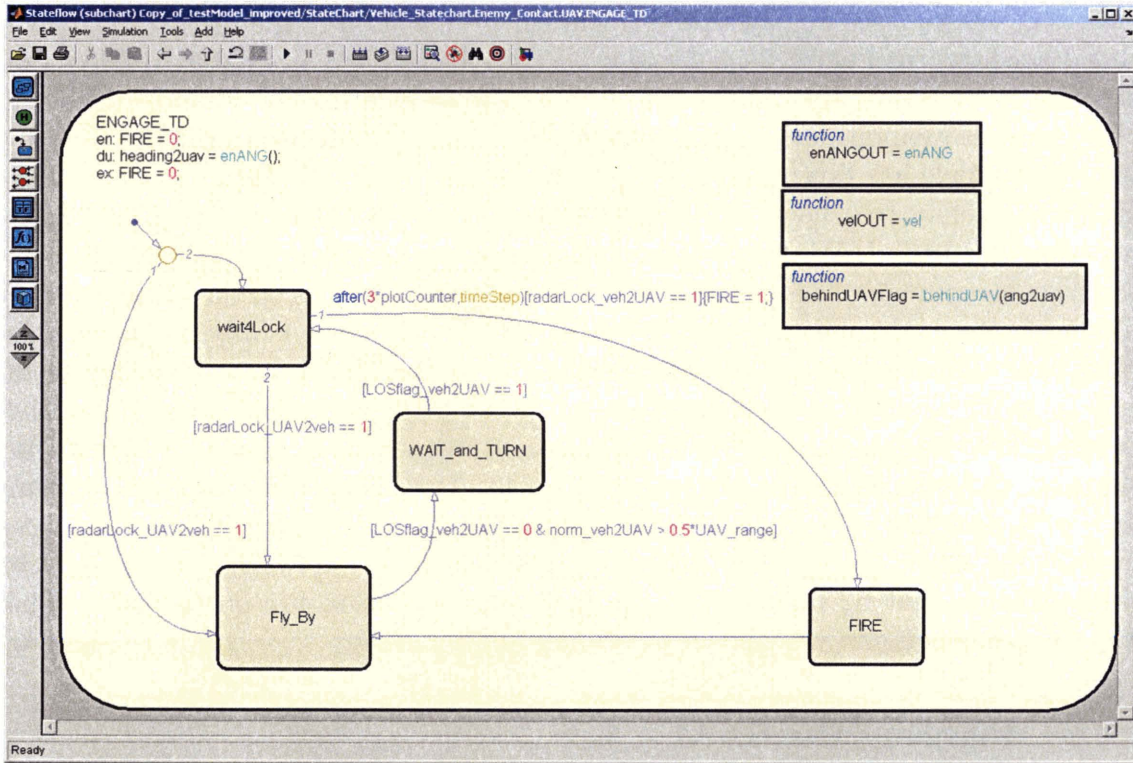


Figure 5-6: Human-inspired statechart to engage enemy UAV.

helped S2's engagement. In this third scenario, S2 spent four fruitless attempts at firing upon the enemy UAV because he never came to a hover. Yet, S2's actions highlight how to combine the superior maneuvering rate of the vehicle with turning inside the enemy UAV's sensor radius to place the human vehicle in a position to fire upon the enemy UAV. S5 made the right decision to stop and shoot, but S5 was able to simply intercept the UAV, which was moving along an obstacle avoidance path. Therefore, by the evaluation of these three cases, we can integrate them into a complete, human-inspired tactic, that draws out the strength or avoids the weakness of each.

There is one final point to discuss. Does the design flaw in the UAV's speed logic negate the inclusion of a tactic that takes advantage of it? We argue that it does not merely because the UAV's speed logic was in no way perceived by the human subject as a flaw. The human subject had no idea that the designers of the UAV's speed logic built in logic that created undesired behavior. In fact, this ability to unearth flaws in the enemy's performance characteristics is exactly the type of problem-solving behavior we desire for the humans to exhibit.

Statechart Representation

The statechart logic for engaging an enemy UAV is displayed by Figure 5-6. In this Engage.TD state, the vehicle defaults to a connective junction. If in the initial stages of detection, the UAV achieved weapons lock on the vehicle, the vehicle transitions to

the Fly.By state. This Fly.By state is identical to that used in the tactic for engaging tanks, as shown in Figure 5-4. If $radarLock.UAV2veh = 1$ evaluates to false, then the vehicle moves to the wait4Lock state. If while waiting for weapons lock, the UAV achieves weapons lock on the vehicle, the vehicle transitions to the Fly.By state. Once the vehicle has flown past the UAV, but while *inside the UAV's sensor radius*, the vehicle comes to a hover and turns towards the UAV. Once the vehicle achieves line-of-sight on the UAV, the vehicle then transitions to the wait4Lock state. If the vehicle achieves weapons lock on the UAV before the UAV does on the vehicle, it then fires upon the UAV and transitions back to the Fly.By state to begin the sequence over. Note that the purpose of always returning to the Fly.By state is to expose the slow speed of the UAV as well as to use the superior maneuvering rate of the vehicle to gain the best tactical position to fire. Because the UAV is more of a threat than the tank, the vehicle cannot just wait in one location to fire twice upon the enemy UAV and then move on. This was the failure of S4, and the strength of S2 to keep on moving, which he did even after damaging the UAV. However, it was only on the fourth try that S2 finally did stop and wait for lock to fire on the UAV. That stop-to-shoot sequence was shown by S5 and applied here.

5.3.3 Fighting the SAM

The enemy SAM has a larger sensor radius, a full 360° weapons cone, and has the greatest probability of hitting another vehicle. However, the SAM is static, it does not move from its initial location. If the human subject runs into a SAM and can successfully evade it without being hit twice by the SAM, then the human subject knows the SAM's location from then on and can adjust his strategy accordingly. Figure 5.3.3 displays the human performance in engaging the SAMs. Out of all three enemies, the enemy SAM was the toughest enemy to kill. All human subjects sustained some damage at least once due to a SAM. In fact, out of the fifteen encounters with SAMs, only two cases, one by S3 and one by S4, resulted in damaging a SAM. Because the SAM has a larger sensor radius and full 360° weapons cone, it makes sense that all *exposure* and *weaponsLock* scores are negative. The human subject must enter into the SAM's sensor radius and thus weapons cone if he desires to engage it. Note that when S3 damaged the SAM, he only scored one hit out of three shots. Also, when S4 damaged the SAM, he only hit the SAM once for two shots taken. The question then, is as follows: does the success of S3 and S4 in damaging the SAM demonstrate positive tactical decisions in how to engage a SAM?

Verbal Reports

After S3 ran into and successfully avoided an enemy SAM, he began thinking aloud about how best to kill it, as given below. In the first line of his comments, "those places" refers to the location of enemy SAM sites. S3 believes the only possible way to kill a SAM would be to move at it at maximum speed, come to a hover right when the SAM is within the human vehicle's sensor radius, and shoot as quick as possible. As he states, this whole plan is contingent upon knowing the exact location




Figure 5-7: Human performance against enemy SAMs.

of the SAM. Because the SAM is static, the exact location can be known once the human vehicle has successfully contacted and evaded the SAM. Note that S3 does not have too much confidence in the plan or in his ability to successfully carry it out. Second, S4 comments reveal a mistake he made in engaging an enemy SAM. S4 ran out of shots, and he was unsure how it happened. Thus, S4's mistake leaves a question of whether or not his engagement tactic would have completely worked. S4 only damaged the SAM, but then ran out of shots and was killed.

SAM Engagements
Case 2, S3 Verbal Reports
<p>It's always good to have identified where those places are at I feel like the only real way to attack it would be to move in hot and fast on it Stop right when it's in range [and] fire off two shots in rapid succession So knowing where it's at is pretty key So I'll keep that in mind and if I have time in the end and I've explored everything I need to it might be helpful for our guys to have a SAM site taken care of But because it's not something I'm very good at doing I may have somebody else do it ... I'm gonna try my strategy, I think, for the SAM site, if he's still there I feel like it might be a good plan Get lined up on him Got finger on trigger and go (unintelligible comment)</p>
Case 5, S4 Verbal Reports
<p>Let's see what we got here Come on shoot, is it really not shooting, shoot Oh, zero shots remaining? What happened? What's going on here? I didn't shoot ...</p>

Surveys

In answer to the survey question on engaging enemy SAMs, S3 stated that he was, "not very likely to engage. It's too dangerous. [The SAM] has a greater range than I do ... I would have to maneuver quickly and precisely to defeat it, [and] I think it better to move on so I can accomplish the mission of clearing out the corridor. My strategy is to note the location and get out of its range as quickly as possible." S3 damaged the SAM, but the SAM killed S3. It appears that this failure in implementing the plan led S3 to believe it too dangerous to try again. He states that to kill the SAM would require the ability "to maneuver quickly and precisely," a trait, it seems, he does not believe he possesses. Therefore, S3 concludes that it is better to focus all the time on searching through the air corridor rather than risk getting damaged or destroyed by the SAM. S4 answered the survey question by writing that SAMs, "had huge targeting areas, and I tried to avoid them until the very end when I had finished everything else. I wasn't aware that you were penalized for dying and may have approached that differently now that I know that." S4 stated that he would only try to engage a SAM site if all other parts of the mission had been completed. Because the scenarios were designed to constrain the amount of time the human subjects had to perform the missions, S4 typically did not engage a

SAM. In the one engagement where S4 ran out of shots, his survey comments imply that he had accepted the risk of attempting to engage a SAM. However, now that he had discovered that to be killed by an enemy in the scenario did not merely end the chance for scoring more points but also carried a penalty, S4 leaves it open as to whether he would ever engage a SAM again. Thus, S4 does not have much confidence in his tactic.

Actions

In case 2, S3 initially detected, evaded, and formulated a plan to engage the SAM, and then he came back to execute it. In executing his plan, S3 first moved back to the SAM's location known from the previous encounter and came to a hover right outside of the SAM's sensor radius. Next, S3 turned until he felt that his weapons cone aligned with the SAM site. Then, S3 applied full throttle to move at maximum speed in a straight line to the SAM. However, two things went wrong. First, S3 stopped too short, and did not have the SAM site in his sensor radius. Second, S3 was not lined up right, and had to turn more. By the time S3 changed his heading and moved in closer to put the SAM within his sensor radius, he had spent too much time in the SAM's weapons cone. S3 could only get off one shot before the SAM killed him. Therefore, it is hard to say whether or not the strategy could be successful due to flawed execution. For S4's engagement with the SAM in case 5, S4 did not enumerate his sequence of actions, but upon observation, S4 performed the exact sequence of steps that S3 attempted. S4 stopped outside the SAM's sensor radius and aligned his weapons cone with the SAM's location. S4 only had two shots going into the engagement. Then, S4 moved at full speed and stopped with the SAM just in his weapons cone. However, S4's mistake was to shoot at the SAM right away before achieving weapons lock. Thus, S4 missed on his first shot. His second shot hit the SAM, and he went to fire a third, but ran out of ammunition. It took too long for S4 to understand his situation of having no more shots left, and the SAM killed him as he tried to make his way out of the SAM's sensor radius. Again, flawed execution prevents a complete affirmation of the strategy's usefulness. Yet, S4 did a much better job of putting the SAM just within his weapons cone as quick as possible than S3. After analyzing the data, it appears that S4 had enough time to achieve weapons lock and fire off two quick shots on the SAM before getting killed. It does seem likely that in this strategy the human vehicle will be damaged by the SAM. Furthermore, there is the limitation again of the *time2shoot* counter for both the AV and the SAM. In Table 4.3, $time2shoot = 4 \times (plotCounter)$ for the SAM and $time2shoot = 3 \times (plotCounter)$ for the AV. Thus, the AV can shoot quicker than the SAM after achieving weapons lock (see Section 4.3.2), but can the AV shoot twice before the SAM can given the distance the AV must travel inside the SAM's sensor radius before the AV can achieve line-of-sight?

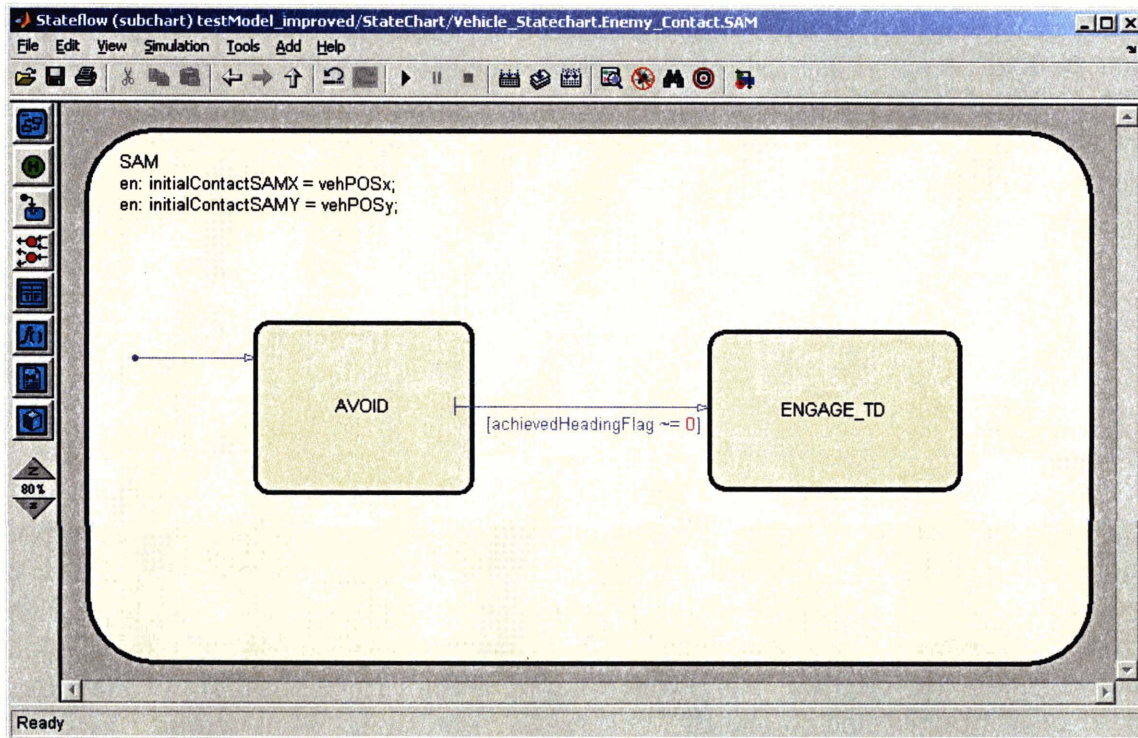


Figure 5-8: Top-level of human-inspired statechart to engage enemy SAM.

Statechart Representation

Figures 5-8, 5-9, and 5-10 depict the overall statechart diagram for the human-inspired tactic of engaging an enemy SAM site. Figure 5-8 shows the two main states in this tactic of AVOID and ENGAGE_TD. Upon first contacting the SAM, the vehicle defaults to the AVOID state, shown in Figure 5-9. In this state, the vehicle turns and moves as quick as possible out of the SAM’s sensor radius in the moveAway substate. To do this, the vehicle determines the heading from itself to the SAM, adds π to this heading, and turns to intercept that path out of the SAM’s sensor radius. After the vehicle has distanced itself from the SAM by an amount equal to the SAM’s sensor radius plus the vehicle sensor radius, it transitions to the turn2SAM substate. In this substate, the vehicle comes to a hover and turns until it directly faces the SAM. After achieving this heading, the vehicle has successfully evaded the SAM and transitions out of the AVOID state to the ENGAGE_TD state, shown in Figure 5-10. In this state, the vehicle simply moves at maximum speed at the SAM until the enemy SAM lies within the vehicle’s maximum effective range. At this point, the vehicle waits for weapons lock and then fires upon the SAM.

5.4 Results of Human-Inspired Tactics

The final step in the process of learning human-inspired tactics is to apply them over a large number of cases to test the solution’s robustness. In this context, a robust

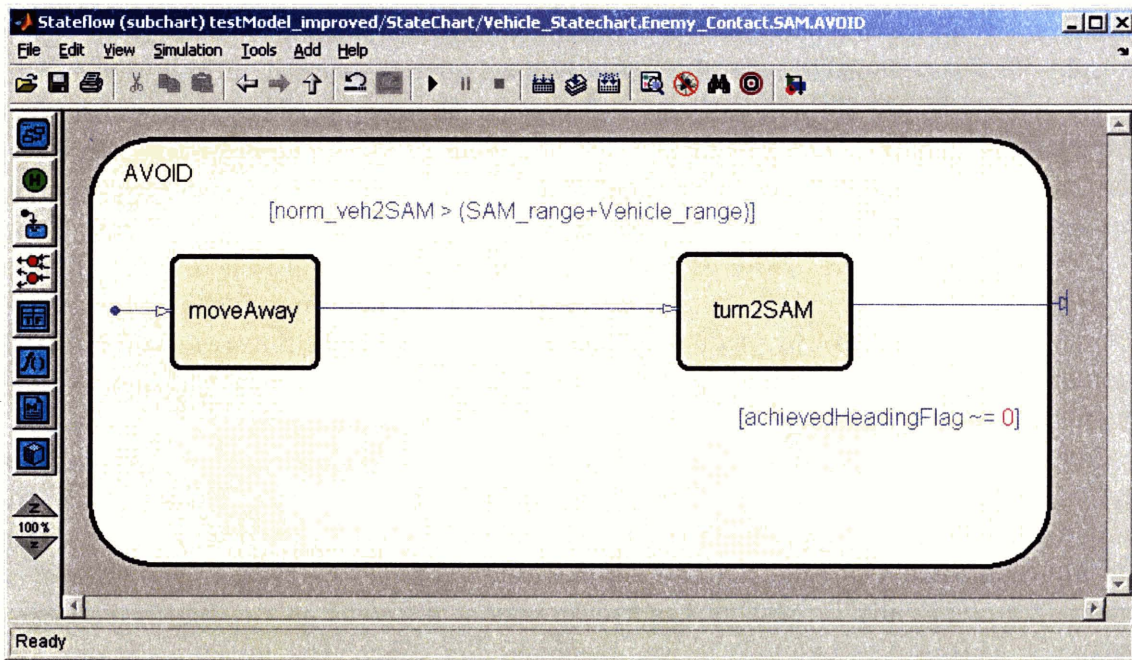


Figure 5-9: AVOID state of human-inspired statechart to engage enemy SAM.

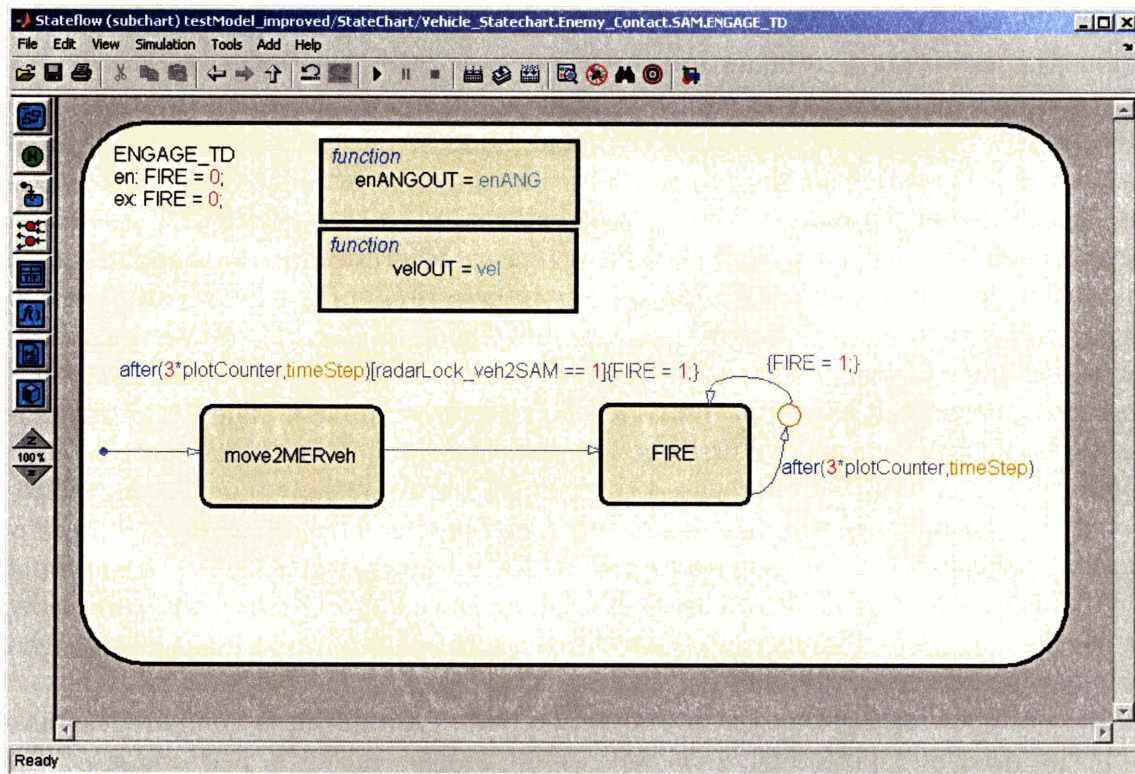


Figure 5-10: ENGAGE_TD state of human-inspired statechart to engage enemy SAM.

solution is where the actions taken by the AV remain true to the human tactic. It is difficult to judge the success of a tactic only by its performance score over a large sample. This is because the training set of human cases is so small. A highly successful tactic utilized by the human subjects over three cases may perform more poorly over a much larger sample. Therefore, there are two goals of Monte Carlo simulation, given here in the form of two questions. First, did the AV's actions consistently replicate the human strategies? Second, did the human-inspired tactic result in consistently high performance scores? Note, that if the tactic scores high but fails to mimic human strategy, the whole goal of the research to learn and apply human-inspired tactics has not been achieved. Also, by evaluating the performance of the improved AV behavior against the baseline AV behavior, we can evaluate the increase in performance given by the human-inspired tactic.

For the purposes of Monte Carlo simulation, a single terrain database and associated air corridor, critical area, and waypoint list was chosen from an earlier case in the human-in-the-loop experiments. In choosing a parameter to randomize, there were two main options. One, the parameters in Table 4.3 could be randomized. However, this would essentially be testing the tactic against different enemy platforms before testing it against the enemies it was designed for. Two, the initial locations and modes of the enemy (see Section 4.3.2) could be randomized. This would randomize how the vehicle and enemy platform initially met for each engagement and is sufficient to test the success of the tactic. Thus, the dominant randomized variable over each run was the enemy's initial location. To ensure that there was a good chance the vehicle would make contact with the enemy, the enemy's randomized initial position was constrained to being in the air corridor. Also, for testing tactics against the tank and the UAV, the dynamic vehicles, not only was the initial position of the tank or UAV varied, but the initial dynamic mode of the tank or UAV was varied. In the cases presented to the human subjects, the tank or UAV could take on the following three initial modes: follow predetermined patrol route, wait-for-contact, or pursue (see Section 4.3.2). Rather than trying to generate a pseudo-random patrol route that at least somewhat guaranteed the intersection of the enemy's and vehicle's paths, we focused on the wait-for-contact and pursue modes. Therefore, for testing tactics against the tank and UAV, a total of two hundred cases were run. The initial positions of the tank and UAV were randomized for each case. The initial mode of the tank and UAV during the first one hundred cases was the wait-for-contact mode, and for the second one hundred cases, the tank and UAV initially pursued the vehicle. Note that the first one hundred initial positions of the enemy randomly chosen for the wait-for-contact mode were *not* reused as the initial positions for the one hundred cases of the pursue mode. However, both the improved and baseline AV were tested on each set of runs. Finally, only those runs which resulted in an engagement were included in the statistical analysis.

In the following figures, the improved tactics performance is compared against both the baseline's performance as well as the human subjects' performance. The only human data shown are the specific cases drawn out in the preceding sections. This is admittedly a very small number of cases to have much statistical significance. Yet, this small number is the training set for the improved tactic, and thus the

mean, one standard deviation, minimum, and maximum is depicted for the human data. The legend identifies the subjects and number of cases per subject used in the training set. Remember that the maximum score for any engagement is a total of five (see Section 5.1.3). For the vehicle and enemy health scores, $healthVEH = health.vehicle_{remaining} - 2$ is a penalty for any damage done to the human vehicle, and $healthEnemy = 2 - health.enemy_{remaining}$ is a reward for any damage done to the enemy, where the enemy includes the tank, SAM, and UAV. For the remaining engage scoring metrics, *exposure*, *weaponsLock*, and *firingEff* all carry a maximum reward of one point.

5.4.1 Against the Tank

Figure 5-11 displays the results of applying the human-inspired tactic of engaging the enemy tank, given by the statechart in Figure 5-4, over 132 randomized cases where engagement actually occurred. Starting from the left side of the figure and moving right, the first metric *healthVEH* shows the weakness in the baseline AV behavior. Only the baseline AV is, on average, damaged by the tank. Neither the humans nor the improved AV were ever damaged by the tank. For *healthTank*, the improved AV behavior killed the tank nearly every time, which is a definite improvement over the baseline. For *exposure*, *weaponsLock*, and *firingEff*, the improved AV, on average, scored very close to the maximum of one point on each of these metrics. Furthermore, the improved AV's *exposure* and *weaponsLock* scores essentially equaled S4's stellar performance and significantly beat the baseline AV's scores in these metrics. Thus, the strength of the human-inspired tactic is utilizing the vehicle's larger sensor radius to keep the tank exposed and in weapons lock without allowing the tank the ability to do so in return. Finally, the tremendous improvement in the overall engagement score equals S4's performance and is much higher than the baseline AV. Thus, it appears that the learned tactic for the improved AV behavior of engaging an enemy tank reproduces the performance of the tactical human decision making used to derive it.

5.4.2 Against the UAV

Figure 5-12 depicts the results of applying the human-inspired tactic of engaging the enemy UAV, given by the statechart in Figure 5-6, over 136 randomized cases where engagement actually occurred. At first glance, it is seen that the improved AV behavior outperforms the human training set for every metric. This performance increase underscores one of the limitations in the experiments, namely that the human subjects were still learning how best to engage enemies (see Section 4.4.7). Furthermore, the performance increase also brings into question the validity of drawing out the strengths across different human subject performances to create a single, unified, and better performing tactic. The reason this is questionable is because we are both the simulation and experiment designers as well as the interpreters of the human strategies, actions, and data. Therefore, we could be prone to interpreting the results and encoding a tactic that we know will perform better because we know the limitations of the simulation design. This is probably more dangerous when no single human




Figure 5-11: Monte Carlo simulation to test human-inspired tactics against enemy tank.




Figure 5-12: Monte Carlo simulation to test human-inspired tactics against enemy UAV.

scenario is used to create the tactic, but multiple human scenarios are interpreted to create the best tactic. Along with this note of warning are two comments. First, note that the humans did perform better than the baseline AV. Thus, there is valuable tactical decision making to be learned by the human subjects' performance. Second, the purpose of employing several different human subjects in the experiments was to observe the strengths and weaknesses of each subject, so that we could combine the best of the tactical decision making skills into AV tactics. Interpretation is desirable but necessarily subjective. To conclude, note that the strength of applying this human-inspired tactic is found mainly in the *exposure* and *weaponsLock* metrics. By designing a tactic which kept the AV continually transitioning back to the FlyBy state (Figure 5-6), the AV was able to successfully exploit the speed logic and inferior maneuvering rate of the UAV to position itself in a superior firing location.

5.4.3 Against the SAM

Figure 5-13 depicts the results of applying the human-inspired tactic of engaging the enemy SAM, given by the statechart in Figure 5-8, over 132 randomized cases where engagement actually occurred. First of all, note that the tactic does not guarantee the vehicle will retain full health after engaging the SAM. In fact, on average, the vehicle will sustain damage, though it will never be destroyed. For *exposure* and *weaponsLock*, the improved AV's mean scores are slightly less negative than the human average scores, but it is interesting that the improved AV's mean scores do not beat the best human performance scores in these metrics. Of course, the AV will always score negative in the *exposure* and *weaponsLock* categories because the SAM's sensor radius is significantly larger than the vehicle's sensor radius. Also, remember that the baseline AV does not engage the SAM, but it always avoids the SAM. Thus, the baseline AV does not sustain any damage, does not damage the SAM, and takes no shots. Overall, the human data would have produced a total score very close to zero if they only sustained a single hit by the SAM and were not both destroyed.

Note that the previous discussion on the improved AV behavior outperforming the human training data applies here. However, much less interpretation is involved in this tactic for engaging SAMs. In fact, no interpretation is involved. The issue is that the human subjects executed this tactic in a flawed manner. S3 was misaligned and stopped short of putting the SAM within his weapons cone. S4 fired too early, missed, and only had one shot left to damage the SAM, even though it appears he had plenty of time to fire two shots after achieving weapons lock. Thus, we believe that the conclusion from Figure 5-13 is that the human subjects did formulate the right strategy to defeat the SAM. However, as mentioned in the limitations Section 4.4.6, there is no noise in the simulation, and S3's and S4's engagements with the SAM emphasize how humans in the loop tend to create noise in the system. One of the major advantages of autonomy over humans is maintaining precise control inputs over long durations. Because the improved AV is maneuvering the vehicle based on exact numbers and sensors with no noise, it can more accurately control the vehicle than the human subject who relies on visual feedback from the display to judge the accuracy of his joystick control inputs. Thus, the human-inspired tactic is extremely successful




Figure 5-13: Monte Carlo simulation to test human-inspired tactics against enemy SAM.

against the SAM, but it requires precise maneuvering, an inherent advantage of the improved AV.

5.5 Addressing Limitations by Human Cognitive Understanding

An important question to ask from the previous section is as follows: where does all the information about human cognitive processes and modeling fit into the derivation of these human-inspired tactics? The mechanisms that regulate human cognitive processes are important because they help to explain a decision-making obstacle found in these experiments. This obstacle in the human subjects' decision making prevented them from quickly learning how to effectively kill the enemy platforms. It can only be explained through human cognitive models, as discussed in Chapter 3, and it has important implications for the procedures of learning human-inspired expertise and the results already presented. As a lead-in to this discussion, we first analyze the application of one of the major elements of Chapter 2 in designing knowledge elicitation experiments.

5.5.1 Harnessing Human Strengths

From Chapter 2, it was noted that human-in-the-loop experimental scenarios should be built with a hierarchy of objectives given to the human subjects, uncertainty in the variables, and flexibility in the solution method so that the strengths of human problem-solving could be displayed. How do these three elements apply to engaging a specific enemy platform? First of all, in forming the human-inspired engagement tactics, we narrow down the hierarchy of objectives down to one, namely killing the enemy contact. The improved AV statechart logic only applies to those actions necessary to reach a specific sub-goal of killing an enemy. It does not show how the AV must decide when it should engage or avoid depending on the fulfillment of various elements of its hierarchy of objectives. Second, even though the human subject can visually see both his vehicle's sensor radius and weapons cone as well as the enemy's sensor radius, there is still uncertainty in the engagements. This uncertainty mainly resides in the tactical characteristics of the enemies. The human subject may see the sensor radius of the enemy contact, but the subject does not know the enemy's speed, maneuvering rate, weapons cone field-of-view, and time requirements to acquire line-of-sight, weapons lock, and to shoot on the subject's vehicle. These variables can only be learned and estimated through experience.

For the third element, there is unfortunately not too much flexibility in engagement actions. The engagement scenarios were discretized in order to more easily conceptualize the design of the simulated entities as well as to help in data analysis by mapping the levels of engagement to specific performance metrics (see sections 4.3.2 and 5.1.2). By discretization, we mean that there were four levels which completely define an engagement: detection, line-of-sight, weapons lock, and firing. These

levels must occur in order the given. The subject cannot acquire weapons lock without first achieving line-of-sight. The tradeoff of discretizing the engagement process was to limit *a priori* the sequence of actions that must be taken by the human subject to kill an enemy. For example, the only way a human subject can shoot and actually hit an enemy is if the enemy is in weapons lock. The enemy cannot be in weapons lock unless it is first within the vehicle's weapons cone for a specified amount of time. Therefore, the human subject cannot utilize any sort of running fire or "strafing" tactic to kill the enemy. Now, there is the question of how realistic or unrealistic is the constraint of having to track an enemy contact with radar for a specified amount of time until the weapon can achieve a lock on the enemy. Yet, regardless of the answer, that constraint appears to have violated the assumptions that all of the human subjects brought into the experiments. The following are excerpts from human subjects' verbal reports that describe the frustration of trying to apply improper engagement techniques during the experiments.

S1 Verbal Reports
<p>Whoa, bad guy, ok I am going to try and go after him, but it looks like he's mobile So I'm going to try and loop around, get him better lined up This guy's persistent, alright, I've had enough of you Turn, turn, turn I'm just trying to get him lined up so I can get a shot off Oh, not doing a very good job So I'm gonna get a little distance in between myself and him Dang't, I can't seem to do it</p> <hr/> <p>But now I'll turn around and come to ... slow ... no, aah, dang Not a very good helicopter pilot here Slow down, slow, slow, slow, slow ... (laughs)</p> <hr/> <p>Ok, tank, slowing down, turning, and then Oh shoot, turn, turn, turn, turn, turn Running out of time, running out of time, turn There we go ... (sigh)</p>
S2 Verbal Reports
<p>Ooh, I guess we can go after that bad guy Oh, he's got me in his sights No, it shot late ... aah ... it's not shooting well (Sigh), there's kind of a time delay on the gun right now</p>
S3 Verbal Reports
<p>Ok, UAV, he's tailing me, ooh Circle out and around Come back at him in this direction ... whoa Ok, behind him, maneuvering This is really tricky Aah, two shots gone</p>
S5 Verbal Reports
<p>Come back around to the top, see if I can get an engagement here And do some running fire on this guy, I'm reporting him to higher I'm going to sweep by as fast as I can on him I don't think that's working right real well so far as targeting</p>

It took multiple failed attempts of moving at full speed at the enemy and trying to quickly fire as the vehicle overtook the enemy for the human subjects to understand the necessity of waiting until weapons lock was achieved. Even after a successful engagement where the human subject stopped, waited, and fired, he still tried the doomed strafing tactic on the next enemy encountered. It could be argued that this is not really that much of a limitation because it did force the humans to find new solutions to killing the enemies, which was exactly the point. However, a quick glance over Figures 5-4, 5-6, and 5-10 shows the similarities between the necessary steps to take in engaging the enemies, namely the necessity to stop, wait, and shoot. The human subjects did have to figure out how best to arrive at the position where they could stop, wait, and shoot, and that did vary between enemy platforms. If there were truly uncertain elements and flexibility in solution methods, it would seem reasonable that some sort of probabilistic decisions would be made in these engagements where the human subjects would help determine the probabilistic thresholds. Yet, this is not present due to the discretization of the engagement sequences.

5.5.2 New Mental Model

As the human subjects continued to learn how best to engage enemies in these scenarios, they were forced to restructure their thinking based on previous experiences. Recall from Chapter 3 that many of the decision heuristics and biases humans exhibit in decision making arise because humans, in the terms of J. Reason, are “furious pattern matchers” [57]. The structure of long-term memory and the limited capacity and temporal decay of the working memory converge to create a default decision making process based on matching current situations to previous experiences. In these experiments, all of the human subjects carried some sort of video-gaming experience in long-term memory storage that was drawn upon to help make decisions in the current simulation. However, it became obvious that their experience with two-dimensional, interactive games included the action sequence of lining up the enemy with the current weapon, moving at the enemy until it is within target range, and firing upon the enemy as soon as possible. If the enemy was within target range, the enemy should be hit. These video-gaming experiences never included the constraint of maneuvering to place the enemy within target range and then having to wait a predetermined amount of time to fire. The need to restructure their thinking can be described by both the RPD [38] and GEMS [57] frameworks (see Sections 3.3.6 and 3.4.2).

Recognition-Primed Decision Model

In the RPD framework, the emphasis is on how the human expert must use a recognition decision making process to quickly and efficiently choose appropriate actions because of time pressure. This model is pictured in Figure 3-9. There are four by-products of recognition: expectancies, cues, goals, and actions. When the human subjects first encountered enemies, these four categories describe their assumption for using strafing against the enemies. First, they matched the visual cues of the vehicle’s sensor radius, weapons cone, speed, and maneuvering capabilities against the

enemy's same characteristics. Second, they recalled the previous goals of first placing the enemy inside the vehicle's target range. At that point, the goal is then to fire upon the enemy as quickly as possible and to keep moving so that the enemy does not have time to fire back. After understanding the goals, the human subjects then chose a set of actions to maneuver the vehicle towards the enemy and fire. Finally, they expected to see the enemy hit. However, this last expectancy was violated time and time again. Rather than trying to change the set of expectancies, cues, goals, and actions to an appropriate set for the given situation, they relied on feature matching and story building to try to infer the presence of the violated expectancy. Two examples of this inference can be seen in the verbal reports given above. S2 blamed the weapon and said it had a time delay on it. S1 blamed himself and stated he was not a very good helicopter pilot. In both cases, it was easier to explain why their actions, which had worked in the past, did not work now than to admit that the situation was no longer recognizable.

Generic Error Modeling System

The GEMS framework categorizes human problem-solving into three levels: skill-based, rule-based, and knowledge-based. When a problem arises, humans prefer to solve it in the rule-based level by applying stored IF-THEN rules from previous experiences. This particular model is shown in Figure 3-10. For these experiments, the human subjects' stored IF-THEN rule was a strafing tactic. Beginning at the Problem block at the top of the rule-based level, the human subjects considered the local state information of tactical capabilities of their vehicle versus the enemy and then had to answer the question of whether the pattern was familiar or not. If they first answered "yes," they then applied the strafing rule, but quickly found out that it did not solve the problem. Thus, they returned to the problem block, reconsidered the local state information, and asked the question again of whether the pattern was familiar or not. If they answered "no," they transitioned to the knowledge-based level and searched for a higher level analogy of the current situation. By finding an analogy of the current situation to previous video-gaming experiences, they quickly left the knowledge-based level and tried applying the strafing tactic again. Therefore, we can see the loop of unsuccessfully engaging the enemy the human subjects followed. Until they admitted that a higher level analogy did not hold, formed a mental model of the problem space, and abstracted relations between structure and function, they would not learn how to appropriately engage the enemy.

5.5.3 Implications

These two decision making frameworks from Chapter 3 describe the reasons behind the difficulties that human subjects experienced in trying to learn how to engage enemy contacts. By understanding the cognitive processes underlying the decisions made, we can now list some important implications. First, because human subjects became frustrated over the course of multiple failed attempts of strafing the enemy, they lost motivation to keep trying to learn the right ways to engage. This hinders

the researchers who are looking for the best tactics. Now that the human subjects have failed so many times, they do not want to keep searching for the best tactics. Second, without sufficient time and training to become experts in the simulation, only a small number of solutions will be found. Such a small training set is one criticism of this research, especially for the engagement tactic against UAVs. In that tactic, only three cases could be analyzed, and the decisions made had to be filtered and collected in an interpretation process that was possibly biased by the knowledge of the experimental designer. Not only then was there a small number of partial solutions to train from, but recall the high variance in engage scores in Figures 5-1 and 5-2. A large set of complete and consistent solutions in how to engage enemies would be ideal and would reflect complete domain expertise. However, in terms of the RPD framework, the human subjects employed in these experiments did not have enough time to abandon their recognition process and reformulate their set of expectancies, cues, goals, and actions to context-appropriate ones. In terms of the GEMS framework, they did not have the time form correct mental models, abstract relations, and learn appropriate actions which could then be stored as new IF-THEN rules. If there is one thing to learn from Chapter 3, it is that to overcome the human bias of pattern-matching takes time and motivation.

5.5.4 Engagement Tactics Conclusions

The solution method of engaging enemies was rigid and violated the previous experiences of the human subjects. Coupled with the lack of time given to experimentation, the human strength of exploring the problem space to find creative, new solutions was not completely harnessed. This accounts for the large variance in engage scores, the small number of cases to analyze, learn, and train from, and the interpretation process required to integrate different decisions into a unified tactic. The success of the engagement tactic against enemy UAVs is therefore presented with this criticism. However, for the engagement tactics against enemy tanks and SAMs, the statechart logic was derived explicitly from complete sets of human solutions. The success of the engagement tactic against enemy SAMs emphasizes the limitation of the simulation which only deals with truth data. The success of the engagement tactic against enemy tanks is truly the strongest proponent of the advantage of learning and applying human expertise.

5.6 Search Problem and Planning

Human subjects had to be prepared to react against enemy contacts while searching through the air corridor and critical area. Over the two rounds of experiments, the subjects learned and applied the right actions to evade and/or engage enemies. This is tactical decision making. The previous sections described the process of evaluating, learning, and applying the best human-inspired, engagement tactics which resulted in a significant increase in AV performance. The previous sections also evaluated the possible criticisms of the methodology and results in terms of human cognitive

processes. However, these engagement tactics are by no means the only ones that can be learned by the experiments. A much more difficult problem exists of how to search through the given terrain and intelligently make tradeoffs between limited time, enemy contacts, and critical terrain.

5.6.1 Plans

At the beginning of the scenarios in the second round of experiments, the human subject was asked to determine and verbalize a plan to accomplish the mission objectives of searching for enemy contacts. The variables present that influenced the plan formulation were as follows. First, the human subject could visually inspect the location of the air corridor and critical area with respect to the terrain. Second, the human subject was given an intelligence report that categorized the probability of enemy contacts in both the air corridor and critical area. Third, the time constraint of four or five minutes forced the human to divide his time between the air corridor and critical area. The human subject then combined these variables with past experiences to construct a plan for the scenario. The following are some examples of formulated plans from the second round of experiments along with the associated map view given to the human subject. Note that the caption of each figure displays the intelligence report for the particular case.

Plan for Case 5, S1 - Figure 5-14

“Very good chance of enemy contact in both areas. I just want to go as efficiently as possible to the critical area, and try to minimize getting killed in the process. Most efficient in terms of coverage, would be to go around the large hill to the west and north around it. But that confines where you can go, plus up north they got a couple good places to hide. So instead of doing that, I’m going to go the southern route and go south and east around it. I’ll probably even duck outside the corridor and go and just center the gap between those two hills to really give myself the best chance of survival. Then, I’ll probably do a north/south weave in critical area. And then . . . hmm . . . shoot, my hesitation there is that I want to do a north/south weave in the critical area from west to east but that kind of puts me where I could catch that last little bit of the corridor. But then I would have to all the way backtrack and go to the unobserved section. So what I might do is go south and east into the critical area and then do a east/west search pattern and forget about the little stretch of corridor off to the right there. And then just hit the south around the gap where the bad guys could be hiding. And then just go back through the critical area go around north and then west around the rectangular hill. I don’t anticipate to live that long.”

There were two distinct characteristics seen here that were a part of all of S1’s plans. First, S1 went into the most detail in terms of search. Not only did he plan out the order of how he would search the air corridor versus the critical area, he also specified in extreme detail the directions of searching in terms of north, south, east,




Figure 5-14: Very good chance of enemy contact in both the air corridor and the critical area.




Figure 5-15: Possible chance of enemy contact in the air corridor and a very good chance in the critical area.

and west. Second, S1 was extremely wary of terrain features that might either offer a potential hiding spot for enemy contacts or constrain his maneuverability. Specific to the above plan, S1 reasoned that he would focus first on the critical area and then on the air corridor because they both had equal chances of enemy contacts. In fact, to move “as efficiently as possible to the critical area” while minimizing the chances of “getting killed,” S1 decided to move outside of the air corridor on his way to the critical area. This would hurt him in terms of accruing points for the percentage of air corridor seen, but he reasoned it was worth it if he stayed alive long enough to reach the critical area. Note that S1 did not have much confidence in his abilities, and he finished his plan by stating, “I don’t anticipate to live that long.” Out of all the human subjects, S1 struggled the most in learning how best to engage enemies (see Section 4.4.7).

Plan for Case 1, S2 - Figure 5-15

“Since the critical area is next to terrain and since there’s a very good chance of running into bad guys in the critical area, I need to be careful in approaching it and also reverse the order [from previous runs where there was higher chance of enemy contacts in the air corridor than in the critical area]. So go through as much of the air corridor first before getting to the critical area, so I at least cover that. Just be cautious and look for a way to escape if there’s one of the really dangerous bad guys in there.”

In this case, the flow of terrain was as follows: it began with the air corridor, the air corridor led to the critical area, and the air corridor also extended past the critical area. Thus, searching along the air corridor in the beginning would naturally lead to the critical area, and then the air corridor continued on past the critical area. Although the critical area was a more important piece of terrain, S2 decided that the greater chance of enemy contacts in the critical area was too much of a risk to cover first. S2 would rather take the sure points of searching through the air corridor where there was only a possible chance of enemy contacts. Also, S2 noted the terrain obstacle adjacent to the critical area, which forced him to be even more cautious and aware of an escape route in case of enemy contact. However, out of all the subjects, S2 used full throttle during the majority of every scenario. Thus, S2 did not equate being “careful” or “cautious” with reducing speed. It was more of a mental reminder to expect enemy contacts.

Plan for Case 6, S3 - Figure 5-16

“I’m still satisfied with the way things are going, in terms of going straight through [the air corridor - the most direct route from beginning to end], spending time in the critical area, and taking the long road home. There’s a slim chance I’ll see enemy contacts anywhere, so I’ll go on my merry way exploring, not even going to have any kind of adjustment. The last time I wanted to fly faster through the narrow area in case there was a SAM site. I probably won’t worry about that this time.”

S3 developed a unique strategy that governed his behavior over all cases. S3 reasoned that the first priority of every scenario was to search through and clear the shortest path from the beginning of the air corridor to the end of it. Thus, he always chose to search through the air corridor first and then to turn towards the critical area. Finally, after searching through the critical area, he would attempt to cover any other missed portions of the air corridor that did not lie along the shortest path. Furthermore, the air corridor was always wider than the vehicle’s sensor radius. To cover the entire air corridor required some sort of weaving or looping through it. Yet, in an effort to clear the shortest path the quickest, S3 chose to not weave but stay one side of the air corridor. The shortest path through the air corridor then did not mean the entire width of the air corridor had been searched, only that portion which fell under his sensor radius. In S3’s above plan, the slim chance of enemy contacts in both the air corridor and critical area allowed him to relax and simply plan to “go




Figure 5-16: Slim chance of enemy contact in both the air corridor and the critical area.




Figure 5-17: Possible chance of enemy contact in the air corridor and a very good chance in the critical area.

on my merry way exploring” without any expected need to make adjustments. Note that S3 decided to be more efficient with his speed (full speed accrued penalty points) due to the lack of expected enemy contacts but not his searching plan. His shortest path mentality was not very efficient, which he knew, but he stuck with it.

Plan for Case 2, S4 - Figure 5-17

“Very good chance inside the critical area, and the critical area happens to be at end of corridor anyways. So I was thinking to try and get the majority of the corridor, and then, when there’s a minute left or so, head for the critical area and see what happens . . . maybe a little over a minute. In the event that I die, I would have at least gotten much of the corridor covered.”

S4’s focus was always on survival. Wherever there was the least chance of enemy contacts, he would move there first. If both the air corridor and critical area had equal chances of enemy contacts, he would follow the natural progression from air




Figure 5-18: Possible chance of enemy contact in the air corridor and a slim chance in the critical area.

corridor to critical area, without really focusing on either. In this plan for case 2, S4 chose to leave the critical area for last for two reasons. One, it had a very good chance of enemies and two, it lay near the end of the air corridor. Also, S4 did not give the critical area an appropriate amount of weighting. For a four minute scenario, he decided that “maybe a little over a minute” was enough time to spend searching the air corridor. He decided this even though he knew that one scoring element was the amount of time spent searching through the critical area versus the air corridor. If more time was spent in the air corridor, S4 would lose points. However, this at least gave him more control over his survival, and if he died in the critical area, he would have at least seen “much of the corridor.”

Plan for Case 4, S5 - Figure 5-18

“I have the circular route around the corridor where there’s a possible enemy contact. I’ll initially want to move towards the critical area. Because I don’t have a lot of time and there’s a lot of area to cover, I’ll choose to go up the corridor this time versus directly to the critical area. Being cognizant of inside the critical area, there’s, in the northeast side, a large mass of terrain which minimizes my maneuverability. So as I scan that area, I’m going to want to scan that area from the south moving north this time, primarily to ensure that I have a turn point where I’m not boxed in by the enemy or I get surprised.”

Out of all of the human subjects, S5 gave the most emphasis to deliberately searching the critical area. He liked to use the expression of “gingerly searching” through a critical area which had a very good chance of enemy contacts. In the above plan for case 4, there was a possible chance of enemy contacts in the air corridor and only a slim chance in the critical area. Thus, he chose to follow the longer path in the air corridor versus moving “directly to the critical area.” S5 was also very conscious of maneuvering limitations given by terrain obstacles. However, where other human subjects simply finished their thoughts by stating they needed to be more careful due to terrain obstacles, S5 spoke of creating exit lanes and buffer zones through his searching patterns “to ensure that I have a turn point where I’m not boxed in by the enemy or I get surprised.”

5.6.2 Partial Plans and Intentions

All of these plans are partial. The human subject cannot specify a plan so detailed and complete that it becomes irreversible. The subject does not know, for example, when or if he might have enemy contact. However, the chance of enemy contact may affect the sequence of searching first through the air corridor or through the critical area. The plan must be partial, and as the scenario progresses, the subject fills in the details of the plan as necessary. The only way to fill in a partial plan with coherence and consistency is to have a strong, governing goal that prioritizes the decisions made. One of the human strengths was to maintain a broad focus in problem-solving, and this is exactly the purpose of Bratman’s BDI framework [7] (see Section 3.4.1). In the BDI framework, the intentions of the subject provide for coherent and consistent execution of partial plans. The human subject can keep a broad focus in problem-solving because of the presence of intentions. If S3 intended to find the shortest path from the beginning of the air corridor to the end and if an enemy SAM lay in the way, S3’s intention governs his decision of whether to avoid or engage the SAM. Intentions are different than desires [7]. If S3 only desired to secure the shortest path, the presence of the enemy SAM could easily lead him to abandon that path because he has a greater desire to survive. Yet, if S3 intends to secure the shortest path, he must decide how to deal with this enemy obstacle. That intention governs the behavior, and these sorts of tradeoffs occurred all the time during the human-in-the-loop experiments. The BDI framework, then, is appropriate to capture

the plans given above. The plans can be formulated through the combination of the subjects' desires, beliefs, and intentions, with the intentions serving as the foundation to keep the partial plan from falling apart due to unforeseen circumstances.

5.6.3 Intention Examples

All five subjects possessed a desire to survive as long as possible. This desire, combined with their beliefs, formed an intention of how to search through the scenario. For example, when the critical area had a very good chance of enemy contacts, S2 and S4 intended to not search the critical area until the very end, despite its greater importance, because they believed being killed early on in the scenario was worse. S1 intended to alter his search patterns due to every concave part of terrain because he believed it was a possible hiding place for enemy contacts. S3 intended to find the shortest path from beginning to end of the corridor because he believed that initial set of actions would fulfill the fundamental requirement to achieve mission success. S5 intended to search very deliberately and slowly in the critical area because he believed it was more valuable than the air corridor. Also, S5 intended to search very carefully along terrain obstacles because he believed that would create a "buffer zone" so that he would not be pinned against the terrain obstacle by the enemy. Finally, S5 intended to always keep "exit lanes" behind him because he believed it would be much better, upon enemy detection, to turn towards already-cleared terrain versus new, unidentified terrain where more enemies could be located.

5.6.4 Intent Function

We propose that the intentions of the human subjects can be separated into three main elements that characterize their searching behavior. These three elements are the search sequence, amount of coverage, and time devoted to searching through the air corridor and critical area. The human subject must decide the sequence of searching the air corridor to searching the critical area, how much coverage is appropriate in each, and how much time should be spent in each of the terrain areas. (Note that the air corridor and critical area could be broken down into smaller search segments depending on, for example, the presence of terrain obstacles. Yet, the smaller the segments the more detailed the plan, because each search segment should still contain the three elements of sequence, coverage, and time. Plans with too much detail are impractical.) These dependent variables represent the subject's intentions. Note that the verbalized plans of the human subjects only explicitly describe one of these three dependent variables, namely the sequence of search. However, by observing the humans' actions throughout the scenario in terms of coverage and time spent searching, the intentions of the human subjects can be derived in combination with the intended sequence of searching. The verbal reports and surveys will also be very useful in deriving intent.

The search characteristics of intended sequence, coverage, and time are a function of the human subjects' desires, beliefs, terrain layout, probability of enemy contact, and time constraint. Hopefully, the human subject has both the appropriate desires

of achieving mission success, which includes survival, and the appropriate beliefs or understanding of the environmental state and the vehicle's tactical performance characteristics relative to any enemy contacts. The intention could then be captured as the following function, where $\overrightarrow{Intention} = \begin{matrix} \text{sequence} \\ \text{coverage} \\ \text{time} \end{matrix}$

$$\overrightarrow{Intention} = f(\text{terrain layout}, P(\text{enemy contact}), \text{total time}) \quad (5.1)$$

Three important considerations must be noted in connection with Equation (5.1). First, the human subject does not intend to cover 67.5% of the air corridor in one minutes and 43.2 seconds. Rather, the human subject may intend to cover "most" of the air corridor using "about half" the allotted time. (Note that for the critical area, overlapping search would be desired.) The goal of human experimentation should not be the identification of a complex, multivariable, nonlinear intention function, because that is not how humans behave. Instead, the intention function should be made up of discrete levels and thresholds. For example, at some combination of independent variables, the coverage scalar in the intention vector switches from covering "some" of the air corridor to "most" of the air corridor. A fuzzy logic controller, then, appears to be a natural candidate for this intention function. Second, an execution function should exist in series with the intention function. After the fuzzy, intended variables of sequence, coverage, and time are calculated based on the terrain layout, $P(\text{enemy contact})$, and total time, they are then combined with the size of the air corridor, critical area, and vehicle sensor radius and finally passed to an execution function. This execution function selects the appropriate maneuver patterns and calculates the required maneuvering speed. The output of this execution function could be a matrix of waypoints and speeds, and then this matrix is handed over to a trajectory-following controller. These maneuver patterns should also be part of the set of tactics learned by the humans. For example, if the human subject intended to cover most of the critical area and if the air corridor width was only slightly larger than the vehicle's sensor radius, the vehicle could move straight down the center of the air corridor. As the width of the air corridor increases relative to the vehicle's sensor radius, there must be some threshold which changes the maneuver pattern from moving along the corridor centerline to weaving back and forth at some angle relative to the centerline. This threshold can be learned by human experimentation. Another central role for the execution function would be to determine at what point during the simulation are the fulfillment of the original intentions no longer possible. For example, if the human subject intends to cover all of the critical area in half the time, but spends most of that allotted time engaging enemies present in the critical area, he must now decide if it is possible to finish covering all of the critical area in the little time amount of intended time remaining and how to proceed. The third and final consideration is that if the execution function determines the original output of the intention function cannot be feasibly implemented, there should exist one more block in this system. This block could be a saturator, where the vehicle still attempts to follow the original intentions but knows that it is not possible. This block could

also be part of a feedback loop which recasts the inputs to the intention function to the current state of the simulation so that the intention function can recalculate the appropriate sequence, coverage, and time for the rest of the simulation. Figure 6-2, given in the next chapter, displays a block diagram representation of the intention and execution functions.

5.6.5 Rational Intent

The BDI framework captures how intentions, combined with beliefs and desires, govern the steps taken to fill in the missing gaps of partial plans. BDI also addresses how intentions provide the consistency and coherence necessary for action coordination. However, BDI does not address whether the intention initially formed was rational from a normative viewpoint. In these experiments, one way to evaluate the subjects' intentions is through the resulting search scores and overall performance. Yet, a performance score helps point to rational action, but does not necessarily describe it. We propose that in evaluating intentions, it is important to keep the decision heuristics and biases in mind. Several of these heuristics and biases were seen in the experiments and are presented as follows.

Framing Effect

In the framing effect from Section 3.3.4, the human subject decides to be risk-taking or risk-averse depending on the subject's frame of reference. In these experiments, the subjects were given a positive frame because the goal was to accrue or gain as many points as possible. The more area seen, the more time spent in the critically weighted terrain, and the more enemies destroyed all increased the total positive return for the human subject. Thus, the subjects were overwhelmingly risk-averse. Consider, for example, S2's plan for case 1, given in Section 5.6.1. There was a very good chance of enemy contacts in the critical area and a possible chance in the air corridor. S2 stated the he would "go through as much of the air corridor first before getting to the critical area, so I *at least* cover that." Now consider S2's plan for case 3, in which the chance of enemy contacts in the critical area and air corridor were reversed from case 1. His plan was stated as follows:

I think I'm gonna go clockwise around the circle, so that I hit the critical area first. Because if there's a very good chance of encountering bad guys in the air corridor, then I'd like to encounter them later in the flying time so I don't die in the beginning and *not accrue any points*.

Nowhere in his reasoning of what search sequence was appropriate, critical area before or after air corridor, did S2 reason, at least verbally, about the importance of weighted terrain. It did not matter to S2 whether the critical area was a more important piece of terrain to search. If it involved greater risk, he would search it last.

Anchor and Adjust

In the anchoring heuristic from Section 3.3.4, the human subject formulates an initial solution to an ambiguous or ill-defined problem. The human subject does not abandon the initial solution even when new information is present, but instead the subject merely adjusts his or her answer from there. There were two distinct examples of the anchoring heuristic in these experiments from S1 and S3. As mentioned earlier, S1 treated every concave piece of terrain as a possible enemy hiding spot. For example, during case 2, as he approached a concave piece of terrain, S1 stated, “Notice I went a little bit up because I was closing in on that mouth. I’m trying to make so I can . . . Yeah, see I turned there to see if I could engage if anybody comes out.” Similarly, in case 1, S1 said, “I’ll just have a narrow weave. This is pretty tight anyways [with] not too many places for bad guys to hide . . . I’ll turn here, be pointed if any bad guys come out.” The amazing fact is that S1 never encountered an enemy proceeding from one of these concave “hiding spots.” Yet, he stuck with this mental image, and still adjusted his searching patterns to accommodate these hiding spots *even when* there was a slim chance of enemy contact for the entire scenario.

S2 displayed another anchoring bias by developing his shortest path solution to every scenario. He valued searching through and clearing a continuous path through the air corridor before turning to search the critical area. This was an inefficient use of time, and on average, S2 missed 12% of the critical area, which was the worst among the five human subjects. Though the critical area was more important terrain and S2 knew he was rushed through searching the critical area at the end of the scenario, he never abandoned his shortest path solution.

Failure to Account for Probability

S1 formed a mental anchor of always changing his search pattern based on ominous pieces of terrain. This anchor kept S1 from making more rational decisions based on probability (see Sections 3.3.4 and 3.3.3). In terms of Bayes Theorem, given by Equation (3.2), A represents the enemy contact, and B represents the terrain hiding spot. Thus, $P(A|B)$ is the probability of contacting an enemy given the presence of a terrain hiding spot. Over all the cases, no enemy ever blind-sided by S1 by attacking him from a hiding spot. Thus, $P(B|A)$, the *a priori* knowledge, should have continually decreased after every case. By Equation (3.2), $P(A|B)$ should have decreased accordingly. Furthermore, S1 maintained his anchor even when there was a slim chance of enemy contacts in the entire scenario. Thus, S1 failed to account for $P(A)$, the base rate of occurrence of enemy contact, which would have also decreased $P(A|B)$. The anchoring heuristic combined with a failure to account for probability resulted in poor search performance for S1. On average, he saw only 60% of the air corridor, which was the worst score, and he missed 11% of the critical area, which was the second worst score. S1 sacrificed efficient search for a mental anchor of being prepared for enemy ambushes.

5.6.6 Search Problem Conclusions

It is a difficult problem to find the right searching techniques based on hard-to-quantify terrain characteristics, likelihood of enemy contacts, and time constraints. Because human subjects can only form partial plans of how to proceed during the scenario, they form conduct-controlling intentions that help keep consistency and coherence to their actions. If the human subject encounters a pop-up threat and must deviate from the search plan to engage, the intentions govern how the human subject now chooses to proceed after the reactive tactical situation has been resolved.

All of the examples of intentions that the human subjects' possessed, described in Section 5.6.3, can be correlated with their resulting search scores from Figure 5-2 to find the best intentions. Next, the rationality of those intentions can be determined through applying cognitive decision making theory. Note the extreme importance of planning verbalization and think aloud reports. These intentions would then be broken down into the categories of sequence, coverage, and time, and the intention function formed. However, the experimentation for this research does not contain enough cases and data to identify both the proper threshold values for the intention function, which can be thought of as the inference process in determining the fuzzy logic rules, as well as a set of maneuver patterns for the execution function which implements the intentions. Further experimentation through simulation runs and surveys would need to be particularly focused on determining these threshold values and executable maneuvers. All in all, given the resources to continue the experimentation, proper search tactics could be learned and applied to the AV through the formation of an intent and execution function.

Chapter 6

Conclusions and Tactical Framework

As stated in Chapter 1, the purpose of this thesis is to address the following two related questions:

1. How can the tactical decision making capabilities of human experts be learned and transferred over to an autonomous vehicle?
2. Can a human expert learn how to exploit vehicle-specific dynamics in tactical scenarios to achieve high levels of performance for goal-oriented missions?

In answer to both questions, we refer to Figure 5-11 and its results as the main achievements of this thesis. S4 was able to learn and consistently apply the right reactive tactics to engage the enemy tank. The answer to question two, then, is “yes,” a human expert can learn how to exploit vehicle-specific dynamics and achieve high scores. In order to learn and transfer this tactical knowledge over the AV, we used S4’s think aloud reports and answers to survey questions to identify his strategies in engaging the enemy tank. Note from Figure 5-11 the strong and consistent performance of S4 over the baseline AV. This disparity emphasizes the need to learn from human experts. Furthermore, note that the incorporation of the human-inspired tactics brought the performance of the improved AV to a level equal with S4. The improved AV did not perform poorly and did not greatly surpass the human performance. Therefore, we present two conclusions from the reactive engagement tactics against the enemy tank as the main accomplishments of the research. One, the actions taken by the improved AV were consistent with the human strategies. Two, the translation from the visually stimulated human domain to the autonomous sensory input domain did not distort the successful performance of the learned tactic.

At the conclusion of this research, we propose the following framework, given by Figure 6-1, to further answer these two questions.

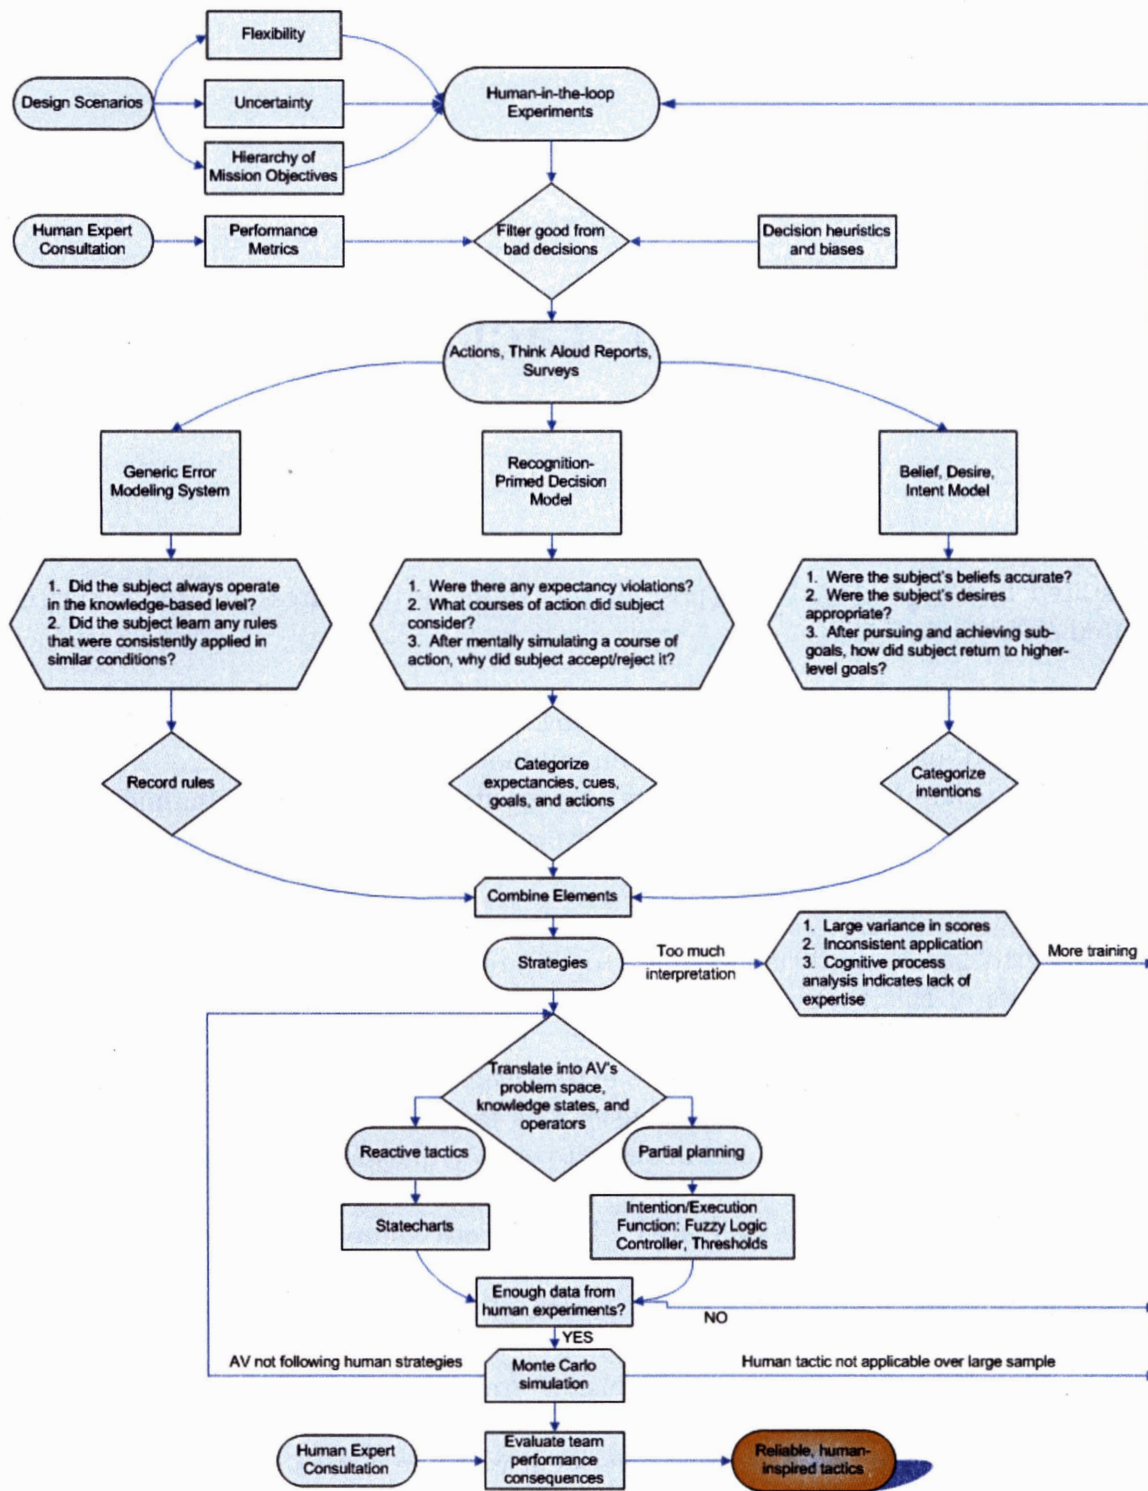


Figure 6-1: Framework for learning and applying human-inspired tactics.

6.1 Human-in-the-Loop Experimental Design

Starting at the top left of the figure are the two most important experimental design and setup steps. First, the design of the scenarios must include uncertainty, a hierarchy of objectives, and flexibility in solution methods. These design elements allow for and encourage the human subject to exhibit those decision making strengths unique to humans that automation does not possess, as discussed in Chapter 2 [9, 11]. Yet, by forcing all engagements to proceed along the same discrete levels in this research, as seen in Chapter 5, the solutions for reactive engagement tactics were not that flexible. Each enemy required slightly different maneuvering on part of the vehicle to place the enemy within the vehicle's weapons cone, but the final few steps of stop, wait, and fire were identical for all three. Second, the design of the performance metrics is also very critical because it is the first step in differentiating good actions and decisions from bad ones. It takes time to learn human strategies, and thus, separating out the right cases early on significantly reduces the time involved. Thus, a human expert should be consulted to derive the right performance metrics. This step did not occur in this research, and as discussed, the engage score metrics correlated exactly with the discretized engage levels to ease the interpretation process of metrics and actions to strategies.

To summarize, the two most important design steps are, first, designing scenarios to encourage human ingenuity. To do so requires dispensing with simulations that are completely quantitative. The second most important experimental design element is defining performance metrics that effectively filter the good from bad decisions. This should happen by consulting human experts who can help focus the performance metrics to bring out the exact nature of expertise the designers wish to learn. Note that there are many other factors that influence human-in-the-loop experimental design. Some of these are the training effect, the need for careful wording and presentation of the experiments to the human subjects due to decision heuristics and biases, and the host of obstacles encountered with designing real-time simulation processes, as discussed in Chapter 4. Yet, Figure 6-1 highlights the emphasis that should be given to scenario and performance metrics design.

6.2 Identifying Expert Performance

After conducting a number of human-in-the-loop experiments, the first step in learning tactical expertise is filtering good from bad decisions. Remember that there are three levels in understanding human tactics - actions, strategies, and processes - and the goal is to learn human tactical *strategies*. The filter is the first step in dividing out and understanding both the actions and the processes in order to ultimately learn the strategies. Thus, this filter includes analysis of the actions by performance metrics and analysis of the processes by decision heuristics and biases. The importance of performance metrics is discussed above. Statistical analysis of the resulting scores helps narrow in on human subject strategies that were consistently and successfully applied. Statistical analysis also helps identify the outliers of successful performance

whose presence might indicate sheer luck or a very good tactic which did not have the chance to be applied over a large number of cases. These outliers should be kept in mind as candidate tactics, and Monte Carlo simulation will help in differentiating between luck or lack of application as the most likely explanation of the outlier's performance. An example from this research is the SAM engagement tactic from Chapter 5.

Also important to the filter process is looking for decision heuristics and biases in the think aloud reports and surveys [15, 36, 87]. Decision heuristics and biases can be fairly rapidly identified, such as the anchoring heuristic displayed by the human subjects, S1 and S3. Furthermore, identifying decision heuristics and biases helps reduce the time required to learn human strategies significantly. For example, say the interpreter of the results is trying to derive a complex, three-dimensional intent function where all three independent variables are altered, in turn, to see how the human subject responds. Moreover, say the human subject is displaying a salience bias, whereby only one of the three variables actually affects the human response because it is always the most attention-getting. The interpreter could become quickly frustrated in trying to derive this function because no matter how two out of the three independent variables are modified, the human subject's response barely changes. Yet, if the interpreter first looked for and found the salience bias, the function could be derived very quickly without frustration. The important point from this example is that decision heuristics and biases are simple shortcuts to complex, multi-dimensional strategies. In linear algebra terms, the interpreter of results was looking for a function that spanned all three dimensions, but due to the salience bias, the rank of A in $Ax = b$ is only one.

6.3 Training

Not shown in Figure 6-1 is the training of human subjects. Initially, there must be some time for the human subject to become familiarized with the simulation environment. These training sessions should be performed exactly like those cases that will be scored, but they are not included in the results. Upon looking back on the results of this research, we propose, with caution, another step in the training process, namely feedback. The human subject should be taken through a few training scenarios, and then those scenarios should be scored and given to the human subject to review. This evaluation opportunity will quicken the learning process of the human subject. For this research, such a training/reviewing process would have helped human subjects to correlate their own strategies with the resulting performance and adjust as necessary. However, there is a major note of caution in order. To break down the human subject's scores across the performance metrics and present it to the subject will narrow all subsequent behavior by the subject towards maximizing those scores. If there is full faith in the performance metrics - if the domain of desired tactical expertise covers only that which maximizes the performance scores - this training/reviewing is extremely useful. In fact, the human subject should be given, time permitting, his or her score after every case to continually learn and adjust strategies to attain maximum performance. However, if the performance metrics only contribute to learning

tactical expertise, allowing the human subject to review his or her own score in any sort of detail will limit the range of future creative solutions. This negates the design of flexibility into the scenarios. In fact, it seems reasonable that one of the goals in the first few rounds of human experiments should be to have to iteratively refine the performance metrics because the human subjects have discovered and displayed ingenious solution methods.

6.4 Analysis through Cognitive Frameworks

Once a subset of cases that appear promising to learn from have been separated out of the experiment results, the corresponding actions, think aloud reports, and surveys are then analyzed by three cognitive frameworks.

6.4.1 Generic Error Modeling System

The first is the Generic Error Modeling System (GEMS), which was designed explicitly for the purpose of capturing the human tendency to pattern-match stored experiences to current situations in decision making [57]. The GEMS framework, given by Figure 3-10, helps determine what performance level of problem solving the human subject is in, skill-based or knowledge-based. In the skill-based level, the human subject applies previous experiences to current problems by applying stored IF-THEN rules. In the knowledge-based level, the human subject has to create a new mental model of the problem space, try new actions, observe consequences, and learn new successful IF-THEN rules. However, recall from Figure 3-10 that even when the human subject moves down to the knowledge-based level, the first action of the subject is still to search for a higher-level analogy that will allow him or her to jump back up to the skill-based level. GEMS, then, can act somewhat as a gauge of expertise. If expert decision making can be distinguished, in part, by the correct application of stored rules, GEMS provides a framework for evaluating both the level of expertise and the identification of the rules. This phenomenon was seen most clearly in S4's strategies in engaging the enemy tank. On the other hand, if the subject always solves problems in the knowledge-based level, this would represent a lack of expertise. Therefore, we highlight two important questions to ask in evaluating the actions, think aloud reports, and surveys with the GEMS framework:

1. Did the subject always operate in the knowledge-based level?
2. Did the subject learn any rules that were consistently applied in similar situations?

6.4.2 Recognition-Primed Decision Model

The second cognitive framework to analyze the results is the Recognition-Primed Decision (RPD) model, displayed in Figure 3-9 [38]. RPD was derived from decisions

made in real-world, time pressure, high risk situations. Therefore, RPD is very applicable to tactical decision making. The main focus in RPD is the categorization of the human subject's recognition of a situation into expectancies, cues, goals, and actions. This view can be seen as an expansion of the IF-THEN rules from GEMS. The simple input/output relation of GEMS would be: if these cues are present, then these actions are taken. RPD states that the antecedent contains both the present cues and the expectancies. Maintaining a set of expectancies of how the situation should unfold is crucial to the decision maker. The main way a decision maker comes to question his or her situation awareness is through an expectancy violation. This is the upper left feedback loop in Figure 3-9, and, in terms of stored rules, the decision maker questions the antecedent - if (???) - and seeks more information. Thus, expectancy violations in RPD appear to be the triggers for the human subject's descent from the skill-based to knowledge-based level in GEMS. Furthermore, RPD declares that the consequent contains both actions and goals. These goals are extremely critical to learning the human's strategies because strategies are typically cast as goal-oriented. As S4 was engaging the tank, he thought aloud that he should stay outside the enemy's sensor radius. This was S4's strategy cast in terms of the goal of maintaining standoff distance from the enemy tank. When the decision maker uses mental simulation to evaluate the consequences of courses of action, the criteria to accept or reject the actions come from the decision maker's goals. Thus, learning why the human subject accepted or rejected a course of action reveals the subject's goals. We propose three important questions to summarize the important features of RPD:

1. Were there any expectancy violations?
2. What courses of action did the subject consider?
3. After mentally simulating a course of action, why did subject accept/reject it?

6.4.3 Belief, Desire, Intent Paradigm

The Belief, Desire, Intent (BDI) paradigm (see Section 3.4.1) is the third cognitive framework to help evaluate the experimental results [7]. BDI describes the human behavior to formulate plans for the future because of the need for coordination in the presence of uncertainty. These plans are partial, and thus, it is the intentions of the subject which guide future decisions and the steps necessary to successfully execute the partial plans. As a comparison, the intentions of the subject in the BDI paradigm correlate with the goals of the subject in the RPD framework. Just as the goals represented the accept/reject criteria of actions, the intentions also act as a filter, admitting those steps which are coherent and consistent with the intentions. Furthermore, just like goals, intentions are directly related to strategies. Yet, if the subject does not have the right beliefs about the environment (the correct mental model of the problem space) or does not have appropriate desires to fulfill the mission objectives, the intentions do not mean much. Finally, if the necessary plan for each scenario contains fundamental elements, such as sequence, coverage, and time of

search for this research's experiments, these elements can be used to categorize the intentions. We emphasize three questions to apply the BDI framework.

1. Were the subject's beliefs accurate?
2. Were the subject's desires appropriate?
3. After pursuing and achieving sub-goals, how did subject return to higher-level goals?

6.5 Combining into Strategies

After analyzing the actions, think aloud reports, and surveys with the three cognitive frameworks, the next step is to then combine the stored rules; recognition by-products of expectancies, cues, goals, and actions; and the categorized intentions into tactical strategies. Although the language of each framework is different, the numerous similarities and complementary elements between them have been highlighted in the above sections. Therefore, integrating the three frameworks into strategies is not difficult unless too much interpretation of the results is required. From this research, it was found that a lot of interpretation was necessary to combine the three cases of engaging an enemy UAV into a tactic. This was for a few reasons. First, in terms of actions, none of the subjects displayed a complete strategy from detection, maneuvering, engaging, and destroying. Second, in terms of processes, all of the subjects unsuccessfully applied stored rules of running fire that violated the simulation's design. Third, in terms of strategies, all of the subjects answered the engagement survey question by stating they would tend to ignore the UAV. Though careful interpretation was necessary to ultimately produce a unified, successful tactic, this success is taken with a note of caution due to the biased nature of the interpreter of the results also being the designer of the simulation. Thus, a feedback loop in Figure 6-1 exists from the "Strategies" block to the "Human-in-the-loop Experiments" block if too much interpretation was involved in forming the strategies. The notion of too much interpretation, as discovered in this research, can imply any or all of the following three items:

1. Large variance in scores
2. Inconsistent application
3. Cognitive process analysis indicates lack of expertise

The limitation in this research is that there was not time to carry out this feedback loop of further experimentation.

6.6 Translation

Once the human-inspired strategies and tactics have been learned, the next step is to translate them into the autonomous vehicle's (AVs) problem space, knowledge states,

and operators. This particular wording of “problem space, knowledge states, and operators” is purposely taken directly from Newell and Simon’s work [47]. Newell and Simon identified invariant characteristics across all human problem-solvers and cast them in terms of information processing, goal-oriented computers. If human strategies are goal-oriented, then the seven characteristics of goal-oriented behavior given by Newell and Simon (see Section 3.2.4) should be a guide in encoding these strategies. Note that the reason for this translation of human strategies into AV-domain language is due to the serial processing nature of computers, the differences in sensing the environment between humans and AVs, and the unique strengths of AVs, such as long endurance, more aggressive maneuvering, and shorter time-delays in response to inputs. Although there are no good answers for how this translation process should occur, the need for translation emphasizes the importance of capturing strategies and not just actions. By learning goal-oriented strategies, the human-inspired tactics act as a template to guide the AVs actions.

6.7 Encoding Tactics

As described by this research, there are two ways to encode human-inspired tactics translated for the AV. The first is through statecharts. The design of statechart logic helps capture the representation of complex, reactive environments through states and transitions. Learning human-inspired, reactive tactics fits naturally into this statechart logic. This research provides one example of learning, encoding, and applying human reactive decision making to successfully engage three different enemy platforms. Statecharts, though, are static. Their structure does not change with time. In contrast, human subjects construct partial plans in determining how to accomplish the reconnaissance mission because it is impossible and irrational to define a rigid structure of actions and goals that is also static. The uncertainty of the future necessitates a plan that is partial. This plan will be dynamically filled in as the scenario progresses, and it is the intentions of the human subject which guides this process. Therefore, it appears that statecharts should not be applied to implement the partial plans and intentions which guide the searching behavior of the vehicle. Rather, as discussed in Chapter 5, forming intention and execution functions are the second way to encode human-inspired tactics translated for the AV.

To further clarify, consider the concept that intentions are different from actions. It is not correct to observe the human subject’s actions and declare that every decision made was also an intention. Recall that intentions are primarily future-directed. When the human subject is intending to search through all of the critical area and he suddenly contacts an enemy SAM, the subject must reactively decide how to evade and/or engage the SAM. It is not as if the subject abruptly forms a new set of intentions in the exact moment of detection to evade the SAM. Rather, in the language of Newell and Simon, the subject must depart the current branch in the heuristic search path over the problem space and pursue the sub-goal of avoiding the enemy threat. After this sub-goal has been achieved, the subject must observe the consequences of the sub-goal’s actions and determine how best to return to the

previous main branch of searching through all of the critical area. This main branch is the consequence of the subject's intention. The pursuit of the sub-goal is the reactive actions taken. Thus, intentions differ from actions. As noted above, statecharts can capture these sub-goals, but they should not capture the main heuristic solution path governed by intentions.

6.7.1 Integrating Reactive and Planning Elements into Tactical Control

Figure 6-2 depicts the proposed block diagram of integrating both statecharts, intention, and execution functions into a tactical control level for the AV. Beginning on the left of the figure in the blue blocks are the input variables. As discussed in Section 5.6.4, the terrain layout, probability of enemy contact, and total time to accomplish the mission are inputs to the intention function, shown here as a fuzzy logic controller. The fuzzy logic rules come from the learned strategies. Outputs from this fuzzy logic controller are the intended sequence, coverage, and time of search for the scenario. These intended variables are combined with the terrain database and vehicle sensor radius and passed to an execution function. This execution function is a library of maneuver patterns learned from the humans that outputs a matrix of waypoints and velocities. As long as the radar sensor does not indicate any enemy contacts, these waypoints and velocities are passed to the low-level, trajectory-following controller of the plant. If the radar indicates an enemy contact, the AV activates its reactive, statechart logic to either engage or avoid the enemy. While the statechart is active, it outputs waypoints and velocities, which are passed to the trajectory-following controller. The feedback loop from the plant to the execution function allows the execution function to monitor if the initially desired intentions are no longer feasible. If they are not, one option for the execution function is to keep trying unsuccessfully to implement them, and its outputs are saturated. Another option is to add another feedback loop from the execution function to the intention function to recalculate intentions for the remainder of the scenario.

6.7.2 Missing Information?

Before submitting these encoded strategies and tactics to Monte Carlo simulation, it is important to ask the question of whether enough data exists in deriving the statecharts, intention, and execution functions. Are the thresholds for the intention and execution functions properly defined from the human experiments? Did important transitions have to be added between states because the humans did not exhibit such behavior? In order to maintain a human-centered, human-inspired approach to learning tactical knowledge, missing gaps in the data should try and be filled by more experimentation, as shown by the second feedback loop in Figure 6-1.

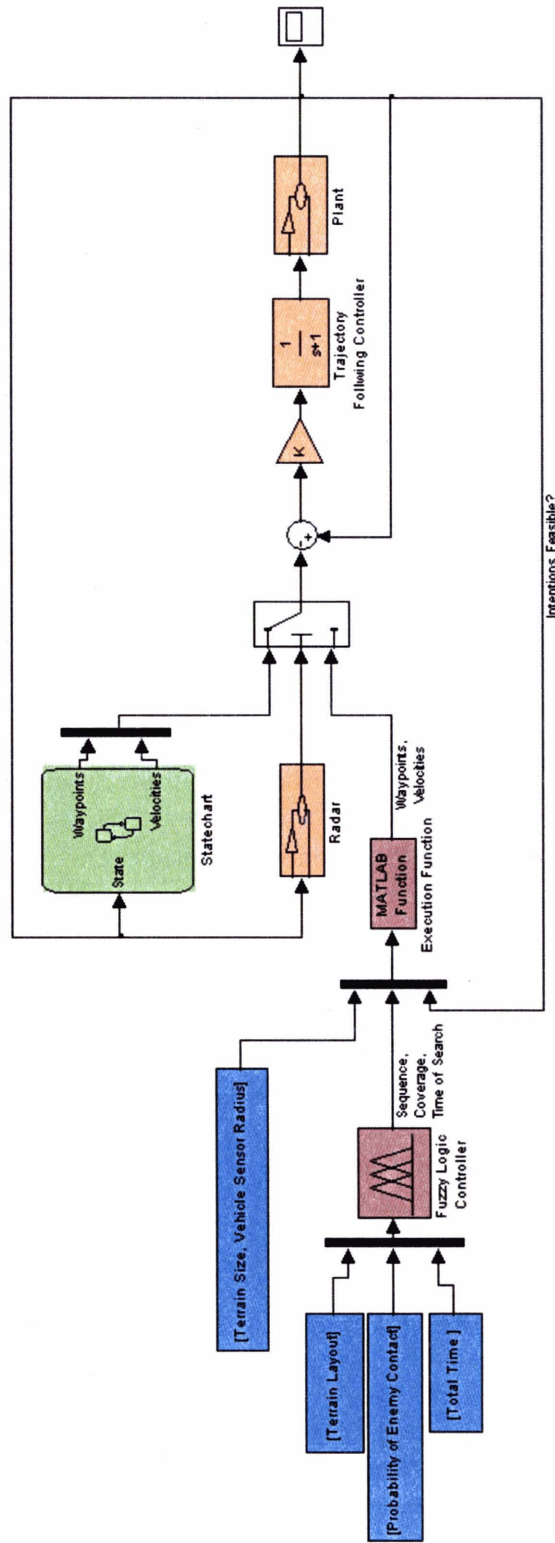


Figure 6-2: Simulink diagram of reactive and planning tactical control for the AV.

6.8 Testing over Large Sample

Monte Carlo simulation should help test for both the success of the tactics over large number of cases as well as verify the correct translation of the human-inspired strategy into the AV domain. In terms of robustness, statistical analysis of the results will verify the applicability of the tactic, which was derived from a small sample of human experiments, over a much larger sample with randomized parameters. If the tactic was unsuccessful over the large sample, further human-in-the-loop experimentation is needed to find a better one. In terms of translation, if the AV begins to no longer behave like the human subjects in executing these encoded tactics, the translation process should be questioned. For example, if the tactic was too successful over the large sample, such as the engagement tactic against the UAV and SAM, the translation process indicates a bias to “curve-fit” the data.

6.9 Testing for Reliability

After testing out new tactics for the AV with Monte Carlo simulation, a final step should occur that did not in this research. The human expert should be brought back in to help grade the AV’s performance. The main question is not necessarily did it achieve the mission objectives. Rather, the main question is whether the human expert agrees with and understands the decisions made by the AV. As the automation design flowchart depicts in Figure 2-1, the primary evaluative criteria of the automation is the effect on human performance. As stated in Section 2.2.5, we propose a second set of primary evaluative criteria depicted here in Figure 6-1 of team performance consequences. As the human expert observes the AV behavior, does the expert believe the AV is making decisions like a team player characterized first and foremost by reliability? Note that this step should occur before conducting a second phase, not shown here, of optimizing the tactics (see Section 1.4). This caution is due to the human expert’s evaluation of whether the AV has truly been augmented with human-like strategies, not just the actions. For this first stage of incorporating human tactics, it is the learned strategies that guided the translation process of human tactics into the AV domain which are under review. It is important to first confirm that the AV is actually pursuing the right goals rather than how the AV is pursuing them. Thus, this final consultation with the human expert is crucial to designing this automation to be a reliable team player.

Chapter 7

Future Work

After concluding the research in this thesis, we present two items for future work. The first is an elaboration on the intention and execution function, which was discussed in detail but not implemented. The second is an extension of this research into the very tough problem of identification of an unknown contact.

7.1 Representing Intent

While this research proposed the formation of an intention and execution function as a method for capturing human-inspired tactics for the search problem, it did not actually create and test them. To do so would be the first, natural extension of this research. This is an interesting problem for two reasons. One, if the human is pressed to truly specify, for example, his or her intended time spent searching in both the air corridor and critical area, the human will very likely reply by classifying the time with fuzzy boundaries. The human intends to spend “most of the time,” a “good portion,” or “about one minute . . . maybe a little more” of the time searching. This kind of fuzzy thinking is extremely common in human cognition, and thus, fuzzy logic controllers offer a natural way of encoding this thought process [17]. Two, the tension in creating the intention and execution functions is due to the presence of a hybrid control issue. The discrete boundaries in the fuzzy variables, such as spending “most of the time” or only “some of the time” searching, have to result in continuous outputs for vehicle control.

7.1.1 Fuzzy Logic

A fuzzy logic controller is made of up three parts: membership functions, rules that act in parallel, and an aggregation scheme to reach a final answer [41]. To illustrate the concept of a membership function, consider the probability of enemy contact as the input to the intention function. In this research, this probability was already given to the human subjects with fuzzy boundaries because it conforms more naturally to human thinking and to “noisy” intelligence information. The probability of enemy contact was classified as “very good, possible, or slim.” The membership function

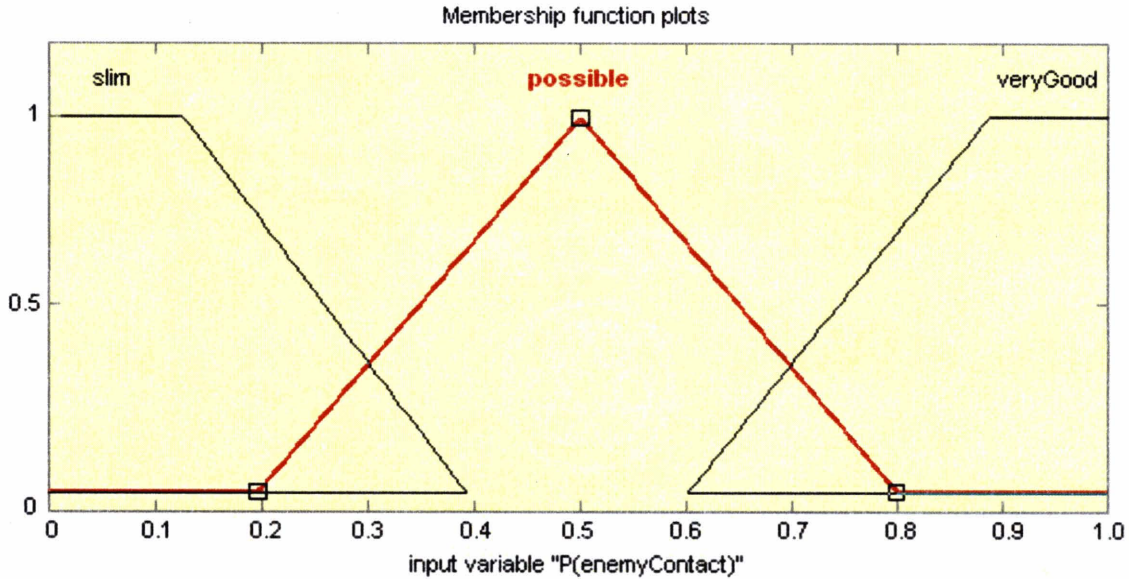


Figure 7-1: Overlapping membership functions.

describes how much an exact, quantitative probability of enemy contact between zero and one belongs to one of the three categories of very good, possible, or slim. These categories are called fuzzy sets. For instance, say the probability of enemy contact was calculated to be 65%. Is 65% a very good chance of enemy contact or only a possible chance? Does 65% have to wholly belong to either the category of very good or possible, or can it partially belong to both very good and possible? The idea of a membership function is to allow an input, such as 65% probability of enemy contact, to have a degree of membership in fuzzy sets rather than a binary yes/no membership. This degree of membership is scaled between zero to one. Note that the degree of membership across the membership functions does not have to sum to unity at every point. In Figure 7-1, a 65% probability of enemy contact would be “fuzzified” by having a membership degree of 0.5 in the “possible” fuzzy set and a membership degree of 0.2 in the “very good” fuzzy set. Note that Of course, there are an infinite number of membership functions, and the designer would have to choose one. This could take the shape of Figure 7-2, where the fuzzified output would have always a unity membership degree in one of the three fuzzy sets. The point here is that the designer is free to choose the membership function for each of the inputs to the intention function. The goal of the human experiments is not to determine the exact shape of these functions, but to derive the fuzzy inference rules.

The fuzzy inference rules are logical statements with antecedents and consequents. For example, if the probability of enemy contact is very good and it is in the critical area, then coverage of the critical area is equal to “all of the critical area.” For this inference rule, “very good” and “in the critical area” are the fuzzified inputs and “all of the critical area” is the fuzzified output. All of these are fuzzy sets represented by membership functions. Also, just like the degree of membership was not a binary yes or no, the Boolean operator “and” in the example inference rule

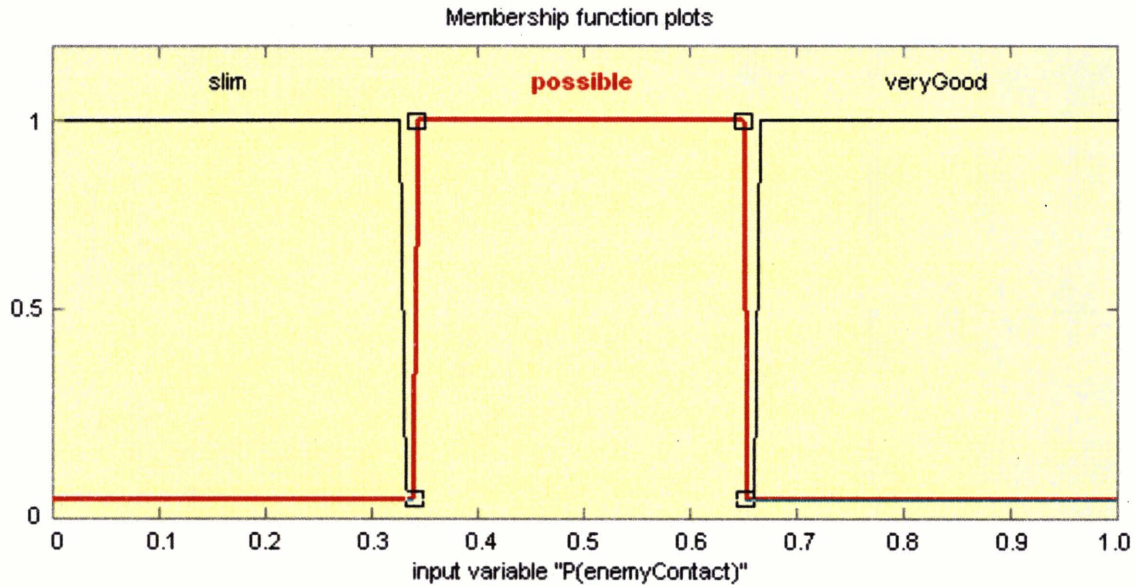


Figure 7-2: Membership functions that always result in 1.0 degree of membership.

is abstracted in fuzzy logic away from simple binary application. Though there are fuzzy operator property axioms that allow the operators of “and, or, not” to be extended differently depending on the designer’s wishes, they typically are defined in fuzzy logic, respectively, as “minimum, maximum, complement.” To clarify, say the probability of enemy contact was given as 65%. This resulted in a 0.2 degree of membership in the “very good” fuzzy set from its membership function. For the variable of “in the critical area,” this is, by definition, either a full membership of 1.0 or not. Therefore, we have two fuzzy variables in the antecedent of the fuzzy rule. The probability of enemy contact fuzzy variable is “very good” to the degree of 0.2. The “in the critical area” fuzzy variable is “in the critical area” to the degree of 1.0. By combining these two variables under the fuzzy “and” operator, the minimum of their membership degrees is taken. The rule’s antecedent - if the probability of enemy contact is very good and it is in the critical area - is true to the degree of the minimum of 0.2 and 1.0, which is 0.2. This 0.2 is then used to clip the fuzzified output’s (“all of the critical area”) membership function at a height of 0.2. Note, however, that the membership function of the output variable in the rule’s consequent is not constrained in its shape. In fact, the membership function could be a simple spike at some value along the abscissa, and no matter what the resulting rule strength was calculated to be, the output variable would always take on the same value.

If the fuzzy logic controller contained multiple rules, then all of these rules would act in parallel. Thus, if a probability of enemy contact of 65% had a membership degree of 0.5 in the “possible” fuzzy set and a membership degree of 0.2 in the “very good” fuzzy set, then any rules that contained in the antecedent the possible or very good chance of enemy contact, as one of the fuzzified input variable, would be activated. As discussed above, the strength of the rule would be calculated by the fuzzy operators, and the output membership function of the rule would be clipped

off at the height corresponding to the rule's strength. The final step, then, in fuzzy logic control is to aggregate the fuzzified output variable's membership functions into a single "crisp" value. Though there are many options for this process, they typically take the form of centroiding, where the centroids for the clipped output membership functions are found and combined.

7.1.2 Strategies and Fuzzy Rules

Although this is a very brief description of fuzzy logic control, the concluding remark is that the purpose of learning human-inspired tactics would be to find the fuzzy rules that resulted in successful human performance. For example, a successful search tactic exhibited by S1 was his focus on slowly and deliberately searching the critical area when the probability of enemy contact was very good. This strategy could be cast in the form a fuzzy rule that states if the probability of enemy contact is very good and it is in the critical area, then spend most of the time searching in the critical area to cover all of it. All the designer must do after identifying these rules from the human experiments is to define the membership functions and the exact nature of the fuzzy operators and aggregation process. Note that Simulink now offers a fuzzy logic toolbox, which has built-in fuzzy inference engines [41]. These inference engines observe the inputs/outputs of a system and build fuzzy rules that mimic the system's behavior. It has not been investigated whether the design elements of uncertainty, hierarchy of objectives, and flexibility in solutions for the human-in-the-loop experiments would hinder the accuracy of these inference engines. Remember the goal was to design scenarios with many elements that are not quantitative, and this qualitative aspect of the scenarios may be too subjective for the inference engines to handle accurately. Of course, this is one of the weaknesses in current artificial learning algorithms and the reason why this research occurred. Therefore, the first extension of this research would be to take the searching strategies learned from the human subjects and map them into fuzzy logic rules to create an intention function.

7.2 Identification

For the interaction sequences between simulated entities, it was assumed that there was no identification phase. The human subjects were told that every vehicle encountered should be assumed to be hostile. Identification is an extremely hard problem, and in tactical situations, to identify another contact incorrectly can result in grave consequences, such as friendly fire incidents. Could the process of learning tactical knowledge from human experts help in finding better solutions for the problem of autonomous identification of contacts? To begin answering this question, a step-by-step process could be envisioned as follows. Say the scenario was an unmanned undersea vehicle (UUV) patrolling the entrance to a major port. An enemy diesel submarine is attempting to covertly pass by the UUV so that it may perform stealthy reconnaissance of the port's activities. First, the simulation environment is built to reflect the desired level and scope of expertise to be learned. Next, a human expert, such as a

former attack submarine skipper, is employed to play the role of the enemy submarine that is trying to deceive and sneak by the UUV. The evaluation of the skipper's actions, strategies, and processes follows along the path proposed by the framework in Figure 6-1. At this point, the human-in-the-loop experiments and evaluations end with a set of expert intentions and strategies for how to deceive and sneak past a UUV into a major port. To accomplish all that would be a tremendous step in the right direction, and the research presented in this thesis fits naturally into the process outlined above. However, this would only be one-half of the solution to the identification problem. The other half is to reverse the application of these human-inspired tactics.

No longer is the UUV trying to implement these human-inspired tactics with its own control logic. Rather, the UUV is attempting to observe these human-inspired tactics in other platforms to assess their intent in moving towards the major port the UUV is patrolling. This is now a pattern-matching, predictive problem. The UUV is trying to predict the platform's intent by pattern-matching the platform's actions to the UUV's knowledge base. Though at face value it seems tremendously complex, there is good reason to believe that the process of learning and applying human-inspired tactics is helpful towards finding a solution. In Chapter 3, it was noted that the human inadequacy to fully understand and abide by probability theory has led researchers to call for the end of human experts taking on the role of prediction under uncertainty [15]. Rather, the role of human experts should be to identify the important predictor variables, show how to measure and encode them, and identify the correct direction of the variable's weighting (i.e. - higher diesel engine noise equates to greater chance of detection). At this point, statistical analysis through computation should make the appropriate prediction.

Through understanding the submarine skipper's strategies of how to sneak past the UUV, the designers can identify and program the UUV to look for the right predictor variables, measure the uncertainty, and calculate a statistical prediction of the other vehicle's intent. Now the human expert can be employed in several ways to verify the reliability of this pattern-matching, predictive process, much like the final step in Figure 6-1. One, the human expert can observe and consult on whether the UUV was able to focus in on the right actions of the other vehicle and interpret those actions correctly. Two, the human expert can continue to play the role of the enemy diesel submarine and actually try to sneak past the improved UUV. Three, the human expert can play the role of the UUV and try to identify another contact. This last use of the human expert might be best made into another process of human-in-the-loop experiments to learn and apply tactics.

Appendix A

Introductory Documents and Presentation of Scenarios to Human Subjects

Tactical Knowledge Elicitation Experiment: Instructions

This is a human-in-the-loop simulated exercise. Your role is to be a helicopter pilot for the U.S. Army. You will be presented with two practice scenarios in order to feel comfortable in the simulation environment. You will then be given five different scenarios. Each scenario will receive a score based on how well you performed. These scenarios are described in the following pages. For each scenario, there will be a concise summary of any prior intelligence before the scenario begins. During each scenario, you will be asked to "talk aloud," to verbalize your thoughts. (explanation given below) After all five scenarios have been completed, you will be asked to complete a survey. Please feel free to ask any questions at any time. The paragraphs below describe all background information and common mission objectives across all scenarios.

The Background:

The U.S. Army has three companies of troops standing by in the Green Zone. They are ready to deploy northeast to help suppress rioting in several towns which have arisen due to insurgent activities. Several Army Blackhawk troop-carrying and Apache escort helicopters will deliver the soldiers to their destinations.

The Mission:

Your task is to scout out a predetermined air corridor and critically important regions. We need to make sure the condition of the air corridor is safe for the passage of our troops. Furthermore, we need to identify any enemy contacts in critically important regions. The mission will be complete when you have searched over as much of the corridor and critical areas as possible, have either declared the air corridor as safe for

passage or not, and returned to base (the original starting point).

The Objectives:

- We only have a small window of time to send you in ahead of the troops to scout out the area. You have a *five minute time limit*.
- During that time, it is important to cover as much of the air corridor as possible, while being careful to avoid all unnecessary contact with enemy vehicles.
- If it is necessary to engage any hostile contacts, you have a total of *five* shots.
- Also, most scenarios will have a critical area for reconnaissance. If a region is defined beforehand as critical, we believe it could either be used by the enemy as an ideal ambush site for our troops, or we hope to use the terrain for landing zones. It is very important to identify any unsafe conditions which might exist in critical areas.
- Assume all contacts are hostile. It is at your discretion whether to engage or avoid and report all contacts. . . whichever helps achieve the mission objectives.

The Intelligence:

Based on satellite imagery and anonymous tips, we will try to quantify whether to expect any enemy contacts or not during the mission.

General Information:

1. Joystick and Simulation Setup - please notify me when you have read to this point.
2. Scoring - the final score to be calculated after each scenario is composed of the following parts:
 - (a) How much area was covered both in the air corridor and the critical area.
 - (b) The ratio time spent in the critical area to time spent in the air corridor. The more time spent scouting the critical area than in the air corridor, the better the score.
 - (c) The total exposure time to the enemy. This includes both the time spent in the enemy's circle, and the time spent in the enemy's weapons cone.
 - (d) The total time the enemy was exposed to you. The same two parts apply as above.
 - (e) The number of hits scored divided by the total number of shots taken.
 - (f) A very small penalty is incurred for the total amount of time vehicle is at maximum velocity. (to simulate limited fuel)

3. Talk (Think) Aloud:

To talk aloud is to simply verbalize your thoughts as they occur. It is very difficult for humans, once they have finished a task, to remember back and verbalize the thoughts they had in the past. Humans cannot consciously articulate or recall all the reasons why they chose a particular course of action or not. Therefore, the purpose of talking aloud is to help me, the interpreter, when I am reviewing your actions. Not only can I replay the game, but by having recorded your statements, I can gain further insight into the motivation behind your actions. Here are some instructions:

- (a) There is a difference between thinking aloud (talk aloud) and explanation. To think aloud is a short-term time stamp on your thoughts. It is to simply speak constantly as if alone in the room. This is what we want you to do; to verbalize your thoughts. Do not worry about coherency. On the other hand, to try to provide an explanation is to try to access long-term thinking. That sort of thinking prevents you from focusing at the task at hand. So don't explain, just talk.
- (b) I will give a simple verbal reminder to keep talking aloud if you fall silent. After twenty seconds of silence, I will say, "keep talking." Otherwise, I will keep silent during each scenario.

4. Symbology

	Normal	Detected	Locked	Damaged	Destroyed
Your Vehicle					
Tank					
S.A.M.					
U.A.V.					

Round 1 Intelligent Reports:

[NOTE: slim < possible < very good]

Practice 1: There is a possible chance of enemy contacts both inside the air corridor as well as in the critical area.

Practice 2: There is a possible chance of enemy contacts.

Case 1: There is a very good chance of enemy contacts inside the air corridor as well as a slim chance of enemy contacts inside the critical area.

Case 2: There is a slim chance of enemy contacts inside the air corridor as well as a possible chance of enemy contacts inside the critical area.

Case 3: There is a slim chance of enemy contacts inside the air corridor as well as a very good chance of enemy contacts inside the critical area.

Case 4: There is a slim chance of enemy contacts inside the aircorridor as well as a possible chance of enemy contacts inside both of the critical areas. Both of the critical areas are equally as important.

Case 5: There is a very good chance of enemy contacts inside the air corridor as well as a very good chance of enemy contacts inside the critical area.

Round 2 Intelligent Reports:

[NOTE: slim < possible < very good]

Practice 1: There is a possible chance of enemy contacts both inside the air corridor as well as in the critical area.

Case 1: There is a possible chance of enemy contacts inside the air corridor as well as a very good chance of enemy contacts inside the critical area.

Case 2: There is a possible chance of enemy contacts inside the air corridor as well as a very good chance of enemy contacts inside the critical area.

Case 3: There is a very good chance of enemy contacts inside the air corridor as well as a possible chance of enemy contacts inside the critical area.

Case 4: There is a possible chance of enemy contacts inside the air corridor as well as a slim chance of enemy contacts inside the critical area.

Case 5: There is a very good chance of enemy contacts both inside the air corridor as well as in the critical area.

Case 6: There is a slim chance of enemy contacts both inside the air corridor as well as in the critical area.

Case 7: There is a possible chance of enemy contacts inside the air corridor as well as a slim chance of enemy contacts inside the critical area.

Appendix B

Surveys

Survey (01 FEB 06)

Note: S3 did not have time to finish the last case and fill out a survey for this first round of experiments.

1. Describe any strategies you used during each scenario. In other words, what were the motivating factors that drove your actions?

S1: Thoroughly investigate the critical areas such that every part of it was completely covered. Accomplished with tight weave that one pass started where previous left off. In the non-critical area, do a weave that covered most of the area, but not necessarily the whole thing. If corridor was narrow enough, I centered myself and then just went straight. If I had less of a chance of getting killed in the corridor, I would do that first, banking on the fact I would still be alive in the critical area. When I got the critical area, I would then thoroughly investigate it.

S2: Quickly scout air corridor to cover all area, then methodically go back and forth in critical zone to rack up points for staying in zone. Went after bad guys a little too much because I got shot twice.

S4: I tried to search/cover the areas in a sequential manner, starting from an edge and working to the other. If there was a higher chance of enemies in one area, I tried to stay out of that area until I saw the other areas. If I stumbled upon a target, I would go for it if it was mobile and within reach. Otherwise, I'd come back for it.

S5: My main priority was to accomplish all my mission tasks. I sorted the missions to ensure I could get as much done as possible; I did low risk tasks first and quicker and tried to go slower and finish with the "riskier" tasks.

2. Describe as best as you can any tradeoffs you had to make between scouting a critical area while anticipating enemy contacts.

S1: If I had just gone into the corridor, I would try to run and investigate more of the critical area and then go back to engage. The tradeoff was, thus, area covered versus eliminating enemies for the guys to come behind. I figured having as much investigated as possible was better because if the guys knew

what was waiting for them, it would be better than if I killed the first enemy I saw but died in the process.

S2: Considered risk a very small factor. Just went about scouting and if I ran into bad guys, I dealt with them.

S4: I started scouting the area as if there were no enemies. Once I came upon one, I'd often pursue it for a little, which wasted some time I could have spent scouting.

S5: I traded off speed with a more deliberate process. My goal was to give myself an area of safe maneuver, and ensure I kept that safe area available to me in the event of contact.

3. Rank the following constraints in the order of importance or influence they had in your actions. In other words, the most important constraint was also the most limiting. (1 = "most important", 5 = "least important")

- (a) Maneuverability of vehicle (turn rate, velocity)
- (b) Ammo limitations of five shots
- (c) Four minute time limit
- (d) The field-of-view of the weapons sensor
- (e) Having to return to base

(Sequence of numbers correspond to sequence of constraints as listed above. For example, S1 ranked the "Maneuverability of vehicle" constraint as 3 and the "field-of-view" constraint as 2.)

S1: 3-4-5-2-1

S2: 3-2-4-1-5

S4: 2-4-3-1-5

S5: 4-1-3-2-5

4. In each scenario, the width of the air corridor was greater than your vehicle's sensing capability. What kind of maneuvering patterns did you find were effective to be able to search over the entire area?

S1: The weaving pattern I described above. If I was done with the critical area, I could "tighten" the weave such that my motion was a more side-to-side motion relative to the corridor. This lessened the missed areas.

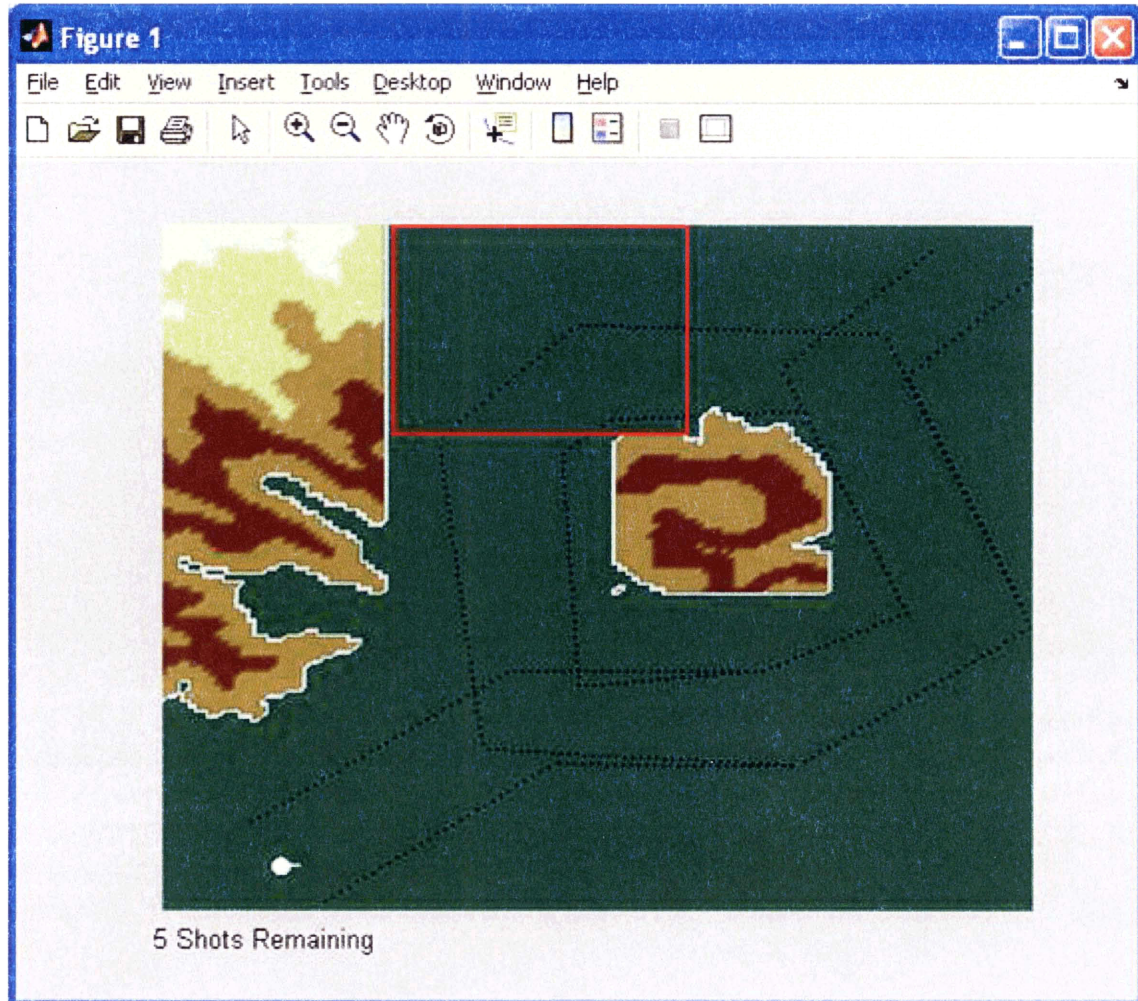
S2: Long stripes - first do the outside loop of a corridor, then the inside.

S4: I just picked a side initially, so that if I had time, I could easily cover the whole thing later by covering the other side.

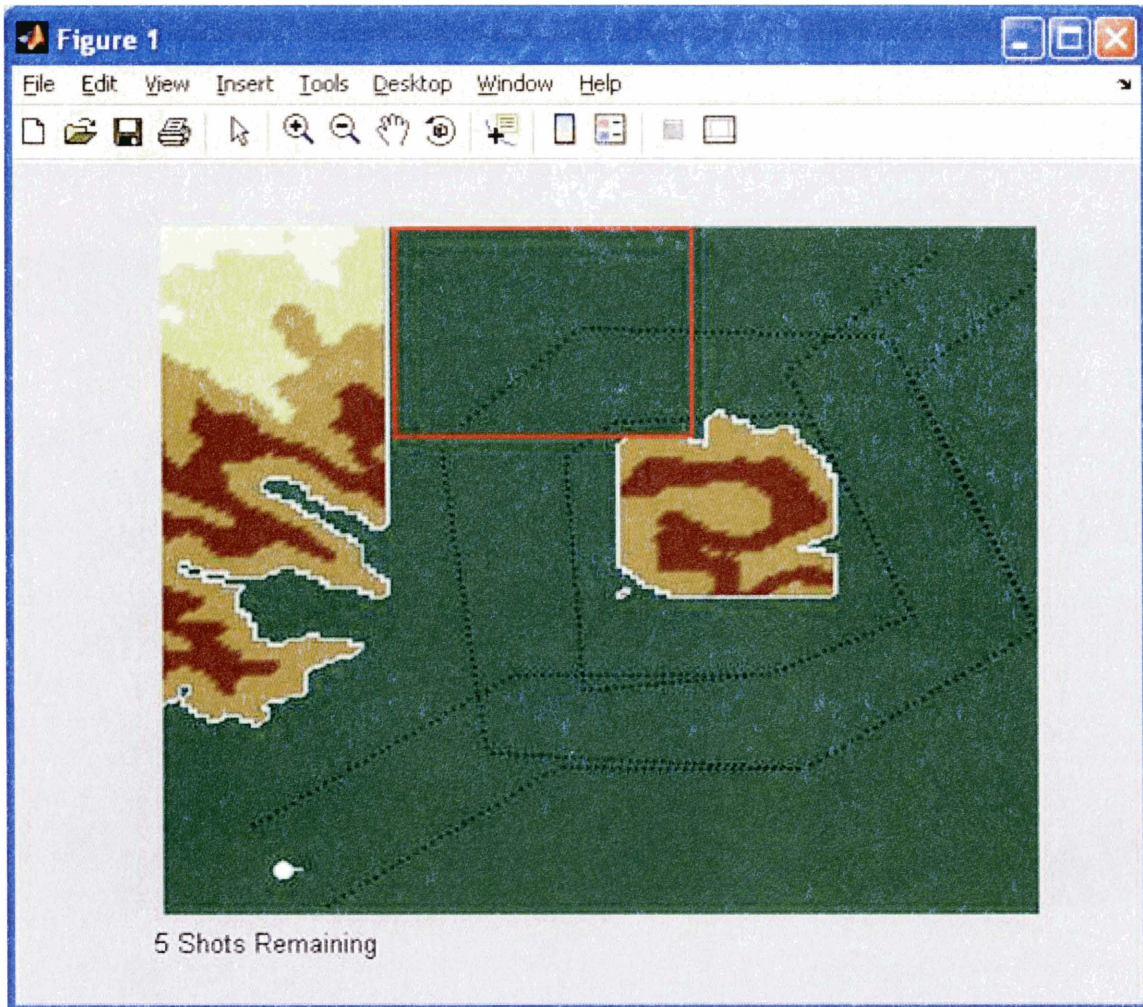
S5: Zig-zag on hasty recon areas and a more deliberate box maneuver on deliberate areas.

Survey (16 FEB 06)

1. Based on your experience in this simulation, draw out how you would plan to search through this entire region (both air corridor and critical area) given the following intelligence report:
 - (a) probability of enemy contact in the air corridor is slim and in the critical area is slim



(b) probability of enemy contact in the air corridor is slim and in the critical area is very good



2. To the best of your knowledge, are you prone to engage or avoid the following three threats? What factors are involved in either engaging or avoiding? If you had to plan your strategy ahead of time, what actions would you take if you unexpectedly ran into each of these three threats in a typical scenario?

(a) Ground

S1: Engage - They had a smaller sphere of influence (SOI) and so that gave me the advantage engaging them. If they were not mobile, I would come to a hover such that I could engage them, but they couldn't engage me. If they were mobile, I would try to maintain separation such that I could engage them, but they couldn't engage me.

S2: Prone to engage. Poor at shooting me. Relatively slow. Pursue if time allows for finishing area coverage.

S3: Very likely - easy target for me to hit, with small probability that it would be able to hit me. Without a weighting factor for which is more important - exploring the corridor and critical area or taking out enemy contacts - I would probably engage the enemy at the cost of exploring as much area as possible. It seems to me that part of the mission was to ensure the passage was safe.

S4: These are easy in that they are slow and have small targeting areas. I would typically engage them because they didn't involve much of a threat or repositioning.

S5: Engage. I have range on him and can expect that he is a threat and I can kill him more often than he will get me.

(b) Air

S1: RUN!! (Actually, not really ...defensive turned into dogfighting.) They had speed and maneuverability that pushed my feeble pilot skills to the test. I tried to be defensive to start and then tried to dogfight them from there.

S2: Prone to ignore. Too fast for me to shoot. Ignore, but don't get shot.

S3: Not very likely to engage: it was a difficult target to hit while it didn't seem to typically be that hostile. If I were to engage I'd have to be in a good position behind it and gaining on it. Because a UAV's position is less predictable than a ground target's, I don't see it as a predictable threat to the ground forces that will move through the corridor.

S4: These were fast and often tried to track you. Initially I engaged them, but this proved to be a waste of time, generally speaking.

S5: Engage if given opportunity, but not a high probability engagement (i.e. - I will have to chase him) and not my mission.

(c) SAM

S1: This one I really did try to avoid contact. They had the big SOI, so that put me at a pretty big disadvantage. Originally, I tried to engage

with speed but strafing led to death. So in the future, I ran.

S2: Run away. Shoots me easily. Avoid completely.

S3: Not very likely to engage. It's too dangerous. Having a greater range than I do - and where I would have to maneuver quickly and precisely to defeat it. I think it better to move on so I can accomplish the mission of clearing out the corridor. My strategy is to note the location and get out of its range as quickly as possible.

S4: These had huge targeting areas, and I tried to avoid them until the very end when I had finished everything else. I wasn't aware that you were penalized for dying and may have approached that differently now that I know that.

S5: Avoid - Don't want to get killed and not be able to accomplish my mission. His range is greater than mine, thus engagement is not in my favor.

Appendix C

Think Aloud Reports

C.1 Round 2

Note that for each subject, the cases are presented in order as seen by the subject. This order was randomized between subjects to diminish unwanted training effects (see Section 4.2.5). In the following reports, “CA” and “Air Corr” are short-hand for “critical area” and “air corridor,” respectively.

C.1.1 S1

Case 3

Plan

given my lack of skill at flying
go to CA ASAP
don't want to be in Air Corr b/c I don't want to get killed
go directly to CA and stay outside Air Corr
tight weaving pattern in CA
all assuming I survive that long, which is a pretty bad assumption
CA first

Thinking Aloud

outside Air corr to begin with to increase probability of surviving
go to top and start weave up there
trying to cover all ground in CA, nothing special, just going back
and forth
not doing as tight of a weave here (Air Corr)
going pretty fast, better slow it down here
there's a guy, I kinda want to get away from him, I'll turn around
and see if he's still there, engage him
two minutes, alright, that's good
speed up while doing it (going through last upper part of Air Corr)

I do want to make it back into the CA
Whoa, not flying too particularly well
looks like a pretty good place for bad guys to hide
spend remaining time in CA
no rhyme or reason what I'm doing here, just trying to stay in CA
I did go along the corridor there, so I guess, there's the logic

Debrief

S1 was asked, "Why did you say, 'this is a good area for the bad guys to hang out at?'?"

S1 answered, "two areas I thought would be very good

- o There you're turning the corner
- o They have inlets/nooks for them to hide
- o You can only approach them
- o There's a good cone for them to view you
- o But pretty tough attack region for you
- o Other area was that strip in southeast corner
- o Going through narrow corridor section
- o To the west there you can bail out outside the corridor
- o But there you're pretty confined
- o And sure enough that's where the bad guy was
- o Maybe I should have investigated that more and tried to find him and engage him
- o But I did want to go back towards the CA
- o So that's why I left"

Case 2

Plan

S1 was asked, "You were talking about terrain earlier, narrow channel, enemies hiding in there, would that necessarily affect how you would search through any parts of the map where you might be confined, would your search pattern change, what might be different?"

S1 answered, "On the previous run, I was more mentally alert...I kinda expected the enemy to be there. I guess I was paying more attention and expected to turn at any point or speed up to get away. But for the actual search pattern, I tended to stick to the same weaving pattern, but then there might be a scenario where I would think to do that last. And once I was there, to do it pretty thoroughly. So the answer is, it kinda made me on higher alert expected to change, and then also, I might put that off to the last to ensure survival and see the other things first, but once I was

there really try and thoroughly investigate it.’’

‘‘on top of northern hill, kinda has same protection area
Seems like that would be a good place for enemy to hide and pop out
So what I might do is take it wide around that corner so it won’t
catch me off guard
Increased chance of survival
Since there’s less of a chance of hitting an enemy in the corridor,
I’ll investigate that first and then go into the CA’’

S1 was asked, ‘‘any idea of how you want to investigate those gigantic lobes in there?’’

S1 answered, ‘‘the bottom one: the southwest corner, since you’re starting off at the very corner, you can weave back and forth north/south, hopefully as you tighten up you can just go north along the narrow part of the corridor...

I was debating east/west vs. north/south at the top one

It might be better to go east/west but then

If you divide that section into two parts you do east/west for the first loop

Do the southwest portion of it

So you can loop around and go into the critical area

If you do north/south you put yourself right at the inlet and if you’re going north at that section, you put your back to the enemy

So if they pop up, they’re on your six and you’re pretty hosed at that point

If you’re east/west you at least have them in your sights

You have a least a chance of engaging them before they engage you’’

Thinking Aloud

just go north/south in this general area

I’m always trying to get with my radar so that it’s just to the edge of the corridor

That’s what I’m trying to do

Since this is just the corridor and not the critical area, there are definitely

Whoa, shoot, bad guy, that’s not good

Well he’s got a way wider influence than I do

I’m just going to skip by him and may come back to him at the end

Because I want to survive to the CA

Now I’ll go up to the top part

And then do east/west search pattern

(I call out time)

I’m gonna speed up in here, kinda cruise through this section

I do want to leave time to try to engage that guy
Notice I went a little bit up because I was closing in on that mouth
I'm trying to make it so I can
Yeah see I turned there to see if I could engage if anybody comes out
Turn towards it so that my
Now I'll go to the CA
Probably die, that's ok
Whoa, ok, almost smacked the hill there
I'm not doing a very good job weaving here
Call this the drunk sailor approach, no offense, jon
You don't really drink
Oh, gotta run, gotta run, gotta run, gotta run, go, go, no, no
Run, no run, shoot oh, I'm gonna die
Ok, now I can go on the offensive here
(laughs) that was interesting
I don't know why, I tend to pull back
I don't mean to slow down, it's like me trying to pull up
I don't why, that's pretty stupid
So I'm going to stay looking for that guy
Cause the SAM is in a known location so I can relay that back to the
guys and they'll know where that is
The UAV is definitely mobile, definitely almost got the best of me
I'll stay in the CA just to search

Case 1

Plan

So in this one, there wasn't a clear, clean-cut, I could go along
the corridor and do this
I definitely want to spend a lot of time in the corridor
There's the terrain area, like you said, that really confines it up
in the northern portion of the CA
Whereas in the southern, you can bail out in all directions
But you're really confined there in the northern half of it
Follow corridor, head west around terrain
That area could be squirrely because of the narrow space
Not too worried about that
Go north and that will put you into the northern portion of
the CA
With your sights set on possible enemies coming in
Very good chance of enemy contact in CA
One thing I need to do a better job at is to have my firing cone
pointed in the right direction
At the dangerous spots

That's why I'm trying to go into the northern section in a good,
strong engagement position
Go into northern portion
Go west to east
There's that inlet right there where a bad guy could be hiding
So I can turn south and be facing that direction so that if some
guy comes out
Swing back and check the other hill area looks like a turtle head
In the turtle neck area, a guy might hiding
But that will put me facing the right direction
Then I can swing by and maybe get that last little bit of the
corridor and go back into the CA
Then that puts me into the area where I can go east and then
follow into the corridor for the rest of it
If I run out of time, I don't really see that as important [last
section of air corridor]
Because you're not really confined in that area, and it's not the
critical area
So the most dangerous portion would be the northern critical area
above the turtle's head

Thinking Aloud

Left my speed way up
There we go, kinda got it better under control
I'll just have a narrow weave, this is pretty tight anyways, not
too many places for bad guys to hide
(I remind him about stick speed control)
So this is kinda, getting into
I went a little wide there just in case
Now, I'll slow down a little bit
Get pointed in correct
Whoa, whoa, Brian, aaah, aaah, there we go
I'll turn here, be pointed if any bad guys come out
Not doing a very good job here [referring to weaving]
Now I'll be pointed south in case bad guys are there
No, good
So I missed a lot of the middlewhoa
Alright, there's a SAM in southeast [he meant southwest] corner
Now I'm gonna loop wide and go in with speed
Aaah, shoot, whoa, I'm getting lit up, gotta get out
Ok, so I'm damaged now, right
Ok, well I do really want to take that guy out
I need to loop around, get the guy lined up
(I tell him, "you got to give it a mississippi or so")

And I'm dead, dang't
So I went, yeah okay, lesson learned on...

Case 5

Plan

Very good chance of enemy contact in both areas.
I just want to go as efficiently as possible to the CA, and try to minimize getting killed in the process.
Most efficient in terms of coverage, would be go to go around the large hill to the west and north around it.
But that confines where you can go, plus up north they got a couple good places to hide.
So instead of doing that, I'm going to go the southern route and go south and east around it
I'll probably even duck outside the corridor and go and just center the gap between those two hills to really give myself the best chance of survival.
Then, I'll probably do a north/south weave in CA.
And then, hmm, shoot, my hesitation there is that I want to do a north/south weave in the CA from west to east but that kinda puts me where I could catch that last little bit of the corridor but then I would have to all the way backtrack and go to the unobserved section
So what I might do is go south and east into the CA and then do a east/west search pattern and forget about the little stretch of corridor off to the right there
And then just hit the south around the gap where the bad guys could be hiding
And then just go back through the CA go around north and then west around the rectangular hill
I don't anticipate to live that long...

S1 was reminded, "Now remember you don't have to attack every enemy."
He replied, "That's true, but so if you don't attack the enemy, they're there for your guys right?"
In reply, S1 was told, "You call in the calvary. You're only a one man mission going through, so"
He said, "So you're pretty much not road-clearing, but more"
S1 was told, "It's at your discretion, essentially."
S1 stated, "Sure, well I'm gonna run for the first portion of it at least, definitely until I get the CA done."

Thinking Aloud

I'm going to go through the corridor here, but then I'm going to duck

out and try and put myself in the best chance of survival here.
Ok, now I'm going to try to go in the CA facing the proper direction
Whoa, bad guy.
But he's outside, so I'm gonna
But it looked like a tank, and he had a smaller sphere
Whoa, bad guy, ok.
I am going to try and go after him, but it looks like he's mobile
So I'm going to try and loop around, get him better lined up
This guy's persistent, alright, I've had enough of you
Turn, turn, turn
I'm just trying to get him lined up so I can get a shot off
Oh, not doing a very good job
So I'm gonna get a little distance in between myself and him
Dang't, I can't seem to do it
(sigh of frustration)
[S1 was told, "You're a helicopter, so you could come to a hover.
It's true I could."']
Good call, come get it, bad guy.
Yes, you are done, go home
(laughs) That trash-talking helps, anyway
Now, just resuming the corridor search
(I call out time - two minutes)
I'm going to go over to this next, and do the remaining part of the
corridor
I missed the far eastern corner of the CA
So I'm going to come back in here and do it now
And then, whoa, bad guy
UAV, not good
Ok, let's see if I can go find him again
Hmmm...decision of whether I go chase him
Not that I can see him right now
I'm going to go for the remaining area
Oh, oh, that's guy's on me, geez
Where'd he go
Ok, yeah, I'm not going to leave the CA if that guy's going to stay
in
Come out, where are you?
Don't, oh dang, oh, did I not hit it? That was an act of God.
I threw the throttle down because I saw it at the last second
But the Lord was looking out for me in a big way there, I'm
surprised I'm not rock food
Where'd that UAV go?
Well, since I'm up here, I might as well
Trying to face the correct direction, to see

Case 7

Plan

Possible chance in Air Corr, Slim chance in CA

Possible and slim, is much better than very good that I've been facing

I still want to do the CA more, but I'm not so concerned about dying

Like I said, the more efficient way to investigate this terrain was northwest around turtle

But now that I look at it, you could do that

There's a greater chance of dying in the corridor than in the CA

So now what I'll do is to go north and around turtle to east, that puts me in the CA faster

I'll go in at southwest corner of CA, do east/west weave up to top portion

Go northwest around turtle

Fly back through southern part of CA again

That puts me back in the corridor

But that's a pretty big area, so what I'll do is a diagonal weave

I'll investigate half of it yeah, that's a better plan, so change that

I'll go around the turtle, north and west

Fly along southern part of corridor, but not do any weave

And then loop around and that will put me back in the CA

Thinking Aloud

this section is pretty narrow so I'm not going to do much weaving, try and center myself

If this was the CA I would, but it's not

Now I'm going to go around the turtle to the east

As I've gone along, I'm trying to get better at being able to shoot in the direction that I am going rather than putting myself in a bad position

So there, I want to slew more than I was earlier

Like I was saying, I can do this last little patch and that will put me up towards the north portion of the map and I can go around the turtle to the north

Which I'm doing now

And then I can go diagonally in around his neck right here just in case there's a bad guy hiding there

I missed a portion so I'm just going to loop back

Not hit the turtle
And then do the western half
And okay, this is what I was talking about, where I'll go through
the southern section
And then (I call two minutes)
Now I'll do roughly a division in half of the corridor
Fly it, no weaving
And then, whoa bad guy
I sped up there to get away from him, just because he kinda
surprised me
But now I'll turn around and come to...slow...no, aah, dang,
not a
very good helicopter pilot here
Slow down, slow, slow, slow, slow, (laughs)
Ok, yeah, there we go, much better
Ok, now, I'm going to go after him a bit
Putting the lever a little bit forward just to go after him
He's running for the hills
Don't hit, d\$*!...(groans)"

Case 4

Plan

pretty much same logic on this guy
I'll take extra caution not to run into the hill
Fly along corridor
Go east around turtle and do same weave pattern, I guess
Go north around turtle and south to cover the western turtle
edge
And south, and do the same thing
But it was pretty dumb, I followed the guy outside the corridor
He was out of danger, and I still had stuff that I could cover
So I'll try not to be so dumb, and especially not to run into
the hill, that sucked
There might be some cutoff where I don't follow bad guys
That was the lesson learned on last one
The whole reason for doing that is that slim chance of dying
in CA and possible
So I'm not too scared, but if it was very good I would try to
take the most direct route to CA
But since it's possible I can stay along the corridor and go
east around the turtle

Thinking Aloud

I need to get used to the helicopter being able to hover thing
That was a good tip on throwing the throttle all the way down,
I don't know why I didn't think of it
Whoa, bad guy, ok
Turning to get away from him, there we go, ok
I was in pretty defensive there, because he was on my six, I
didn't really have a shot
I turned away from the CA to get away from him
So now he's in here, so now I deviated
I didn't go north and then south, I didn't care
I was more worried about not dying and finding the guy and
being defensive
Probably benefits me
So I know the guy's probably still in the CA, but I know I'm
going to be coming back to it
I'll turn there just so if anybody's hiding there, I'll have
a shot
(I call out two minutes)
Now I'll go back in the CA
I might stay around here and see if the guy is anywhere I can
find
I wouldn't have done this originally
Alright, I'm gonna go, follow the corridor down
(I call out one minute left)
I'm gonna cut it a little short, I did most of the corridor,
so I'm going to turn back so I can still have time to get back
Whoa, bad guy, ok
Whoa, and he's cookin' on me
So really I'm just trying to stay out of his sphere
Ok, now I'm trying to loop back around so I can go after him
Well, nope, ok, I'll learn from the last encounter not to
follow him
Ok, tank, slowing down, turning, and then
Oh shoot, turn, turn, turn, turn, turn
Running out time, running out of time, turn
There we go, (sigh)
(I tell him to really squeeze the trigger button)
Yeah, well I'm not doing too good of a job in here anyways

Case 6

Plan

slim chance in both

It seems natural thing to go along corridor just to start with

and go north
From the south, through the east to the north, counter-clockwise
around the hill
While I'm doing that I want to investigate the CA
I'll try and stop the investigation to the north so that I can
go directly east and continue the counter-clockwise rotation to
hit the corridor
Well, my hesitation right there
It seems like there's that confined geography where you're coming
into the corridor from the south
There's an increased chance of dying
If it was possible or very good or above slim essentially, I
might go north around the hill clockwise instead
Anything other than slim I would do that
But I'm not too worried about it
The problem with going the other way
If you went clockwise, you would have to shoot up to the far
northeast corner and there's no corridor or CA really there
and then you would have to loop back around
It's not that much of a deviance
But if you went counter-clockwise you could go through the CA
and then come back and shoot up and that would be the last
section you covered
Seems a bit more efficient to do it that way
And I'm not too worried about dying
Like I said, any other scenario I would probably go clockwise

Thinking Aloud

trying to center myself, so that if there is anybody
I'm at least not right up against the hill, maximize
maneuverability
I want to come into this area facing the inlet there, just in
case
Yeah, nobody there, which is good
I want to end up over there so I'll cut the search a little
short
A little short of the northernmost border, because I want to
hit that last
I come into this area facing the proper direction, I'll turn
A little east/west weave here
I probably could have planned this a little bit better
(I call out 2 min 30 sec)
I'll do a little bit of a north/south weave to cover this
section that I haven't

Now I'll go back around the corridor here
I'm kinda investigating the southern portion of it
Because that seems more dangerous
Like right here, looks like a sea lion's nose projection
You could bail out to the north
I'm gonna try and cook through this section because this a
very safe area in terms of geometry
Now I'm back in this area, so I'm going pretty much as fast
as I can because I already did investigate this
I'll just continue up to do the section that I didn't quite
get yet
Ok, I'll just turn around get the last little portion of it
And that will put me back in the CA
I'll go back to what I thought was the most dangerous
The southwest corner here
To see if there's anybody there
Slowing down so I don't smack it though
Whoa, and I'm going to, whoa, nice

C.1.2 S2

Case 3

Plan

I think I'm gonna go clockwise around the circle, so that I hit
the CA first
Because if there's a very good chance of encountering bad guys in
the air corridor
Then I'd like to encounter them later in the flying time so I
don't die in the beginning and not
Accrue any points

Thinking Aloud

I'm gonna fly the air corridor
Whoa, this one's really fast
Go clockwise first
And scan the entire CA first
So that I get the points for doing that
And then go around and look at the rest of the area where I might
encounter the bad guys
Mosy around in the CA for a little bit
Then go around
Is that running into terrain just the dot in the middle or any of
your whole circle there?

It's the white part but not the dotted circle around it?
So I'll be careful over here, (chuckles)
Alright, so I think I've covered all of the CA's area
So I'll go out and fly the corridor
(call out time)
This way and maybe run into some bad guys out here
I think I'll go up and make sure I look at this little tail area
I don't see any so I'll go back
Still don't see any so I'll go in this between area where you said
they might have been last time
And that one is a square and I don't remember what a square is
Oh, ground vehicle, he's probably pretty bad at shooting me
Slow down
I think I'm wasting a lot of time chasing the guy around, but
How do you know when he's already red does it just have to be in
the square when he's like that can I shoot him?
Or does he light up again?
There we go, got him
And I think I'll spend the rest of it in here
Because I kinda wasted a lot of time out here
But I didn't do this outer edge up here when I scanned this thing
Or maybe it was the inner edge, I don't remember
So I don't think I'll do very well on the score for amount of area
covered
Because I didn't do it very systematically this time
Because I got distracted by the bad guy
And I'll zip back around

Case 4

Plan

Initially, I'll fly the air corridor directly to the CA and
search that first
Just so I don't expose myself before I'm able to search that
to enemies
And I think what I learned from last time if you see a SAM or
whatever the long-range shooter is, you should just run away
because they're pretty good at killing you

Thinking Aloud

Wow, that's fast
Going quickly to the CA
Coming into it now
Mosy around in here until I've seen all of it

And if I see a bad guy I'll probably chase it around
Like I did last time but like I said the SAM's seem to be the ones
that are dangerous
Whoa, getting close
I don't know if I'm making very good lines on this one
Getting really close to the edges
And now I've pretty much seen everything here
I'll go down and do this right end of the Air Corr first and I'll
just kinda go out and back and do that loopy part second
Whoa, there's a bad guy, I think I'll go get him
I'm making circles, let's try it the other way, haha
Aaah, bad shot
Yes, got him
So I'll continue going around the air corridor
Just to make sure I've seen all the area and maybe look for that
other bad guy
But I don't see him so I'm going to go back this way and go
around the rest of the Air Corr before time runs out
Probably do the inside part first
Then do the outside part
And no other bad guys here
Going back around
And probably just hang out in the CA and just make sure it's clear
I don't know where that other bad thing went
I don't think he's in here though
I guess I'll go back out here and see the rest of this part that
I didn't see the first time
And take it all the way back to the beginning
Oh, there's a bad guy I wonder if I could get him
Oh man

Case 6

Plan

Fly directly to CA
Scan it slowly as usual
Then go quickly through AirCorr
Since there's a slim chance of enemy contact I don't expect
we'll see anything but I think it's still smart to do the CA
first so I don't die before seeing anything and run away from SAM's

Thinking Aloud

Going fast, whoa, hard to control, I'm gonna slow down because
I can't really do it that fast

Wow, that's kind of hard to control
I think maybe I should slow down or something
Kind of missed
(I remark about lever versus pulling back on stick)
Well, I think I'm getting most of this area
It's kind of hard to control on this one
I'm going to go out here
And since it's going pretty fast and jumpy, I don't think I'll
push it forward much to go faster
I'll just let it cruise
We'll see if there's anything out here
It doesn't seem like there is
There's not
I'm going to take the loop back down this way
See if there's anything in the other direction
(I call out 1:45 left)
That's not much time, guess I'll go a little faster
No bad guys out there
Maybe there are some over this way
And go fast this way see if we encounter anything
I don't think we do
So I'm just going to go hang out in the CA
So here I am patrolling around for nothing
(I ask, "So if this was your first time searching through it [CA],
what would be the most effective way to cover all of it in the
forty seconds you have left?")
What do you mean?
(I reply, "In terms of weaving back and forth")
I think doing straight lines across the area works and longer
paths because it's hard to do the turns
Doing it this way is not as effective as the longer direction
of the triangle because it's easier to get lined up and just...

Case 5

Plan

Since there's a very good chance of enemy contact I might just
avoid the air corridor going to the CA, and scan it first, and
go really fast and not stop to see who I run into
It doesn't really change the rest of plan
Go straight to the CA and then secondarily go through all of
the Air Corr

Thinking Aloud

Going fast
I'll kinda stay on the outside and maybe avoid bad guys
So I'm doing the longer direction because you have to do less
turning
Because it turns around pretty quick like that
Ooh, I guess we can go after that bad guy
Oh, he's got me in his sights
No, it shot late
Aah
It's not shooting well
(Sigh), there's kind of a time delay on the gun right now
I'm out of shots so I think I'll just kind of ignore him
So now that I've covered pretty much all of the area over there
We'll just go around
Go back this way and go back here
See this part of the air corridor and swing around
And go back up through here, whoa, big triangle down there,
which is dangerous
I'll stay away from that
I think I'll come..
Whoa, that guy's got me in his sights
He's a fast one
For the last minute, I think I'll avoid all these bad guys
I think zig-zagging tends to help
Then I'll go back into the CA, which seems a little safer
Except for the SAM down on the bottom
Kinda stinks I only have five shots
They're on to me

Case 1

Plan

Since CA is next to terrain and since there's very good chance of
running into bad guys in the CA, I need to be careful in approaching
it and also reverse the order...so go through as much of the air
corridor first before getting to the CA, so I at least cover that,
just be cautious and look for a way to escape if there's one of the
really dangerous bad guys in there

Thinking Aloud

Wow, it's flying better on this one
Going along and gonna search around the air corridor first
Whoa, that's one of the diamonds, that's another UAV
Don't know how dangerous they are, I think I'll leave him alone

for a minute
Come back to him later
I'm going to go around
Since there's a very good chance of running into bad guys in here
I think I'm gonna work my way down
Opposite of what I've been doing
Now that UAV
I don't know how cautious I need to be around him
It's going very slow
Is he running away nowah, coward
Ok, so I only have two shots and I haven't seen all of the CA yet
I want to finish looking at this part down here
Go around it one more time
Whoa, run away, run away, run away
So I'm not going to go back over there
And actually I'm going to run away from these guys too, because
I'm damaged
Oh man, he's fast
Probably going to have to zig-zag around here
I think I've lost him
Hmm..(I call out 45 sec)
So I never got over to that other side but since I might die if
I go over there because I've already been damaged, I may not
even attempt it
I'll just hang out here
Another 45 sec, it'll be up soon anyway
(Whistles)

Case 7

Plan

So this is the same layout as the last one I just did
But the last one had a very good chance in CA and this one says
slim
So I'm gonna do the CA first instead of second, unlike last time
Try to make paths through it on the long axis

Thinking Aloud

Running pretty quickly on this one
And, oh, just flew through a bad guy
I'll ignore him for a little bit
Because that seems to be un-productive, maybe I'll goof around
later
I'm not very good at shooting

And since I did that one scenario before where I almost kept on running into the terrain
This way tends to be easier too, I can go over the edge and it's not dangerous
Looking at everything
Going back up and now that we've seen all this
Now I'm going to check out the air corridor and maybe find that bad guy
Since it said possible chance of finding or encountering enemies here
Get him on the way back
This one's really wide so I think I'll have to do some little strips because it's a little hard to judge where you have or haven't been
Whoa, he's off the screen, not gonna try it
So now that I'm up here I guess I'll do this outer part of the corridor really quick and come back down
And do the outside now and try and get this guy
Whoa, he got me, no he didn't
Aah
Seems like you have to kinda come out from a ways off or else they just turn away from you
Got him, no, I don't think I'm holding him in sights long enough
Ooh, two shots remaining, gotta get it
Don't know why that's not working
Oh well
And he's off the board, coming back

Case 2

Plan

That's a really wide flight corridor so there's a lot of area to cover there
And there's also no really direct route to the CA
So I guess fly as much of the flight corridor on the way to the CA and then
The CA has a very good chance of enemies, so I think what I might do is a lot of the corridor first
And then hit the CA
The CA, you're kinda boxed in
If you get in over there and there's a SAM, there's no where to go and he's pretty much got you
So I guess I'll do the flight corridor first

Thinking Aloud

Doing this section, go make kind of a border run around the whole thing and then concentrate on these two fat areas
Whoop, that guy's pretty good at shooting me
So I'm going to try and avoid him
Because I'm not very good at shooting
Gotta remember I kind of got away from the board over there
Ooh, not flying so well
Go back this way
Go down
Now this time I'll stay one swath away from the edge
And finish up doing this area down here
Alright and now that that's sort of done, hopefully
I'll take the middle road here going back up
And come up here
Get that section I kinda missed
Kinda do some of the middle of this guy
Probably need to get into the CA because there's not a whole lot of time left
Just be cautious in here, fly slowly as I look around
Oh, there's a bad guy, he's got me in his sights
I think I confused him
And I'll go back and get this edge here
Probably run into him again here somewhere
Hmm...he's gone, look for him
So I think I covered the majority of the CA
So I'm going to come back out and try and get this big fat area here
Which I didn't do very well the first time
And there was another UAV out here somewhere, but I don't know where he went
Hmmsso I think I covered that decently, I guess I could go back in here and look for that bad guy, maybe shoot him in the last second here
Is he around here?

C.1.3 S3

Case 7

Plan

Looking at the corridor, I'm starting at the lower left
I've got an exit point at the upper right
I'll circle around the island at the upper left
Hit the CA
Go up to the upper right, come back

I guess it's like formulating a trajectory of where I want to go
So I can cover as much of that area as I can
I'll probably spend more time in the box
Because there's less of a chance of me getting shot down there
and it's a critical area for the mission
I also feel like sometimes I have an attitude
The first set of go arounds I had an attitude
Where I was making sure I stayed as much in the air corridor as
possible
This time I'll be a little bit more liberal, in terms of missing
the edge of the Air Corr
While still being to hit a complete path through the air corridor
Making sure there's not things there
That it'll be safe for the troops

Thinking Aloud

So it's easier to remember, I think I'm going to follow one part
of the corridor
I want to cover as much area of the corridor as I can
I'll stay up on the left hand side
So I'm coming up on the island
I'll make my left to go up and search more
Ok, so we're just going to stick to the plan
No reason to deviate
Coming up on the CA
So I'll take time to zig-zag to make sure I cover as much area as
I can
Yeah, taking my time
Most of the CA has been scoped out
I've spent more time in here than probably anywhere else
This pass I'm gonna exit and make sure the rest of the corridor is
clear
At least the path of the corridor
There doesn't seem to be any reason why the left, the right, or the
middle of the corridor would be any better to explore so I'm just
doing something to remember where I was
OK, that was another UAV
I saw the direction he took off in
He didn't shoot me and I didn't shoot him
I'll look for him as I come back around this way
He's usually faster than I am if he's running away I won't catch him
There he is
Now I've got him
Circling behind him

And he's off the map
Got a slight glimpse of him above
I've mapped the entire route
So at this point if there's an opportunity to get an enemy, maybe
I'll take it
He's too fast for me, I'll move on
A tank, uh-oh, gotta get outside of his area
Circle around, come back in on the backside
I have a larger area than him so I'll have a pretty good shot

Case 4

Plan

The last time I took the long way around to get there [CA] initially,
maybe it will be better to have a shorter path to get through the
air corr
So instead of going around the island
I'll go directly straight up, go ahead and clear the entire air
corridor all the way to the top
Then on my way back spend time in the CA
And then if I have time, I'd like to stop in the CA going back around
Come back down

Thinking Aloud

I know I'm going to make a right hand turn up here
So I'll stay to the right side of the corridor
Be on guard
Early in the game, I got time to go after him
But I missed him and I have three shots remaining
He's a tank on the run, I'm gonna save my ammo
Yup, jumped the gun a little bit on that
Need to spend more time in target practice
There's the UAV
Gotta get outside his cone and then come back
He's too fast for me
I'm not going to bother with that too much
Look for the tank
Slow down a little bit
This is where I saw him come up
Covering the CA and now for the upper pass
There is a tank
Coming in for a shot, oops, get in range
Alright, now I'm out of ammo
Any chance I get to run, I will run

Whoa, looks like I hit someone's
Someone's target targeting me for a second
And here comes somebody I'd rather not face
On the move
I covered this area, so I was maneuvering to avoid whatever it
was that I ran into last time
And there's a slim chance there's anyone in here, so I'll just
spend the rest of my time in here

Case 1

Plan

I'm gonna stick to the same plan I had last time, staying to the
right
That will put me in the CA for a short amount of time before I
cover at least a path through the corridor
And then when I come back, I'll scope out the CA
Try not to waste any shots until I get there and have explored
that area
And then move on

Thinking Aloud

I think it's getting faster every time
Oh, I'm in someone's...ah, got hit
I know there's a SAM site there
Whoa, there's another one
Yeah, I'm toast

Case 3

Plan

There's a very good chance of enemy contacts within the corridor
and possible chance in the CA
So I'm going to scoop around to the left of the corridor
Spend some time in the CA first where there's less chance of
running into enemies
Once I've scoped that out and cleared it, then I'll move on
So if I do die, then at least I've hit the CA
Since I died last time...stinkin' SAM sites

Thinking Aloud

Sticking to the plan so far
I can get three passes, so I'll do that
This one vertical pass
Second vertical pass
(not sure what he says in here)
OK, scoped out the CA, moving on to the rest of the corridor
Moving up to the left, so I do have a path through the corridor
that will work
Now I'm going to go basically straight back to where I started
If I have more time maybe I'll just cover some more area
(not sure what he says here)
There's one
Whoa, I'm gonna need to slow down
Alright
Ok, one shot to go, means I can't kill him, maybe I can wound him
There was one more tank over here
I'm not too afraid of tanks
Maybe there was one
I thought there might have been another
Now I feel like I've covered most of the air corridor
Sweeping it one more time to cover as much area as I can
And then I'll be out of time
Did a lot of the air corridor...

Case 5

Plan

There's a very good chance of enemies everywhere, so I'm pretty much going to stick to the plan that I've done before where I'll carve out the shortest path through the air corridor, brushing through the CA, then come back and spend some time in the CA, and take the longer path home, that's my plan

S3 is asked, "Do you think a very good chance of enemies would necessarily affect how you would particularly search through the corridor? Are you more just more attentive when you go through? Would you change speeds or how weave or anything?"

S3 replies, "Not until I think I either see them or I know they see me. At that point, I'll either try to avoid or engage on whether...if it's a SAM site, I'm out of there...if it's a tank, they're pretty easy targets, so I'd probably circle back and get it...if it's a UAV, if I have a good shot at it and I'm close enough that I can catch it before it flies away, I'd probably take that

opportunity, but if I'm not in a good position as I go by the UAV,
I'll just keep on moving straight through"

Thinking Aloud

I think also the whole following the path kind of goes out the
window at that point too

My sole focus would be on whoever it is and either engaging or
moving away

Ok, so I've got a tank

One shot, he had me in his sights for a second

Out of there

Ok, I'll make sure to avoid that corner next time through

I'll bump into the CA here

Steering clear of the SAM site

Alright, UAV, he was behind me, (??) I saw him

Keeping an eye out otherwise I'll be out of business (??)

UAV, circle around back

Didn't see him

Ok, come back to my path

Check out this top part

It's kind of a narrow area, if there's a SAM site I don't have
anywhere to run, I'll scoot through

Ok, I've scoped out everything

Ok

Whoa, got someone behind me

I'll circle around

See if I can't get behind him

He disappeared

Ok, remember the SAM site's up there

I was going to ask how much time I had I might have gone for it

Case 6

Plan

I'm still satisfied with the way things are going, in terms of
going straight through spending time in the CA and taking the
long road home

There's a slim chance I'll see enemy contacts anywhere

So I'll go on my merry way exploring

Not even gonna have any kind of adjustment

The last time I wanted to fly faster through the narrow area
in case there was a SAM site

I probably won't worry about that this time

Thinking Aloud

It was definitely an area of less maneuverability last time, I felt more vulnerable going through, but not this time
Sticking to the plan so far
Yeah, 25 deg of bank
And the intelligence makes the difference too, when I started flying this mission, I had my hand at the bottom of the joystick and then I realized, well you should always be ready
Whereas last time I started off with an attitude of being more ready
So I've seen a good portion of CA
I have a path from start to finish
I'm going to go the long road around
Just kind of on autopilot in terms of executing the plan, not thinking about a whole lot
There would be though, I think if I was in a real mission, thoughts or aspects of where would an enemy be, would he be there hiding behind that terrain
Should I go out and further around?
To avoid him first before I have a good view of it
1 minute left?
30 seconds
And now I'll just carve out the other half of the area I missed inside this
And I should be running out of time pretty soon

Case 2

Plan

In terms of feeling constrained, I know I probably wouldn't have any enemies coming from the left, which would be off the screen, so if I start on that side, if I see an enemy I can turn, because he'll likely be to the right of me or ahead of me...So that would be reason for starting at the upper right and scooting down the right hand side first I'll probably exit the CA, since the probability decreases out there, turn around with my sights forward right next to the area I've already cleared, so I continue to have a better idea of where I think the enemies will be coming from

S3 is asked, "Now these big lobes of air corr sections, any idea of how best it would be to cover that area?"

S3 replies, "I was thinking my plan might be totally different this time, in terms of covering the shortest path and making the trip back, the CA doesn't coincide with a trip back, maybe I would

actually spend some time wandering, and plan to leave myself a minute or so to spend in the CA near the end, because that will be at the end of my flight path through there, that way I don't have to go back over territory, besides that there's a very good chance of enemies in CA, so I want to put that last, I want to be able to cover the path through the corridor, make sure the troops can get clear, get through there, but since I'm a UAV (??) really accomplish at least one objective of the mission before running into enemies"

Thinking Aloud

To cover the most area in here I need to zig-zag pattern
I'm moving pretty quick right now
It may help me cover more area
Also makes it a little harder to control
If I'm moving at this pace I may not have a problem
Whoa, SAM, I'll get out of there
Ok
It's always good to have identified where those places are at
I feel like the only real way to attack it would be to move in hot and fast on it
Stop right when it's in range, fire off two shots in rapid succession
So knowing where it's at is pretty key
So I'll keep that in mind and I have time in the end and I've explored everything I need to it might be helpful for our guys to have a SAM site taken care of, but because it's not something I'm very good at doing, I may have somebody else do it
So I'll make one more quick pass through here
Come to the end, execute my plan in the CA
Ah, trigger happy there
Little herky, jerky so, I'm a little all over the place because it's moving so fast
And there he is
Ok, UAV, he's tailing me, ooh
Circle out and around
Come back at him in this direction
Whoa
Ok, behind him, maneuvering
This is really tricky
Aah, two shots gone
Ok, he's left on his way, at least as far as I can tell, maybe he's coming back, but I am moving on
I'm gonna try my strategy I think for the SAM site
If he's still there

I feel like it might be a good plan
Get lined up on him
Got finger on trigger and go
(?? Something about not tight enough)

C.1.4 S4

Case 7

Plan

Slim chance of enemy aircraft within CA
So in order to diminish the likelihood of being struck by an enemy
and completing the entire task as planned
Go to CA first and take the right branch of the corridor
Cover the entire CA
After which we'll continue with the corridor which goes to the far
right of the viewing area
And if we have time, go back to visit the loop in the corridor

Thinking Aloud

Sticking to the right of the corridor here, trying to make things
a little easier for me to remember
Coming up to the split
Head to the right like I said
Here's the CA
Sort of sweep through it
Trying not to miss anything which is kinda hard
Not sure that this is helping, but alright
It's better if I slow down before I hit the sides
And I definitely missed a spot there, but it's ok, cause heading
back, and that's probably it right there
Continue on and hit corridor down here
It's pretty big, so I might just hang around here the rest of
the time
Oops, there's an enemy
Looks like its moving, so I'm going to just ignore it for now
Let's see, how are we going to do this
And, going for final sweep
Looks like that's a stationary object there
That was easy, ok
Finish up by hitting this little corner that I missed
Alright, well, that's about all I have

Case 1

Plan

Same map as we used from last time except now there's a possible chance of enemies in the CA

From experience of the last scenario...we were short on time and couldn't finish the loop area and go through the entire corridor, but just visiting the corridor on the right side of the screen
This time I think we can fit it in by trying to go first for the loop

And then going toward the CA

If you have time after the CA, going for the corridor on the right side of the screen

See if that works

Thinking Aloud

Alright, I'm not controlling very well this time

It might just be me

So I'm gonna go ahead and go the left this time

What I'm gonna do is go to the right hand side

Do another sweep of the loop so that I can get the whole thing

And come back underneath here

Whoa, alright, so that looks like an enemy

And it looks like a stationary one which means I don't have to worry about hitting it right now

I would to actually get the CA done

Since I'm going way fast

Oh no...did I just crash? Sweet...oh...did I just...no way

Case 2

Plan

Very good chance inside CA and CA happens to be at end of corridor anyways

So I was thinking to try and get the majority of the corridor and then when there's a minute left or so head for the CA and see what happens...maybe a little over a minute

In the event that I die, I would have at least gotten much of the corridor covered

Thinking Aloud

Let's see if I don't crash into something this time

This weird region up here, I'll try and cover it and then head

back around
(I remind him about how to slow down)
Oh, pulling back slows down, ok
Okie-dokie
(I call out time)
Alright, one minute, looks pretty good
Before I take this section, I'd like to do a quick(groans)
Alright, that works
Quick sweep of this area in the corridor
And do a left turn here
Get the center that I missed
(loud sigh)and now this area
I'll start at the top actually
Ok, and right turn
And that's way too far down
(I call out 1:45 left)
Ok, let's just do one more sweep after this one
(sigh)...love those turn radius
Yeah, well, I'm not going to get to the rest of this, which is
fine
I definitely want to hit the CA
Get ready for some bogeys
Ok, now I don't want to hit the side, I'm guessing
That's much better, except for that
Whoa, get out of there, whoa buddy, whoa buddy, he's gunning
for me
Yeah, he's definitely targeting me
Whoa, what just happened
Well I got hit first of all
Ok, well, not successful
Well, that's ok

Case 5

Plan

Similar overview or layout as I've seen before, except now there
is a possible chance of enemies both inside and outside
Knowing that, it would make sense to first do the loop like last
time and then go for the CA and then see what's left
Because of that wall there, and the fact that I ran into the wall
last time
I'd like to in my search pattern instead go alongside the wall
parallel to the wall
So that I have less chance of running into it in my search of

the CA

And do that one row parallel to the wall and continue on with the search

Thinking Aloud

Off we go

While I'll never get around to it, I'm just going to stick one side, one portion of this corridor, just because

Whoa, whoa, whoa (laughs)

Wow, so that wall is actually sticking out in several spots, which is not friendly

And the plan, hey, oh goodness

Of course, I could be sneaky and crash as quickly as possible so I don't have to stay here and fly Jon's dumb sim (laughs)

Ok, that's not good

Alright, looks like I (??)

Oh gosh, maybe you just shouldn't have made it as jumpy as it is

Either that or I'm not a helicopter pilot

Whatever this has to do with helicopters

Alright, so here's the plan

(I call out a little over two minutes)

Plenty of time, plenty of time

Stay as close to the wall as possible

There is a bit of variability going on with the throttle

It looks like I'm going faster if it's all the way up

Right? Or am I just imagining that?

No...I guess not

(I call out 1:20)

Oh, what the heck, alright I'll be back for

Does that mean I didn't get hit?

Cause I was red for a little bit?

Let's see what we got here

Come on shoot, is it really not shooting, shoot...oh, zero shots remaining, what happened?

What's going on here, I didn't shoot...

Case 3

Plan

Very strong possibility of seeing enemies outside of the CA and I'd like to get the CA done and survive through as much of the scenario as possible

So I'll go directly to the CA using the left fork

Because it's a little odd-shaped there, what I'll do is use an

up and down pattern for that rectangular region and then take that small sub-rectangular region afterwards and then continue on with the rest of the corridor if I get the chance to

Thinking Aloud

Alright, stick to the left-hand side here
Up and down pattern I was talking about
Seeing as much of the corridor (??) as possible
And one more sweep
Actually, I'll go back to what I did and what I'll now do
Come down here since it's such a small area
And do that and whoa
One more sweep
And I think that got it all
Alright, on now
What I'll do is travel the loop until I've hit where I've been
and then come around
Stay centered
Ok, nice...I want to stay out of the range the whole time
I'm gonna do this section here...whoa
Thirty seconds, I might as well go to where I know I definitely
haven't been
Stick to that side
(Making noises)

Case 4

Plan

Possible chance of enemies inside corridor and slim inside CA
I'm willing to take the chance of being in the possible air
corridor and see enemies just for the sake of efficiency
What worked well was doing the top of the loop and then coming
down into the CA and finishing that
And then doing the right hand side of the corridor
Because of the big obstacle on the right side and the one on the
left side of the CA, I think what I'd do is the up and down pattern
in the search to decrease my chances of hitting the wall

Thinking Aloud

Stick to the top here
I think I'm getting better at this, oops discretion
Now I'll start this...goodness, this thing's crazy, it's got a mind
of its own

Ok, so this is the kind of search pattern I want to do
Of course, the fact that's it at the top screen is probably similar
to being an obstacle...but that's ok
I'll just be careful
(I tell him he could go off the screen, he just has to find his
way back)
Oh...got it.
You could kinda use those corridors in there as reference points,
I suppose
Actually, I think I might want to do one more run
Ok, well I feel like I've gotten most of that
(I call out two minutes)
Two minute warning
Ok, so I'll go here and get ready in case there's any enemies
Whoa...ok that's one shotfour shots remaining, in case we see him
again
Yeah, well, let's try and go for him this time..ok
Ok, well I'll do the center, that's what I missed because of the
disturbance
I've already covered this, so I'm going to try and speed through
it so I can cover the other section
Looks like this is all I need in here

Case 6

Plan

Alright, so there's a slim chance of enemies in all areas in
which case I should go about it without changing my strategy
because of the enemies
What I'll do is go up the top fork, around the top of the object,
cover the corridor
At which point I will do the CA, finish that
And continue to the bottom portion of the corridor
And whatever I can get on the top part that I missed

Thinking Aloud

Alright well I'll stick to the left side here
Whoa...yeah, not doing very good am I
Ah, it's easier when you have it up here
Ok, so I got that top part, now I want to hit the CA
Wow, that's bizarre
Dude, this thing's going wicked fast
Alright, well, that whole wall thing is kind of here
So maybe I'll modify my search and do this

These are remarkably closer to each other, than...
So I think after this one I'll be pretty much done
Alright, that's about good
Can I (I call out two minutes) remember...
Slow down here for a second, get my bearings
You know, the corridor gets a little narrower here, (??)
Do the center here
(??)
I still have to do the top portion
(I call out one minute left)
(Big sigh)
Pretty much got that whole thing here and so I'll figure what
I do is just kill myself
No scratch that...see if I missed anything here...it's fairly
invariable
What really would have been nice if there had been a trail left
behind to see what you've already seen
(I reply by telling him that'd be cheating)
No it wouldn't

C.1.5 S5

Case 1

Plan

Looking at the terrain here
I have a circular corridor around some terrain
Then there's the CA
I know I have a possible enemy contact in the corridor
So I can expedite somewhat through the corridor with a very good
chance of there being an enemy in the CA
My initial plan will be to focus rapidly on as much corridor as
possible to accomplish that mission
Move slowly through the CA providing myself exit lanes and be on
the lookout for the enemy
And then focus the remainder of my time in that CA
And then only after I'm comfortable with the CA I'll continue
with the last third of the corridor on the upper right hand side

Thinking Aloud

As per my pre-mission plan, I'm going through the corridor
Enemy contact is possible so I'm being somewhat guarded
Lost my wingman, so I'm being careful since I don't have a wingman
Approaching the juncture here, critical terrain

I'll go to the left real quick, analyze that portion of the terrain
Actually I'm going to come into the right and do an initial glance
into the CA here
That turn didn't seem realistic
Ok, oops I got engaged
Ok, so I know there's somebody potentially in the bottom left
corner of the CA
I'll come back in from the top knowing that there is a potential
for a threat in the bottom corner
Which doesn't necessarily mean the only place, but
Now I'm approaching the CA, I'm gonna move slow, gonna do the top
search here, move slow through here
Give myself an exit, right now I...ooh, crap, that's some tough
terrain right there
And back down, you put me right up against the terrain...that's not
fair
Ok, now I'm searching the top portion, so I'm working my way south
Since I think there's enemies in the south, by searching from
bottom...er, from top to bottom
I know that if I get in contact with the enemy, I can move up to
the top and get somewhat safe
I don't like the terrain here at the edge of the perimeter, so
I'm going to search that
Ooh, crap, I almost hit it right there
And move back up, give myself a little bit of buffer area next
to that terrain
Continuing to search the CA, trying to go slow
Get a good, detailed search
Moving in almost like a movement to contact type fashion because
I know there's some enemy at least somewhere down there
Where exactly he is, I don't quite know, but I know that I'm
getting close to him
Since I did get contact somewhere in this vicinity, I'm going to
come to a hover and milk my way over this direction, give myself
a higher chance of
Ok, I got engaged, I know where the enemy's at now
Come back around
I know where he's at, I know that he's a SAM...oops, sh*\$...ok
there's two SAMs in there, one in each corner
Cut here through the middle of them
Come back around to the top, see if I can get an engagement here
And do some running fire on this guy, I'm reporting him to higher
I'm going to sweep by as fast as I can on him
I don't think that's working right real well so far as targeting
Aah, this sucks

Case 3

Plan

Ok, because I know there's a very good chance in the corridor, I'm going to move more gingerly in the corridor and since it's possible in the CA, I'll move more quickly through the CA trying to examine all of the CA, but I'll be more careful than ginger. I'm actually going to move to the CA first, staying outside the corridors and then move my way through the corridors. That way I can accomplish the CA first and then do the higher risk mission second

Thinking Aloud

I'm in the CA now, moving at a somewhat quick pace but not blazing through there

I want to make sure I cover the whole area and provide good intelligence to what's inside the CA

Make sure I cover every square inch of it with my sensors

I also know that this portion I'm in the corridor also, the reports for inside the corridor so I'm being cognizant of a potential enemy contact and moving at an average rate of speed

And continuing to work through the CA

I'll do one more upsweep through the CA, then I'll move to the right hand side of the corridor and work my way around the terrain feature

Knowing that enemy contact is very likely, very good

And I'm done, I'm comfortable with my analysis of the CA

Since I failed to do any of the corridor, I'll work my way counter-clockwise around the center feature

(I call out a little under two minutes)

Halfway through, so I gotta speed up a bit, wasted too much time in the CA

So it doesn't look like I'll be able to accomplish all of my mission objectives

I'll move through the center of the corridor where my friendly units are more likely to travel

Keeping in mind a tactical exit strategy in the event of contact so as not to hit the terrain

You're channelizing me a lot in these corridors

Focus on this junction here, get a good sweep of the terrain

Approaching my s.p. [starting point], still no enemy contact

I'll call that portion of the corridor clear, as much as possible

I'll do a more deliberate search of the corridor leading up to the CA because this seems like a highly probable route of travel for us to move to the CA

I'm making an assumption that the paths are designed for other aviation assets to come up later
I'm trying to scan both sides of the corridor
Trying to get as much overlap so any enemy that might be in a position to affect the area along the corridor but might be outside the corridor
And there doesn't look like there's any enemy contact anywhere

Case 4

Plan

I have the circular route around the corridor where there's a possible enemy contact
I'll initially want to move towards the CA
Because I don't have a lot of time and there's a lot of area to cover
I'll choose to go up the corridor this time vs. directly to the CA
Being cognizant of inside the CA there's in the northeast side a large mass of terrain which minimizes my maneuverability so as I scan that area, I'm going to want to scan that area from the south moving north this time
Primarily to ensure that I have a turn point where I'm not boxed in by the enemy or I get surprised

Thinking Aloud

Whoa, ok, and I'm moving through the corridor
It's a lot touchier this time for some reason
I know there's possible enemy contact in the corridor but I don't have a lot of time so I have to move somewhat expeditiously and efficiently through the corridor
Approaching the corridor I'm going to choose to go right up through the CA this time
Then around the terrain feature
I'll check the southern portion of the CA now
There's a slim chance of enemy contact so I can be slightly less cautious
Now as I approach this terrain feature, I'm going to move slowly up along it
In the event I have contact, I can move south because I have limited maneuverability to the right side
With a successful pass through here I'm creating a larger buffer zone of maneuverability for myself
Knowing also that the enemy is highly likely to be up on that high terrain with air defense weapons to try and use them against me

Moving back through CA moving expeditiously because I know that intel reports are never wrong
Hence there's a slim chance of enemy contact in the area
And it appears on face value that the enemy information was correct that there is no enemy activity inside the CA
So I'll focus now on this portion of the course, of the corridor that we're trying to cover
Going faster than I would normally want to given the chance of enemy contact
But I have a large amount of terrain that I have to cover to meet the mission goals
I'm going to focus on this junction here because it's a critical point in the path
Where a lot of aircraft might be operating
Passing back through the CA and into the corridor
The corridor has widened up a bit so I'll do a little more back and forth
Since obviously my sensor doesn't range the corridor width
Trying to get some overlap
Trying to get as much coverage of the sector as I can
But also in a time efficient manner
Whoa, ok, I have enemy contact
He's moving, I'm gonna finish up the path, the corridor up in here
Report the presence of an enemy tank at this time, moving back towards the tank
This is where I got contact with the tank, it looked like he was moving towards the northwest
There's the tank, I don't know if that's the same tank or a different tank so I'll just gingerly move in his direction
That's a fast f*#\$in tank
Chasing him for a stupid reason
I'm suckered in, I really want the kill so I can get my air medals
Lost myself, I'm in the clouds, still following him
Still chasing him, even though he's outside the corridor, outside the CA, so it makes no sense

Case 7

Plan

Again this time I'm going to move through the corridor in a somewhat time-efficient manner
Because there's a possible enemy contact
I'll examine the CA a little more quickly because there's a slim chance of enemy contact

So I'll focus a lot of the required time in the CA but a majority of time where the threat is
Done

Thinking Aloud

Searching through the corridor now
Move back to the CA
Examine this portion of the corridor so I don't get surprised by the bad boys
Now I'm inside the CA
Being cognizant I don't have much maneuver room because of this terrain feature
Just continuing to search the CA
Focused more on the CA because it's critical by its definition and spelling
And moving back up north
And one more sweep south
And since I'm down here, I'll just go this direction
No real explanation of why, just because
Corridor's wide there's a lot of area in there
Looking for hiding enemy that might need to be reported on or a threat to a follow-on mission
Ok, I have enemy contact there
Destruction of enemy tank, I'll report, at this point it's too early and I don't want to risk getting killed and not accomplishing my primary mission, so I'll finish up this sector
I won't fall in the trap this time of chasing him down
Ooh, crap, that's a SAM site there...that's of crucial intelligence, he's chasing me
Which I didn't know happened
Now I'm running, I'm scared, got some tears
Well, I'm injured, not going to choose get in a fight with these...ooh, sh#&
I'll move back up, report the enemy contact
There's another air defense weapon, I'll report that location
He's moving into the CA so I'll follow him into the CA and report to my higher headquarters that he's fast as crap
There he is there
I'll chase him, get as much reporting on him as I can
But choosing not to get engaged and risk losing my own life
And not accomplishing the range of missions

Case 5

Plan

There's a very good chance of enemy contact throughout the corridor and CA so I have to use efficient maneuvers both in the corridor and CA

I'll move through the corridor directly to the CA focusing a lot of my attention to the CA and then move counter-clockwise around the terrain feature

And back through the CA to finish up the remainder of the corridor

Thinking Aloud

I'm moving through the corridor, I'm being somewhat...I'm trying not to be reckless with my maneuver because I know there's a good chance of enemy contact, and I'm scared

I have no wingman, I'm out there by myself

Searching the corridor area as much as I can

Moving up towards the CA where reports of enemy contact are high

Search the corridor somewhat more efficiently

Being deliberate based off the likelihood of enemy contact

Approaching the CA

I'm more aware of my surroundings

Enemy tank...he engaged me...I'm choosing to bypass and report based on that I'm not far along in my mission and its not beneficial for the mission for me to sit there and get in a gun fight with the guy

Since my mission is reconnaissance

Moving through the CA, trying to get this terrain feature so it's not such an obstacle against me

There's the enemy tank, I'll see if I can get a quick shot off on him

Well, he's dead and I have no more bullets

It took me five shots to kill him, so now my choices about getting into a gun fight later on are pretty limited

So now I'm in survival mode since I'm winchester

I'm gonna focus on the CA as much as I can

Knowing there's some enemy contact

He probably called his buddies and they know that I'm here

Being very deliberate of my maneuver through the box

Trying to stay adjacent to areas that I've already scanned so that I have an exit area to move to in the event of further enemy contact

There's the dead enemy tank, there's some bodies laying around

Keeping the northern sector of the CA as my examined area so that in the event of contact I'll turn up into that area

Passing by the tank again...what's up dudes

Now that looks like it's it for the CA, I'll finish up a little more of the corridor with my remaining time

Oops, there is a SAM site, I have no more weapons so I'll report on his position
Where there's one there may be more, so I'll try and do a somewhat ginger scan in this area
Right where the corridors are entering and exiting the CA
I'll scan over here staying outside of his weapons range
Again looking at the corridor and where it intersects with the CA

Case 2

Plan

There's a very good chance for enemy contact in the CA and since it's at the end of the corridor I'll do a hasty search in the corridor and since there's a very good chance in the CA when I approach that area last, I will search it much more deliberately being cognizant of the large western terrain in the CA and how that can limit my maneuverability...I take it you did that on purpose?

Thinking Aloud

I'll search this large corridor area, which was obviously developed by a freakin' flight school student
Somebody's who has never developed a corridor before in their life, that is pretty obvious
The corridor is 27 km wide, but hey, whatever suits your boat
This isn't a corridor, it's an area, a freakin' flight area, but hey
Now it's starting to turn into a corridor here
Being more deliberate here, because it's a little smaller and reasonable
Have search boundaries on, pretty rough terrain here, doesn't lend itself to hanging out with the enemy and getting shot at here
Now I'm back up in Wyoming here, the Wyoming corridor
Twice the size of Wyoming, searching the area as best as I can
Turning fuel into noise at this point, no real enemy contact nor do I have a high belief that there's enemies in the area, but you never know
Getting ready to set up my approach into the CA
I'll stay away from terrain this time initially
Slow down a little bit and scan the eastern portion first
Whoa, sh##, I turned totally the wrong way
Mistake of mine, I turned into the enemy, rather than away from him
There's enemy contact already in the CA, so I'll go a little bit slower inside that area

I'm mad at myself now for bad reaction
Move back up through the CA here
Working myself towards the terrain feature
So in the event, when I move towards the terrain feature, if I
have enemy contact, I can move to the east given the situation
I have reported the enemy contact to my higher-ups at headquarters
He was a UAV or flying something...
Approaching the terrain features, so I'm more cognizant of my
maneuver desires
I know that I can go to the left in the event of contact in the
area, but not to the right
Corner here, whoo, ok, have enemy contact
Looks like the same enemy, UAV
Approaching now back up into here
Plucked him, plucked him dead
I'll continue the reconnaissance up the terrain feature, again
knowing that it forces me to the right hand side

Case 6

Plan

On this scenario here, I have a circular corridor around terrain
feature and a CA towards the end of the route corridors and
where they intersect
There's a slim chance of enemy contact
So I'll be more hasty in my movements and focus more of my
time in the CA where the enemy doesn't appear to be a factor
in this entire mission

Thinking Aloud

Since the location of the CA is towards the end of the corridor,
I'll move up towards the CA at a fast rate on one side of the
corridor, turn around and work myself back along the other corridor
and approach the CA from the north to the south Do a quick southern
sector scan, do an initial look-through into the CA, get some points
Thinking about the CA Again, moving as fast as I can through the
corridor at this point knowing I already covered the other half or
most of the other half Since I'm not going to come back through this
corridor, I'll do some center weaves Randomly search the corridor
but also conserve some time I can spend inside the CA Sort of doing
a hasty reconnaissance here of the corridor No enemy detection thus
far So I have three minutes of fuel remaining, while I'm still
flying around, I don't know...brave And approaching the CA where

I'll do a very deliberate search Go from one corner to the next and then move back based off of its shape I know I've already covered the southern portion Still reconning the CA, trying to cover every square inch of it as I can And I'll still just searching the CA, weaving back and forth along it I'm running out of time, so I'll finish up the little bit of corridor which is remaining right here I'm confident there is no enemy activity inside the CA Or anywhere on this whole map, at least that I've found, at least And back through the CA By nature I'll focus more on the CA and where the corridors intersect, because that's where my friendly forces will be, coming up to intersect up in the CA Now I'll go back in these woods, search these mountains, looking for Al Qaeda caves

Bibliography

- [1] Lisanne Bainbridge. Ironies of Automation. *Automatica*, 19:775–779, 1983.
- [2] F.C. Bartlett. *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press, Cambridge, 1932.
- [3] Ladislau Bölöni and Dan C. Marinescu. A Component Agent Model - From Theory to Implementation. In *Proceedings of the Second International Symposium From Agent Theory to Agent Implementation*, pages 633–639. Austrian Society of Cybernetic Studies, 2000.
- [4] M. Ted Boren and Judith Ramey. Thinking Aloud: Reconciling Theory and Practice. *IEEE Transactions on Professional Communication*, 43(3):261–278, September 2000.
- [5] Mark Bowden. *Black Hawk Down: A Story of Modern War*. Atlantic Monthly Press, New York, N.Y., 1999.
- [6] Michael S. Branicky, Vivek S. Borkar, and Sanjoy K. Mitter. A Unified Framework for Hybrid Control: Model and Optimal Control Theory. *IEEE Transactions on Automatic Control*, 43(1):31–45, January 1998.
- [7] Michael E. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.
- [8] Kevin Burns. Mental Models and Normal Errors. In *New Directions in Research on Decision Making*, Berndt Brehmer, editor, chapter 2, pages 15–21. Elsevier Science Pub. Co., New York, July 1986.
- [9] John S. Canning. A Definitive Work on Factors Impacting the Arming of Unmanned Vehicles. Technical Report NSWCCD/TR-05/36, Dahlgren Division Naval Surface Warfare Center, Dahlgren, VA, May 2005.
- [10] John S. Canning, O. Thomas Holland, USMC Major G.W. Riggs, and Carolyn J. Blakelock. A Concept for the Operation of Armed Autonomous Systems on the Battlefield. In *Proceedings of the 2004 Association for Unmanned Vehicle Systems International North American Symposium*, Anaheim, CA, 2004.

- [11] A. Chapanis, F.C. Frick, W.R. Garner, J.W. Gebhard, W.F. Grether, R.H. Hennehan, W.E. Kappauf, E.B. Newman, and Jr. A.C. William. Human Engineering for an Effective Air-Navigation and Traffic-Control System. Technical report, Ohio State University Research Foundation, Washington, D.C., March 1951. A report prepared for the Air Navigation Development Board by the Ohio State University Research Foundation under the auspices of the NRC Committee on Aviation Psychology.
- [12] Aaron Cohen. Interfaces. Lecture in Aircraft Systems Engineering Design class (16.885J) at the Massachusetts Institute of Technology, Fall 2005.
- [13] Presidential Commission. Report of the PRESIDENTIAL COMMISSION on the Space Shuttle Challenger Accident. Technical report, United States Government, June 1986.
- [14] Leonardo Daga. RT Blockset. <http://digilander.libero.it/LeoDaga>, 2004.
- [15] Robyn M. Dawes. *Rational Choice in an Uncertain World*. Harcourt Brace Jovanovich, San Diego, CA, 1988.
- [16] Robyn M. Dawes. A Message from Psychologists to Economists: Mere Predictability Doesn't Matter Like it Should (Without a Good Story Appended to It). *Journal of Economic Behavior and Organization*, 39:29–40, 1999.
- [17] David W. Dorsey and Michael D. Covert. Mathematical Modeling of Decision Making: A Soft and Fuzzy Approach to Capturing Hard Decisions. *Human Factors*, 45(1):117–134, Spring 2003.
- [18] K. Anders Ericsson. Superior Decision Making as an Integral Quality of Expert Performance: Insights Into the Mediating Mechanisms and Their Acquisition Through Deliberate Practice. In *New Directions in Research on Decision Making*, Berndt Brehmer, editor, chapter 10, pages 135–164. Elsevier Science Pub. Co., New York, July 1986.
- [19] K. Anders Ericsson and W. Kintsch. Long-term Working Memory. *Psychological Review*, 102(2):211–245, April 1995.
- [20] K. Anders Ericsson and Herbert A. Simon. *Protocol Analysis: Verbal Reports as Data*. MIT Press, Cambridge, MA, 2 edition, 1993.
- [21] Xiacong Fan, Shuang Sun, Michale McNeese, and John Yen. Extending the Recognition-Primed Decision Model to Support Human-Agent Collaboration. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi Agent Systems*, pages 945–952, July 2005.
- [22] Hans K.G. Fernlund and Avelino J. Gonzalez. Using GP to Model Contextual Human Behavior. In *Proceedings of Genetic and Evolutionary Computation Conference*, June 2004.

- [23] Paul M. Fitts. Functions of Man in Complex Systems. *Aerospace Engineering*, 21(1):34–39, January 1962.
- [24] Laura M. Forest. Effective Design of Planning Systems to Support Human-Machine Collaborative Decision Making. Submitted for publication in 2006 Proceedings for the Human Factors and Ergonomics Conference, San Francisco, CA.
- [25] David H. Freedman. *Corps Business: The 30 Management Principles of the U.S. Marines*. HarperCollins, New York, N.Y., 2000.
- [26] V. Gavrillets, Bernard Mettler, and Eric Feron. Human-Inspired Control Logic for Automated Maneuvering of Miniature Helicopter. *Journal of Guidance, Control, and Dynamics*, 27(5):752–759, September 2004.
- [27] Michael Georgeff, Barney Pell, Martha Pollack, Milind Tambe, and Michael Wooldridge. The Belief-Desire-Intention Model of Agency. In *Proceedings of the 5th International Workshop on Intelligent Agents V: Agent Theories, Architectures, and Languages*, Jörg Müller, Munindar P. Singh, and Anand S. Rao, editors. Springer-Verlag, 1999.
- [28] Avelino J. Gonzalez and David L. Gross. Learning Tactics from a Sports Game-Based Simulation. *International Journal in Computer Simulation*, 5(2):127–148, 1995.
- [29] Thomas R. Gordon, Michael D. Covert, Dawn L. Riddle, Donald E. Miles, Kimberly A. Hoffman, Thomas S. King V, Linda R. Elliott, Samuel G. Schifflett, and Scott Chaiken. Classifying C2 Decision Making Jobs Using Cognitive Task Analyses and Verbal Protocol Analysis. In *Proceedings of the 6th International Command and Control Research Technology Symposium*, 2001.
- [30] David Harel. Statecharts: A Visual Formalism for Complex Systems. *Science of Computer Programming*, 8(3):231–274, 1987.
- [31] Mark M. Hickie. Behavioral Representation of Military Tactics for Single-Vehicle Autonomous Rotorcraft via Statecharts. Master’s thesis, Massachusetts Institute of Technology, June 2005.
- [32] Douglas L. Hintzman. Judgments of Frequency and Recognition Memory in a Multiple-Trace Memory Model. *Psychological Review*, 95(4):528–551, 1988.
- [33] J.H. Holland. *Adaptation in Natural and Artificial Systems*. University Michigan Press, Ann Arbor, MI, 1975.
- [34] Hui-Min Huang, James Albus, Elena Messina, Robert Wade, and Woody English. Specifying Autonomy Levels for Unmanned Systems: Interim Report. In *Proceedings of the SPIE Defense and Security Symposium, Conference 5422*, 2004.

- [35] Randolph M. Jones, John E. Laird, Paul E. Nielsen, Karen J. Coulter, Patrick Kenny, and Frank V. Koss. Automated Intelligent Pilots for Combat Flight Simulation. *AI Magazine*, Spring 1999.
- [36] Daniel Kahneman, Daniel Slovic, and Amos Tversky. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, New York, N.Y., 1982.
- [37] David Kinny, Michael Georgeff, and Anand Rao. A Methodology and Modelling Technique for Systems of BDI Agents. In *Agents Breaking Away: Proceedings of the Seventh European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, Lecture Notes in Artificial Intelligence 1038. Springer-Verlag, January 1996.
- [38] Gary Klein. *Sources of Power: How People Make Decisions*. MIT Press, Cambridge, MA, January 1998.
- [39] Raanan Lipshitz and Adi Adar Pras. Not Only for Experts: Recognition-Primed Decisions in the Laboratory. In *New Directions in Research on Decision Making*, Berndt Brehmer, editor, chapter 6, pages 91–105. Elsevier Science Pub. Co., New York, July 1986.
- [40] Jeremy Malasky, Laura Major Forest, Alexander C. Kahn, and J. Ramsay Key. Experimental Evaluation of Human-Machine Collaboration in Resource Allocation and Planning for Multiple UAVs. In *Proceedings for IEEE Conference on Systems, Man, and Cybernetics*, Waikoloa, Hawaii, 2005.
- [41] The Mathworks MATLAB and Simulink for Technical Computing. Fuzzy Logic User’s Guide. MATLAB documentation, 2005.
- [42] The Mathworks MATLAB and Simulink for Technical Computing. <http://www.mathworks.com>. Commercial Software, 2005.
- [43] The Mathworks MATLAB and Simulink for Technical Computing. Stateflow User’s Guide. MATLAB documentation, 2005.
- [44] George A. Miller. The Magical Number Seven Plus or Minus Two: Some Limits of Our Capacity for Processing Information. *Psychological Review*, 63:81–97, 1956.
- [45] Henry Montgomery, Raanan Lipshitz, and Berndt Brehmer. Naturalistic Decision-Making Research. In *New Directions in Research on Decision Making*, Berndt Brehmer, editor, chapter 1, pages 1–3. Elsevier Science Pub. Co., New York, July 1986.
- [46] D. Morales and M.L. Cummings. UAVs as Tactical Wingmen: Control Methods and Pilots’ Perceptions. MIT aero/astro department, to be published in Unmanned Systems.

- [47] Allen Newell and Herbert A. Simon. *Human Problem Solving*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1972.
- [48] Emma Norling and Liz Sonenberg. Creating Interactive Characters with BDI Agent. In *Proceedings of the Australian Workshop on Interactive Entertainment*, 2004.
- [49] Joint Strike Fighter Programs Office. Overview of Joint Strike Fighter Program. Presentation by Lockheed Martin and attended by the author, 2003.
- [50] UAS Roadmap 2005. Technical report, Office of the Secretary of Defense, United States of America, August 2005.
- [51] M.M. Omodei, J. McLennan, G.C. Elliott, A.J. Wearing, and J.M. Clancy. A Bias Towards Overuse of Resources. In *New Directions in Research on Decision Making*, Berndt Brehmer, editor, chapter 2, pages 38–41. Elsevier Science Pub. Co., New York, July 1986.
- [52] Committee on the Use of Humans as Experimental Subjects. Consent to Participate in Non-biomedical Research. <http://web.mit.edu/committees/couhes/forms.shtml>, 2005.
- [53] Raja Parasuraman and Victor Riley. Humans and Automation: Use, Misuse, Disuse, and Abuse. *Human Factors*, 39(2):230–254, June 1997.
- [54] Raja Parasuraman, Thomas B. Sheridan, and Christopher Wickens. A Model for Types and Levels of Human Interaction with Automation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 30(3):286–294, May 2000.
- [55] Jean-Charles Pomerol. Artificial Intelligence and Human Decision Making. Foreign Conference in Jerusalem, July 1995.
- [56] Lawrence J. Printzel. Team-Centered Perspective for Adaptive Automation Design. Technical Report NASA/TM-2003-212154, Langley Research Center, NASA, Hampton, VA, February 2003.
- [57] James Reason. *Human Error*. Cambridge University Press, New York, N.Y., 1990.
- [58] S. Ian Robertson. *Problem Solving*. Psychology Press, Philadelphia, PA, 2001.
- [59] Karol G. Ross, Gary A. Klein, Peter Thunholm, John F. Schmitt, and Holly C. Baxter. The Recognition-Primed Decision Model. *Military Review*, pages 6–10, July-August 2004.
- [60] W.B. Rouse. Models of Human Problem Solving: Detection, Diagnosis, and Compensation for System Failures. In *Proceedings of IFAC Conference on Analysis, Design, and Evaluation of Man-Machine Systems*. Baden-Baden FRG, September 1981.

- [61] Alan C. Schultz and John J. Grefenstette. Improving Tactical Plans with Genetic Algorithms. In *Proceedings of the 2nd International IEEE Conference on Tools for Artificial Intelligence*, pages 328–334, November 1990.
- [62] Thomas B. Sheridan. *Telerobotics, Automation, and Human Supervisory Control*. MIT Press, 1992.
- [63] Thomas B. Sheridan. Allocating Functions Rationally Between Humans and Machines. *Ergonomics in Design*, 6(3):20–25, July 1998.
- [64] Thomas B. Sheridan. *Humans and Automation*, chapter 7, pages 130–142. Wiley Series in System Engineering and Management and HFES Issues in Human Factors and Ergonomics Series. Human Factors and Ergonomics Society, 2002.
- [65] Taha A. Sidani and Avelino J. Gonzalez. A Framework for Learning Implicit Expert Knowledge through Observation. In *Transactions of the Society for Computer Simulation*, volume 17, pages 54–72, June 2000.
- [66] Herbert A. Simon. *Models of Man*. John Wiley & Sons, New York, 1957.
- [67] Herbert A. Simon. *Reason in Human Affairs*. Basil Blackwell, London, 1983.
- [68] Herbert A. Simon and W.G. Chase. Skill in Chess. *American Scientist*, 61:394–403, 1973.
- [69] Captain Tyler Smith. AV Application in Military Scenarios. Personal conversation with Captain Smith, Summer 2005.
- [70] John A. Sokolowski. Can a Composite Agent Be Used to Implement a Recognition-Primed Decision Model? In *Proceedings of the Eleventh Conference on Computer Generated Forces and Behavioral Representation*, pages 473–478, May 2002.
- [71] Terry Stanard and Robert J. B. Hutton. A Computational Model of Driver Decision Making At An Intersection Controlled By a Traffic Light. In *Proceedings of the International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design*, 2001.
- [72] Brian L. Stevens and Frank L. Lewis. *Aircraft Control and Simulation*. Wiley, Hoboken, New Jersey, 2 edition, 2003.
- [73] Camtasia Studio. <http://www.techsmith.com/camtasia.asp?CMP=KgoogleCStm>. Commercial Software, 2005.
- [74] John A. Swets, Robyn M. Dawes, and John Monahan. Psychological Science Can Improve Diagnostic Decisions. *Psychological Science in the Public Interest*, 1(1), May 2000.

- [75] Glenn Taylor and Robert E. Wray. Behavior Design Patterns: Engineering Human Behavior Models. In *Proceedings of the Behavior Representation in Modeling and Simulation Conference*, May 2004.
- [76] Telex. <http://www.telex.com>. Commercial Hardware, 2005.
- [77] The Navy Unmanned Undersea Vehicle (UUV) Master Plan. Technical report, The Department of the Navy, United States of America, November 2004.
- [78] Stephen Thrasher. A Reactive/Deliberative Planner Using Genetic Algorithms on Tactical Primitives. Master's thesis, Massachusetts Institute of Technology, June 2006.
- [79] Peter Thunholm. Planning Under Time Pressure: An Attempt Toward a Prescriptive Model of Military Tactical Decision Making. In *New Directions in Research on Decision Making*, Berndt Brehmer, editor, chapter 4, pages 43–55. Elsevier Science Pub. Co., New York, July 1986.
- [80] United States Army. *Army Regulation 95-1: Flight Regulations*, September 1997. Chapter 5, Section 5-2.b.
- [81] United States Marine Corps. *Field Training Manual - Student Handout*. reference taken from Wikipedia Online Encyclopedia.
- [82] Michael van Lent and John Laird. Learning by Observation in a Complex Domain. In *Proceedings of the Banff Knowledge Acquisition Workshop*, 1998.
- [83] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, N.J., 1953.
- [84] Scott A. Wallace and John E. Laird. Toward Automatic Knowledge Validation. In *Proceedings of the Eleventh Conference on Computer Generated Forces and Behavioral Representation*, pages 447–456, May 2002.
- [85] Scott A. Wallace and John E. Laird. Behavior Bounding: Toward Effective Comparisons of Agents and Humans. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2003.
- [86] Walter Warwick, Stacey McIlwaine, and Rob Hutton. Developing Computational Models of Recognition-Primed Decisions: Progress and Lessons Learned. Looking at Hintzman's Memory Model and Encoding Expectancies with RPD framework.
- [87] Christopher D. Wickens and Justin G. Hollands. *Engineering Psychology and Human Performance*, chapter 8, pages 293–325. Prentice Hall, Upper Saddle River, New Jersey, third edition, 2000.
- [88] Christopher D. Wickens and Justin G. Hollands. *Engineering Psychology and Human Performance*, chapter 13, pages 528–550. Prentice Hall, Upper Saddle River, New Jersey, third edition, 2000.

- [89] Christopher D. Wickens and Justin G. Hollands. *Engineering Psychology and Human Performance*, chapter 7, pages 241–284. Prentice Hall, Upper Saddle River, New Jersey, third edition, 2000.
- [90] Earl L. Wiener. Beyond the Sterile Cockpit. *Human Factors*, 27(1):75–90, February 1985.
- [91] Robert E. Wray and John E. Laird. Variability in Human Behavior Modeling for Military Simulations. In *Proceedings of the Behavior Representation in Modeling and Simulation Conference*, May 2003.