Toward The Development Of A Predictive Computer Model Of Decision Making Under Uncertainty For Use In Simulations Of U.S. Army Command And Control Systems

Sam E. Middlebrooks
Human Research and Engineering Directorate
U.S. Army Research Laboratory
Fort Hood, Texas
Sam.Middlebrooks@otc.army.mil

Brian J. Stankiewicz
Department of Psychology & Center for Perceptual Systems
The University of Texas at Austin
bstankie@mail.utexas.edu

The 13th Annual U.S. Military Academy / U.S. Army Research Laboratory
Technical Symposium
2-4 November 2005
West Point, New York and Aberdeen Proving Ground, Maryland

KEYWORDS: partially observable Markov decision processes (POMDP), Bayesian statistics, decision making under uncertainty (DMUC), command and control systems (C2S), computer simulation.

1. Executive Summary.

In today's increasingly complex world of digital command and control, it is seldom obvious or intuitive how the introduction of new automation systems will affect the overall performance of battlefield command and control systems. Field observations can account for performance factors that are directly observable such as rates of communication flow, rates of flow and quality of incoming intelligence. However, what the human mind does under the influence of all of these factors is not directly observable and is the subject of considerable experimentation.

This research addresses this limitation through the development of predictive quantitative models of decision making under conditions of uncertainty such as exist in many aspects of human performance and certainly in battle-space management. Utilizing Bayesian statistical approaches implemented through Partially Observable Markov Decision Processes (POMDP) that describe experiential decision processes moderated by Monte Carlo effects to account for performance variability, a series of computer simulations are being developed with the goal to predict the quality of decisions possible from a given set of input conditions.

These simulations are based on cognitive models being developed in a collaborative effort through a series of empirical studies that investigate human performance in a sequential decision making with uncertainty task using human subjects. Through this collaboration, the

results of these studies are being applied at each stage of the research to predictive computer simulations of Army battlefield performance where battlefield automated command and control systems are involved. These simulations, when operational, will allow cognitive effects, such as predictive levels of effective decisions possible from a given set of circumstances, to be assessed as a battlefield metric. The usefulness of these simulations will be realized in their ability to predict cognitive performance improvements that can potentially be realized through modifications of the work system such as organizational changes, new system components, and changes in training levels of the team members.

2. Introduction.

The time is now. The place is a U.S. Army command center in a far away land conducting combat operations. To the uninitiated observer this would seem like a place of uncontrolled chaos. Computers, computer generated tactical projection displays, and communications devices fill every available space. People are running back and forth yelling at each other with each one having a seemingly singular role to do whatever they are focused upon with no regard to other activities around them. Multiple radios are blaring from all sides of the workspace. Individual groups of 2-3 people are engrossed in huddles in various corners of the area and are oblivious of everything else going on around them. The commander is yelling at a computer operator to get a computer generated tactical map projection display updated and current with the admonition at a high fever pitch "you are killing me!!!". The computer operator is feverishly and frantically calling for help from adjacent operators and NCOs… While it is not obvious to the casual observer, there is structure to this chaos and order to the disorder resulting in a carefully balanced mix of people, machines, and weapons conducting an orchestrated performance on the battlefield.

Into the midst of all this seeming pandemonium a new communications device is brought on line that has been developed by the best and brightest minds that the Army Acquisition community can bring to bear on the problem. Its design promises to increase the rate of communications flow into the tactical operations center by 150% with an improvement in data quality by 90% (sample numbers for an imaginary system). However, there has been no chance to validate its design promises because of a rapid fielding initiative that put it into the field in an accelerated timeline. Even if an opportunity had been afforded to put the system under field testing conditions the results of the test would typically validate that the system is or is not working as designed and whether it does provide greater communications flow at a higher fidelity. Even if these field tests had been conducted, there would still be the unanswered question of the overall effect on the battlefield that the introduction of this new system would have. This effect would be the result of the introduction of a new system component into the command and control system with the potential to change the overall performance of the total system. Furthermore, while the physical parameters of the new system can be measured and quantified, the effect of these changes on the cognitive performance of the work group and their ability to adapt to changes in operational paradigms is much harder to estimate. In fact, even in optimal laboratory conditions it is a complex undertaking to assess cognitive performance factors such as situation awareness, individual and team performance and the effects that these factors have on decision-making performance. Furthermore, many laboratory based evaluations are conducted as a part of a basic research effort and are left to future research to apply to specific application areas.

Thus, the development of predictive computer-based models of optimal performance has a significant potential to aid in the evaluation of overall work systems where the human is an integral component. This project has an objective to develop computer models of optimal performance in the area of decision making under conditions of uncertainty that is the result of a coordinated effort that allows basic research utilizing empirical investigations to be directly applied into the structure of a predictive computer simulation of decision making.

2.1. Modeling the human ability to make efficient decisions under conditions of uncertainty.

One of the most fundamental aspects of human cognition is the ability to make decisions. Humans can make decisions in a broad range of domains. For example, everyday decisions are made as to what time an individual will depart for work, where or what will be eaten for lunch, what will be worn to work in addition to thousands of other decisions. These decisions often seem mundane because the ramification of a "poor" decision are not significant. However, other decisions appear to be more critical and the ramification of a "poor" decision can appear, and often are, more detrimental. For example, military decisions, medical decisions, and fault detection can all have significant ramifications.

It is important to recognize that most decisions that are made are not "one off" decisions in which the decision is made and then the rewards reaped or the punishment endured. Instead, most decisions that are made have future ramifications and affect the options and decisions that are available later. One challenge faced by any decision maker is the uncertainty that the decision maker has about the **true** state of the system. In most circumstances the true state of the system is unknown or hidden. That is, it cannot be directly observed. For example, in military decisions often there is uncertainty about an enemy's position, strength and morale. Given that the true state is hidden, there are things that can be done to reduce the decision maker's uncertainty about these states. For example, the decision maker may attempt to determine the enemy's position by sending reconnaissance to a location where the enemy is believed to be located. When the reconnaissance returns with either an "enemy sighted" or "enemy not sighted" report the decision maker must update their belief about the location of the enemy.

If the observations and actions were all deterministic, updating a belief would be relatively simple. However, in almost all conditions, the observations and actions are probabilistic. That is, the probability of getting an observation given the true state of the environment is not necessarily 0.0 or 1.0. Or, in the example above, there is a certain non-zero probability that the reconnaissance mission was sent to the right location and will **miss** the enemy and send a report of "enemy not sighted". Furthermore, there may be a non-zero probability that the reconnaissance mission falsely sent a report of "enemy sighted" (or false-alarmed) when the enemy was not actually at the location.

Given that the observation and actions are probabilistic, updating a belief given an observation and an action can become cognitively difficult. Furthermore, evaluating the added benefit of a specific piece of equipment that changes these probabilities can also become difficult. The current research focuses on a task that is commonly faced by decision makers in the military — namely, a seek-and-destroy task. In this task the decision maker is trying to localize and destroy an enemy within a specific region. At the decision maker's disposal are actions that allow him to gain information about the true state of the system (i.e., the location of the enemy) in addition to changing the state of the system (for example, moving the enemy from being at a specific location to the state of Destroyed). The former actions are reconnaissance actions and the latter

are artillery actions. The outcomes of these actions are probabilistic. That is, reconnaissance actions will not always detect the enemy when a sensor is sent to the enemy's location. Furthermore, the reconnaissance may also falsely report that the enemy is seen at a location in which the enemy is not located. Furthermore, the artillery will not always move the enemy to the "Destroyed" state when it is sent to the right location.

2.1.1. The Optimal Observer.

To best evaluate performance in a task that leads to uncertainty and probabilistic actions, it is useful to define the optimal performance within the task. The optimal performance can be calculated using Bayesian statistics. However, due to the nature of the current type of task, simple Bayesian statistics are insufficient. That is, with simple Bayesian statistics the likelihood of the true state of the system can be optimally estimated, but this likelihood does not indicate what action should be selected. In order to do action selection, not only must the current state be calculated given the previous actions and observations, but also the optimal action to be performed in a given belief state must be calculated where a belief state is a particular probability distribution across all of the possible states in the environment.

A variation on classical Bayesian statistics that may well add some additional predictive power for sequential decision making under uncertainty is the *Partially Observable Markov Decision Processes* (POMDP) (Cassandra, 1998; Cassandra, Kaelbling, and Kurien, 1996; Cassandra, Kaelbling, and Littman, 1994; Kaelbling, Littman, and Cassandra, 1998; Sondik, 1971). By defining the *State Space, Observation Vector, Transition Matrix* and the *Reward Structure*, the expected reward for a particular action can be computed. In the following sections a description of these actions will be provided. In addition, a description of how to optimally update an individual's belief (*Belief Updating*), given these definitions, is provided.

An Ideal Observer Model provides optimal performance given the information available in the task. Typically ideal observers are not proposed as models of human cognition. Instead, the ideal observer provides a benchmark by which to compare human performance. More specifically, these models illustrate what optimal performance should look like. When human performance matches that of the ideal-observer model, one can conclude that the human is making use of all of the information in the task. When the human under-performs the ideal observer, specific discrepancies between the human data and the ideal data may illuminate the constraints imposed by the human information-processing system.

Ideal observer analysis has been used to understand perceptual functions from the quantum limits of light detection (Hecht and Shlaer, 1942) to many forms of visual pattern detection and discrimination (Geisler, 1989), to reading (Legge and Hooven, 2002; Legge and Klitz, 1997) object recognition (Liu and Knill, 1995; Tjan and Braje, 1995; Tjan and Legge, 1998) eye movements (Najemnik and Geisler, 2005) and also in reaching tasks (Trommershäuser and Gepshtein, 2004).

2.1.2. Defining the State Space.

In all problems that are solved using a POMDP architecture, there is a set of *states* that the problem can be in. In a POMDP problem, the true state ($State_{True}$) is not directly observable (i.e., it is hidden). For the problems used in this project, the hidden state was defined as the enemy's current position within the 5*x*5 state space area plus an additional "Destroyed" state that the

enemy could transition into following an artillery strike at its current position for a total of a 26 state space.

2.1.3. Defining the Observation Vector.

Although the true state is hidden, the observer typically has actions and observations that provide information about the true state of the problem. In the current problem, the observer can fire artillery at a specific position or reconnaissance can be sent to a particular location within the environment (i.e., one of the 25 locations). The current problem has three possible observations: *Enemy Sighted*, *No Enemy Sighted*, or *No Information*. When the observer decides to send reconnaissance to a particular location, one of two observations will be received: 'Enemy Sighted' or 'Enemy Not Sighted'. In the current problem, the artillery returns only one possible observation: 'No Information'. This replicates the fact that the artillery firing unit does not see the effects of its fires because it is an indirect firing unit and is not able to see where the artillery rounds fall. It must rely on forward observer assets to report back what is termed 'battle damage assessment (BDA)' in military jargon. In these models it is the UAV that provides the BDA.

2.1.4. Defining the Transition Matrix.

Thus, in this Seek & Destroy problem, the commander has 51 possible different actions. There are 25 reconnaissance actions (one to each of the 25 locations in the environment), 25 artillery actions (again, one to each of the 25 locations within the environment) and the 'declare destroyed' option. The transition matrix defines the probability of the resulting state if the observer generates a particular action in a specified state (i.e., $p(s'|s,a)$). In the static form of the Seek & Destroy problem there is only one state transition that could occur. When the commander fires artillery to where the enemy is located, the enemy will transition into the "Destroyed" state with a probability of 0.75. Probabilities that are assumed for the purpose of this analysis, that are merely estimations for the purposes of this discussion, are shown in Tables 1-3. These values are estimates only and are not to be construed as factual. The determination of valid probabilities is left for future field and empirical work.

Table 1 – The set of actions and their observations for the current Seek & Destroy task.

(The observations for the reconnaissance action are dependent upon whether the enemy is actually within the viewing region of the reconnaissance. Thus, the two possible states are "Enemy Present" and "Enemy Not Present".)

Action	Observation	State	Probability
Recon	Enemy Sighted	Enemy Present	0.75
Recon	Enemy Not Sighted	Enemy Present	0.25
Recon	Enemy Sighted	Enemy Not Present	0.2
Recon	Enemy Not Sighted	Enemy Not Present	0.8

Table 2 – Probabilities for Observation from Artillery Strike

Action	Observation	State	Probability
Strike	NoInfo	Enemy Present	1.0
Strike	NoInfo	Enemy Not Present	1.0

Table 3 – Probabilities for Killing Enemy from Artillery Strike

Action	Result	State	Probability of Dead
Strike	Probability of Enemy being killed.	Enemy Present	0.75
Strike	Probability of Enemy not being killed.	Enemy Present	0.25
Strike	Probability of Enemy being killed.	Enemy Not Present	0.0

2.1.5. Belief Updating.

Given an initial probability distribution over the state space, the Observation Matrix and the Transition Matrix, hypotheses can be generated about the current state of the problem following an action and the returned observation. Equation 1 provides the Bayesian updating rule.

$$p(s'|b,o,a) = \frac{p(o|s',b,a)p(s'|b,a))}{p(o|b,a)} \tag{1}$$

where

s' ≡ true state (of the condition being present within the total of all states, S), represented as: s' ∈ S

b ≡ prior belief

o ≡ observation

a ≡ action that was generated

Equation 1 specifies how the ideal observer would update its belief that *s*' is the true state given the prior belief (*b*), the observation (*o*) and the action that was generated (*a*).

To illuminate the process of belief updating, a simple example of a smaller Seek and Destroy problem is provided. The Transition and Observation Matrices used in the empirical studies will be utilized here. But to simplify the process a three-state problem will be used instead of a 26-state problem. More specifically, the enemy will be in one of three states: $State_1$, $State_2$ or Destroyed. The initial (prior) probability of the state will be: $State_1$=0.5, $State_2$=0.5 and Destroyed=0.0 (to simplify, this is represented as [0.5, 0.5, 0.0]) meaning that there is a 50% probability of the enemy of being in $State_1$ a 50% probability of the enemy being in $State_2$, and a 0% probability of the enemy being in *State Dead*, i.e., the enemy is alive.

Assume that the enemy is in $State_1$ and that the observer decides to do reconnaissance to $State_1$ and receives a "Enemy Sighted" observation. What is the likelihood of the belief that the enemy is in $State_1$, $State_2$ or *State Dead*?

Using Equation 1 the likelihood that the enemy is in $State_1$ can be updated. That is, the desire is to compute: *p*(State1| [0.5, 0.5, 0.0], "EnemySighted",Recon1).

First, compute $p(o|s',b,a)$ or p("EnemySighted"|State$_1$, [0.5, 0.5, 0.0],Recon$_1$). To do this the likelihood of obtaining an observation of "Enemy-Sighted" if $State_1$ was the true state is needed.

In the Observation Matrix section above, the likelihood of correctly identifying the enemy as 0.75 is defined. The likelihood of the true state being $State_1$ given the previous belief and the action of $Recon_1$ is needed. Because there is no transition possible, these remain at the prior probabilities of 0.5. Finally, the likelihood of receiving the observation "EnemySighted" when Reconnaissance is made at $State_1$ or p('EnemySighted'|[0.5,0.5, 0],Recon$_1$) needs to be computed.

$$p(o|b,a) = ((0.5 \times 0.75) + (0.5 \times 0.2) + (0.0 \times 0.2)) \quad (2)$$
$$= 0.475 \quad (3)$$

$$p(State_1|[0.5, 0.5, 0.0], \text{“}EnemySighted\text{”}, Recon_1) = \frac{0.75 \times 0.5}{0.475} \quad (4)$$
$$= 0.7895 \quad (5)$$

Furthermore,

$$p(State_2|[0.5, 0.5, 0.0], \text{“}EnemySighted\text{”}, Recon_1) = \frac{0.20 \times 0.5}{0.475} \quad (6)$$
$$= 0.2105 \quad (7)$$

and Finally,

$$p(Destroyed|[0.5, 0.5, 0.0], \text{“}EnemySighted\text{”}, Recon_1) = \frac{0.2 \times 0.0}{0.475} \quad (8)$$
$$= 0.0 \quad (9)$$

Thus, if the first action is to observe at $State_1$, the new belief vector would be [0.7895,0.2105,0.0].

Now, imagine that the observer decided on the action $\boldsymbol{Strike_1}$ following the action $\boldsymbol{Recon_1}$. To update the belief that the enemy is in $State_1$ the conditional probability p(State1|[0.7985, 0.2105, 0.0], "NoInfo", Strike1) is computed.

First, compute $p(o|s',b,a)$ which is the probability of receiving the "NoInfo" observation given that the true state is $State_1$, the current belief ([0.7895,0.2105,0.0]), and the action $Strike_1$. The probability of receiving this observation is 1.0. Regardless of the state of the problem, a Strike always returns the observation "NoInfo" (see Table 2). This same probability and logic holds for computing $p(o|b,a)$.

The conditional probability $p(s'|b,a)$ also needs to be computed. That is, what is the probability of the true state being $State_1$ given the current belief and the action $Strike_1$. As described in Section 2.1.4 on the Transition Matrix, the probability of transitioning the problem into the Destroyed state is 0.75 if the enemy is at the location where the artillery strike occurred. This means that there is a probability of 0.25 that the enemy's state will not change, or that the enemy will remain in $State_1$ if it was initially in $State_1$.

$$\begin{aligned} p(o|s',b,a) &= 1.0 \qquad (10) \\ p(s'|b,a) &= 0.7895 \times 0.25 \\ p(o|b,a) &= 1.0 \\ p(State_1|[0.7895, 0.2105, 0.0], \text{“}NoInfo'', Strike_1) &= \frac{1.0 \times 0.7895 \times 0.25}{1.0} \\ &= 0.1996 \end{aligned}$$

$$\begin{aligned} p(\text{“}NoInfo''|Destroyed, [0.7895, 0.2105, 0.0], Strike_1) &= 1.0 \qquad (11) \\ p(State_2|[0.7895, 0.2105, 0.0], Strike_1) &= 0.2105 \\ p(\text{“}NoInfo''|[0.7895, 0.2105, 0.0], Strike_1) &= 1.0 \\ p(State_2|[0.7895, 0.2105, 0.0], \text{“}NoInfo'', Strike_1) &= \frac{1.0 \times 0.2105}{1.0} \\ &= 0.2105 \end{aligned}$$

$$\begin{aligned} p(\text{“}NoInfo''|Destroyed, [0.7895, 0.2105, 0.0], Strike_1) &= 1.0 \qquad (12) \\ p(Destroyed|[0.7895, 0.2105, 0.0], Strike_1) &= 0.7895 \times 0.75 \qquad (13) \\ p(\text{“}NoInfo''|[0.7895, 0.2105, 0.0], Strike_1) &= 1.0 \qquad (14) \\ p(Destroyed|[0.7895, 0.2105, 0.0], \text{“}NoInfo'', Strike_1) &= \frac{1.0 \times 0.7895 \times 0.75}{1.0} \qquad (15) \\ &= 0.5989 \end{aligned}$$

Thus the belief vector, following Recon1 with the observation of "EnemySighted" followed by Strike1 with an observation of "No Information", is [0.1996,0.2105,0.5989]. Or

$$\begin{aligned} p(State_1) &= 0.1996 \\ p(State_2) &= 0.2105 \\ p(Dead) &= 0.5989 \end{aligned}$$

2.2. Applying the Decision Making Under Uncertainty Model.

Using cognitive models such as those described above now allows computer simulations of command and control systems configured around task performance analysis based on previous work (Middlebrooks, 2001; Middlebrooks, 2003; Middlebrooks, 2004; Middlebrooks et al., 1999a; Middlebrooks et al., 1999b; Middlebrooks and Williges, 2002; Wojciechowski, Plott, and

Kilduff, 2005), to now be structured to incorporate cognitive decision making as a performance metric using this belief updating model. The steps in this process resemble the well known Observe-Orient-Decide-Act (OODA) model (Belknap, 1996; Salas, Morgan, Glickman, Woodard, and Blaiwes, 1986). The decision actions here described consist of gathering information, updating the belief about the environment or state space, taking an action to accomplish an objective in the state space, and then making a decision of whether to continue the mission or terminate it with an assessment of either mission success or failure. The example in this military command and control scenario employs a unmanned aerial vehicle (UAV) to gather the intelligence, artillery to take an action to destroy an enemy somewhere within the state space, and belief updating to evaluate the situation after each action and either continue or declare 'mission complete'.

To structure this scenario in a computer simulation, the programming environment of Command, Control, and Communications: Techniques for the Reliable Assessment of Concept Execution (C3TRACE) (Kilduff, Swoboda, and Barnette, 2005; Plott, 2002; Plott, Quesada, Kilduff, Swoboda, and Allender, 2004) is employed. C3TRACE, developed through funding by the Human Research and Engineering Directorate of the U.S. Army Research Laboratory, is an adaptation of the commercial discrete event programming language MicroSaint™ (Schunk and Plott, 2004). While the basic MicroSaint™ programming language allows for the development of task based computer simulations of real world systems and processes to be represented, C3TRACE has embedded data structures that augment MicroSaint™ to allow for representation of Army command and control systems.

2.2.1. Simulation Design.

C3TRACE programs are implemented using discrete event language constructs common to any MicroSaint™ simulation program. The top level of a command and control sub-workgroup within a sample organization is shown in the example depicted Figure 1. Here messages received by the radio operator are distributed according to their subject content. Situation Reports (SITREP) are passed to the S3 Operations officer, Logistics Reports are passed to the S4 Logistics officer for action and so on. If, for example, a mission directive such as seek out and destroy an enemy, is received it is passed to the commander for action. There are different reactions that might be experienced to such a directive. The commander might communicate back to the originating authority to clarify information, an initial estimate of the situation before taking action might be performed, an updating of the situational awareness before taking action might be performed, or, the mission might be undertaken as directed. In this case, as depicted in the green box in Figure 1, what is referred to as the Decision Making Under Uncertainty (DMUC) process would be initiated to execute the mission.

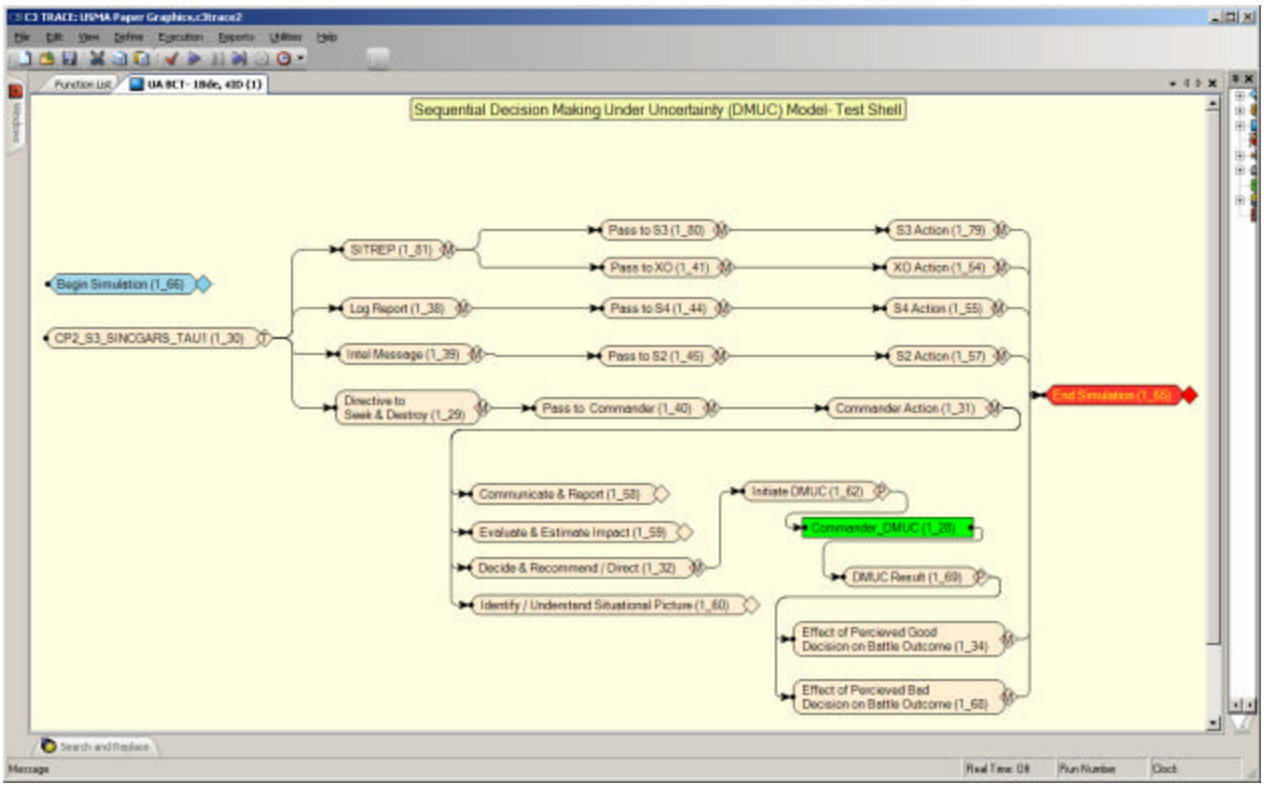

Figure 1 – C3TRACE Command and Control Simulation Vignette

Figure 2 illustrates a modeled decision process that is very similar to the OODA model. This diagram represents an iterative process where the decision maker makes an initial estimate of the situation and then begins an iterative process of either gathering additional information (flying a UAV mission) or taking an action to destroy the enemy (firing artillery). When the commander believes that the enemy has been destroyed, a mission complete decision is made and the results of the decision are realized. If the enemy was destroyed and the decision maker made that correct assessment, then a positive reward resulting from a good decision is applied to the performance of the overall system. If the enemy was not destroyed and the decision maker believed that it was destroyed, then a negative battlefield outcome is applied to the simulation. Likewise, if the enemy was destroyed but the decision maker believed it was not, then the results of poor decision making are applied. This process of iterative action can be generalized to similar scenarios where information is gathered (Observe), belief updating occurs (Orient), decisions are made for mission success (Decide), and actions are taken to accomplish the mission (Act). The examples of employing a UAV and firing artillery are used here to simply provide a tangible example of how this type of activity might occur.

2.2.2. The DMUC Decision Loop.

Referring to Figure 2, the top level logic for this model can be examined. After initiating the decision sequence and performing an initial estimate of the situation, the commander updates his belief vector, defined as the belief about the current situation regarding the enemy, and then begins an iterative process of looking for information or taking an action to accomplish the mission. When this process has reached some level of belief that the mission is accomplished,

the commander terminates the action and completes the decision process by either declaring a mission 'success' or 'failure'.

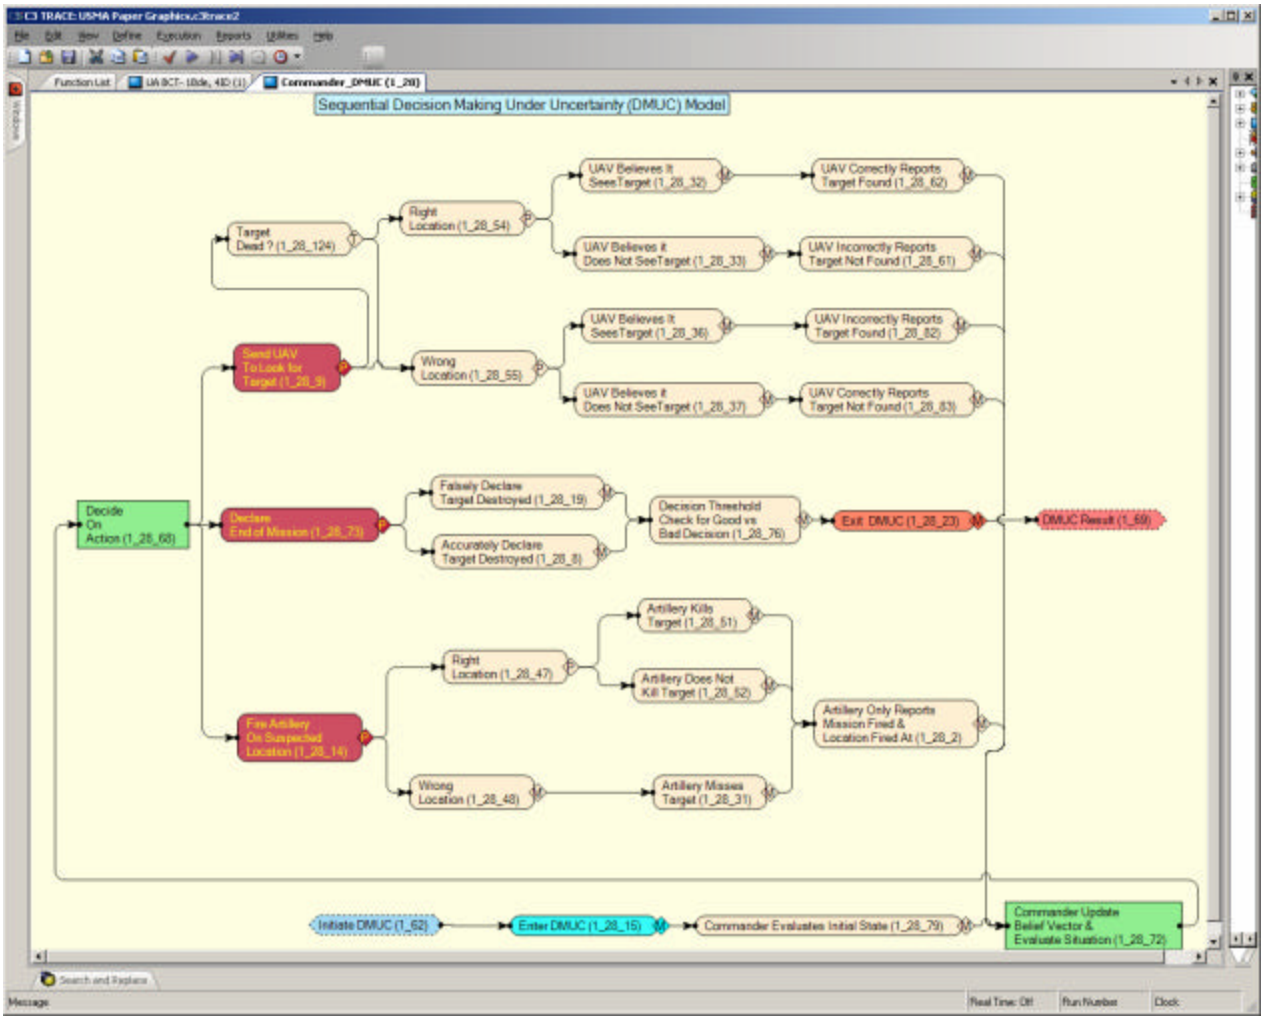

Figure 2 – Decision Making Under Uncertainty Model

If the initial desire is to obtain additional information, a UAV is sent to a specified location to attempt to locate the enemy. The UAV is the information gathering or battle damage assessment (BDA) tool available to the commander to update his belief vector about the enemy. If the target is already dead from previous artillery action, then there is no correct location for the enemy because it does not exist as it is dead. If the enemy is alive and the UAV is sent to the correct location then it has a probability, according to Table 1, of either detecting or not detecting the enemy representing the accuracy of the UAV. From this it will either correctly or incorrectly report that the enemy was found. Likewise, if it is sent to a location where the enemy is not located, or if the enemy is already dead, it may correctly or incorrectly report the enemy sighted again according to Table 1. The values in Table 1 are only sample estimates for use in the development of this model and do not represent any actual system currently in existence. After the UAV mission is flown the commander evaluates the report from the UAV through the

process of updating his belief vector (described below) and, using this new information, decides what process to invoke next.

If the commander decides to fire artillery (which is representative of taking a positive action to do something to accomplish the mission) then the probability exists that either the right or wrong location will be fired upon. If the artillery fires on the wrong location then the only outcome will be to miss the target. If the correct location is fired upon then the artillery will either kill or not kill the enemy according to the circular area of probability for the type of artillery fired. Independent of where the artillery is fired, the only report that is sent back to the commander is that the artillery fired upon the location directed. This represents the fact that artillery is an indirect fire weapon and the firing unit never actually sees the target. It is up to the forward observer, in this case the UAV, to report back as to the actual target situation, i.e., to provide the BDA. It is then up to the commander to evaluate the firing data and information from previous UAV reconnaissance missions to decide if to continue the mission or declare the enemy is dead and end the mission.

When the commander believes that the enemy has been destroyed, then mission complete is declared and the commander is faced with either the rewards of a successful or good decision sequence where the enemy was killed, meaning that the mission accomplished, or the effects of a bad decision where the mission was not accomplished.

2.2.3. Evaluating the Current Information.

Figure 3 illustrates the inputs feeding the sequence of evaluating the current situation, and updating of the belief vector and the resulting choice for the next action.

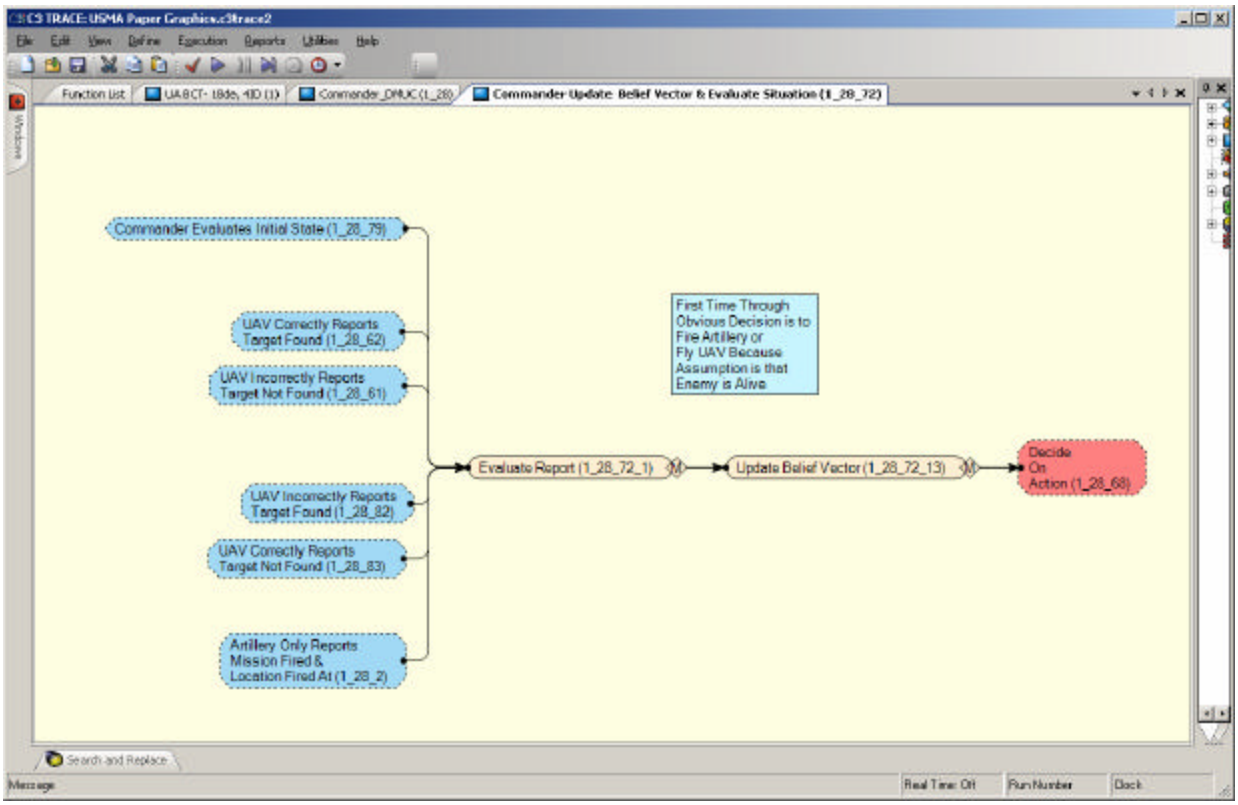

Figure 3 – Belief Vector Updating

The actual logic for each of these activities occurs within the C3TRACE program using C Sharp (C#) code statements embedded within the beginning and ending effects sections of each of the task blocks in the diagram.

A verbal description of the computer logic for the 'Evaluate Report' task takes the form:

```
For UAV Mission Cases:
        If (UAV mission ordered) & (Enemy Located at specified location)
                Then, There is a positive probability that the enemy will be reported sighted.
        If (UAV mission ordered) & (Enemy Not Located at specified location, or is dead)
                Then, There is a small probability that the enemy will still be reported sighted.
For Arty Mission Cases:
        If (Arty mission ordered) & (Enemy Located at specified location)
                Then, There is a positive probability that the enemy will be destroyed.
        If (Arty mission ordered) & (Enemy Not Located at specified location)
                Then, There is a zero probability that the enemy will be destroyed.
        For all Arty mission cases,
                Report = "No Info".
```

2.2.4. Updating the Belief Vector.

Figure 3 illustrates the Update Belief Vector task that follows the Evaluate Report Task. The pseudo logic that is implemented in C# for this task is:

```
For UAV Mission Cases:
        If (UAV mission ordered) & (Enemy Located at specified location)
                Then, Update Belief Vector with positive information for State Reconed.
                Then, Update Belief Vector with negative information for State Not Reconed.
                Then, Update Belief Vector with negative information for State Dead.
        If (UAV mission ordered) & (Enemy Not Located at specified location, or is dead)
                Then, Update Belief Vector with positive information for State Not Reconed.
                Then, Update Belief Vector with negative information for State Reconed.
                Then, Update Belief Vector with negative information for State Dead.
For Arty Mission Cases:
        If (Arty mission ordered) & (Enemy Located at specified location)
                Then, Update Belief Vector with positive information for State Fired Upon.
                Then, Update Belief Vector with negative information for State Not Fired Upon.
                Then, Update Belief Vector with positive information for State Dead.
        If (Arty mission ordered) & (Enemy Not Located at specified location)
                Then, Update Belief Vector with negative information for State Fired Upon.
                Then, Update Belief Vector with negative information for State Not Fired Upon.
                Then, Update Belief Vector with negative information for State Dead.
```

3. Discussion / Results.

Current work has focused on the simplest possible state space which consists of 2 location states and a state representing the status of the target enemy, either Dead or Alive. While a location state of only two conditions appears to be trivial and unrelated to any actual human performance condition, even this simple arrangement can relate to actual performance. While this simple 3 state model is not sufficient for actual analysis, it forms the basis for future, more complex,

models to become predictive of selected decision processes. When faced with a decision choice a human decision maker many times has one of two selections to make. Referring to Figure 4, whether it is right versus left, good versus bad, right versus wrong, high payoff versus low payoff, cheap versus expensive, or whatever the criterion, a two state response choice augmented by one result state (in this case Dead), can be descriptive of actual conditions. Also, even in this most simple of state environments, the conditional probability logic can become exhaustive. Computer runs with these conditions serve to illustrate the conditional probabilities resulting from relatively simple actions in this state space.

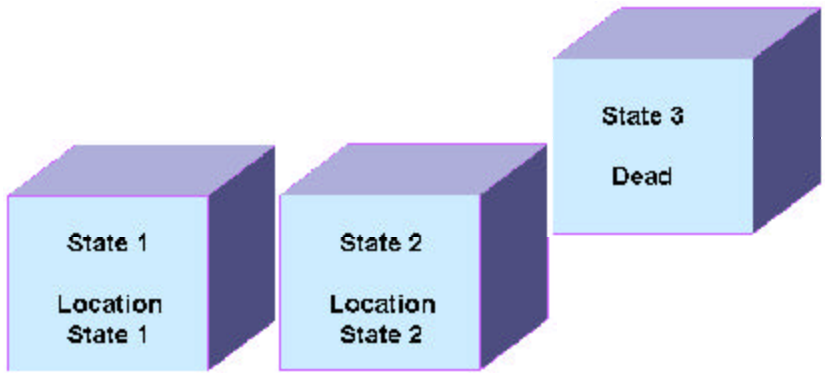

Figure 4 – 3 State Space

As shown in Figure 4 in this simple state space case, the enemy is either located in $State_1$ or $State_2$ for a binary location condition, with a condition of either Dead or Alive as indicated by $State_3$.

3.1. Action Sequence Assessment: Recon1, Strike1, Recon2-N

In order to examine the conditional probability logic associated with actions in this state space, some examples of actions and resulting belief vectors will be examined. The assumptions are that the enemy is located in State1 and that it is static, i.e., not moving. Also, there is an equal probability in the belief of the commander that the enemy could be in either of the location states and that the enemy is alive. The initial belief vector is thus [0.5, 0.5, 0.0], meaning a 50% chance of being in location State1, a 50% chance of being in location State2, and a 0.0% chance of being in State Dead, i.e., the enemy is alive.

Assume the following sequence of actions:

1. Conduct a UAV mission into State1, called Recon1
2. Fire artillery strike into State1, called Strike1
3. Conduct continuous UAV missions into State1, called Recon2 to ReconN.

Using the POMDP methodology in the C3TRACE DMUC model, examine the resulting beliefs that the decision maker has following the decision to perform this series of actions. The output from this simulation run is illustrated in Table 4, and Figure 5.

Table 4 – Probabilities for Sequence Recon1, Strike1, Recon2-N

NumDMUC_ Iterations	Action	State_Space[1,1] - Cell1	State_Space[1,2]- Cell2	State_Space_ Dead
1	Recon1	0.5000	0.5000	0.0000
2	Strike1	0.7895	0.2105	0.0000
3	Recon2	0.1974	0.2105	0.5921
4	Recon3	0.4797	0.1365	0.3838
5	Recon4	0.7757	0.0588	0.1655
6	Recon5	0.9284	0.0188	0.0528
7	Recon6	0.9799	0.0053	0.0149
8	Recon7	0.9945	0.0014	0.0040
9	Recon8	0.9985	0.0004	0.0011
10	Recon9	0.9996	0.0001	0.0003
11	Recon10	0.9999	0.0000	0.0001
12	Recon11	1.0000	0.0000	0.0000
13	Recon12	1.0000	0.0000	0.0000
14	Recon13	1.0000	0.0000	0.0000

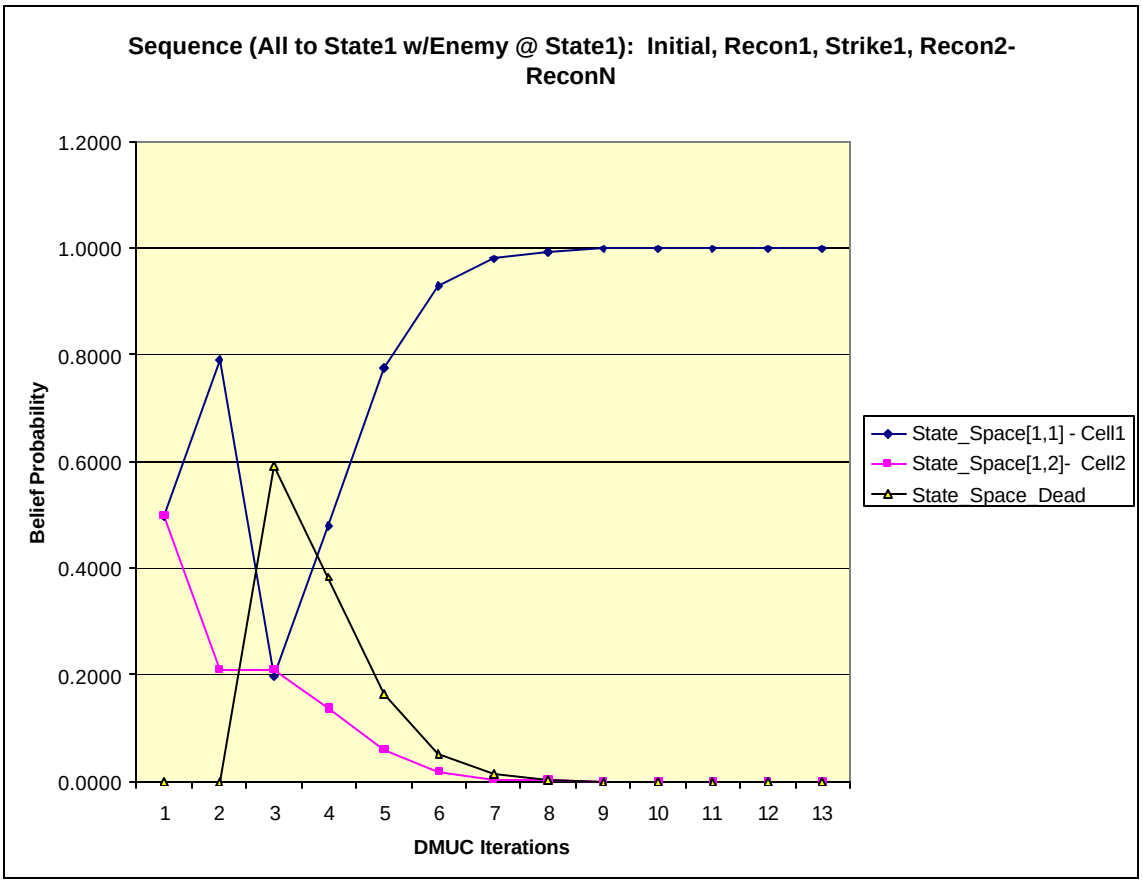

Figure 5 – Belief Graph for Action Sequence Recon1, Strike1, Recon2-N

These actions (first a recon, then a strike, and then only recons) produce a predicted belief pattern of almost 80% after the first recon that the enemy is in fact at $State_1$. After the strike the belief of being in $State_1$ drops to 21% while the belief that the enemy is dead goes from 0% to almost 60%. However, if the only succeeding actions at this point are to perform recons to $State_1$ which have no ability to harm the enemy because only an artillery strike can cause damage, then, because of successive incremental probabilities that the enemy might still be at $State_1$, and therefore not Dead, the belief value for State1 asymptotes at 100% while the values for $State_2$ and $State_3$ Dead go to 0%.

3.2. Action Sequence Assessment: Recon1, Strike1-N

Now consider the action sequence of an initial recon, Recon1, followed by multiple artillery strikes, Strike1-N. The previous assumptions for enemy location and initial belief vector are the same.

Table 5 – Probabilities for Sequence Recon1, Strike1-N

NumDMUC_ Iterations	Action	State_Space[1,1] - Cell1	State_Space[1,2]- Cell2	State_Space_ Dead
1	Recon1	0.5000	0.5000	0.0000
2	Strike1	0.7895	0.2105	0.0000
3	Strike2	0.1974	0.2105	0.5921
4	Strike3	0.0493	0.2105	0.7401
5	Strike4	0.0123	0.2105	0.7771
6	Strike5	0.0031	0.2105	0.7864
7	Strike6	0.0008	0.2105	0.7887
8	Strike7	0.0002	0.2105	0.7893
9	Strike8	0.0000	0.2105	0.7894
10	Strike9	0.0000	0.2105	0.7895
11	Strike10	0.0000	0.2105	0.7895
12	Strike11	0.0000	0.2105	0.7895

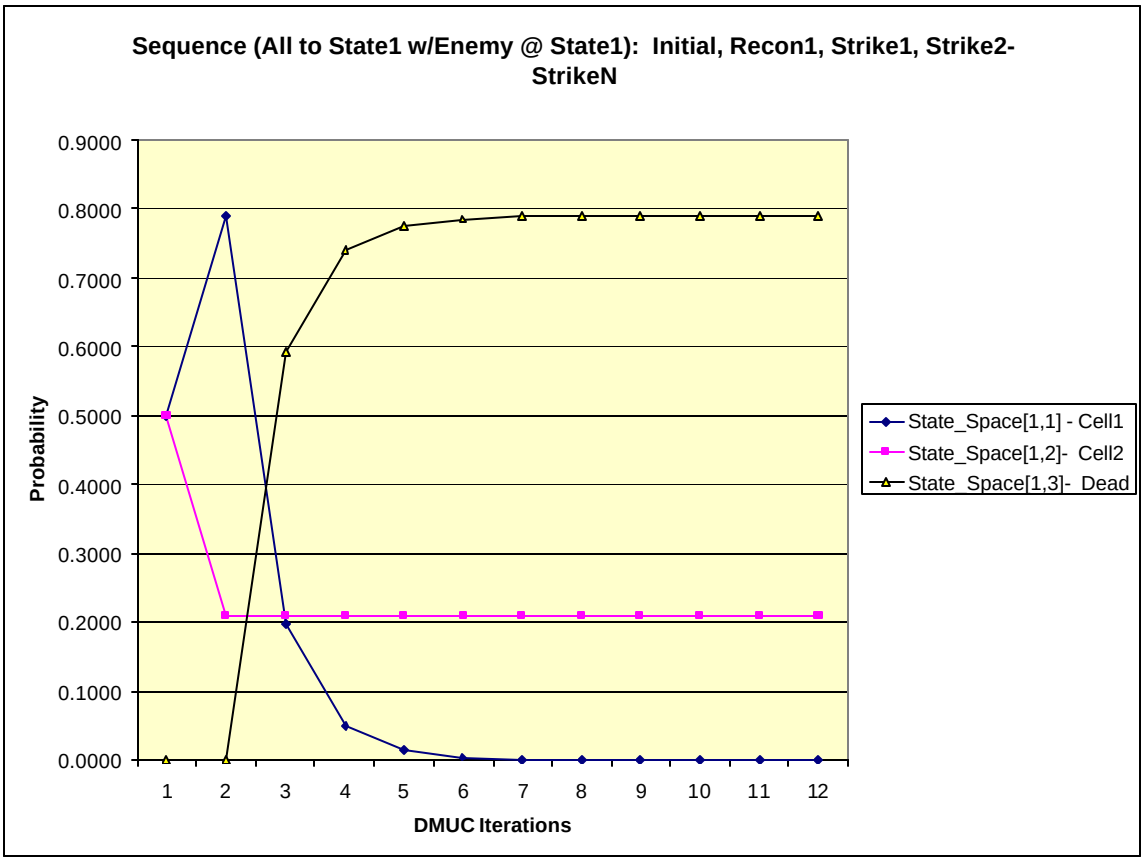

Figure 6 – Belief Graph for Action Sequence Recon1, Strike1, Strike2-N

After the first two belief updates, which are the same as in the previous runs, the belief values for the enemy being dead continue to rise with successive artillery strikes but asymptote at 78% while the belief value for the enemy still being in $State_1$ goes to zero. This is offset by a suspected belief condition for the enemy being at $State_2$ which is 21%. Since the belief probabilities for all three states must sum to 100% representing the complete belief condition, this causes the predicted belief for the enemy actually being dead to only max out at 78% even with successive artillery strikes each of which have the ability to completely kill the enemy. Actually, the C3TRACE model had the enemy killed after the first artillery strike, but the predicted belief by the commander has the values shown.

3.3. Action Sequence Assessment: Recon1-N

To check the response of the model, additional runs were made of only recon missions and only artillery missions. Table 6 and Figure 7 show the results for the case of only flying the UAV for multiple recons into $State_1$. Initial conditions remain the same as before.

Table 6 – **Probabilities for Sequence Recon1-N**

Action	State_Space[1,1] - Cell1	State_Space[1,2]- Cell2	State_Space_ Dead
Recon1	0.5000	0.5000	0.0000
Recon2	0.7895	0.2105	0.0000
Recon3	0.9336	0.0664	0.0000
Recon4	0.9814	0.0186	0.0000
Recon5	0.9950	0.0050	0.0000
Recon6	0.9987	0.0013	0.0000
Recon7	0.9996	0.0004	0.0000
Recon8	0.9999	0.0001	0.0000
Recon9	1.0000	0.0000	0.0000
Recon10	1.0000	0.0000	0.0000
Recon11	1.0000	0.0000	0.0000
Recon12	1.0000	0.0000	0.0000

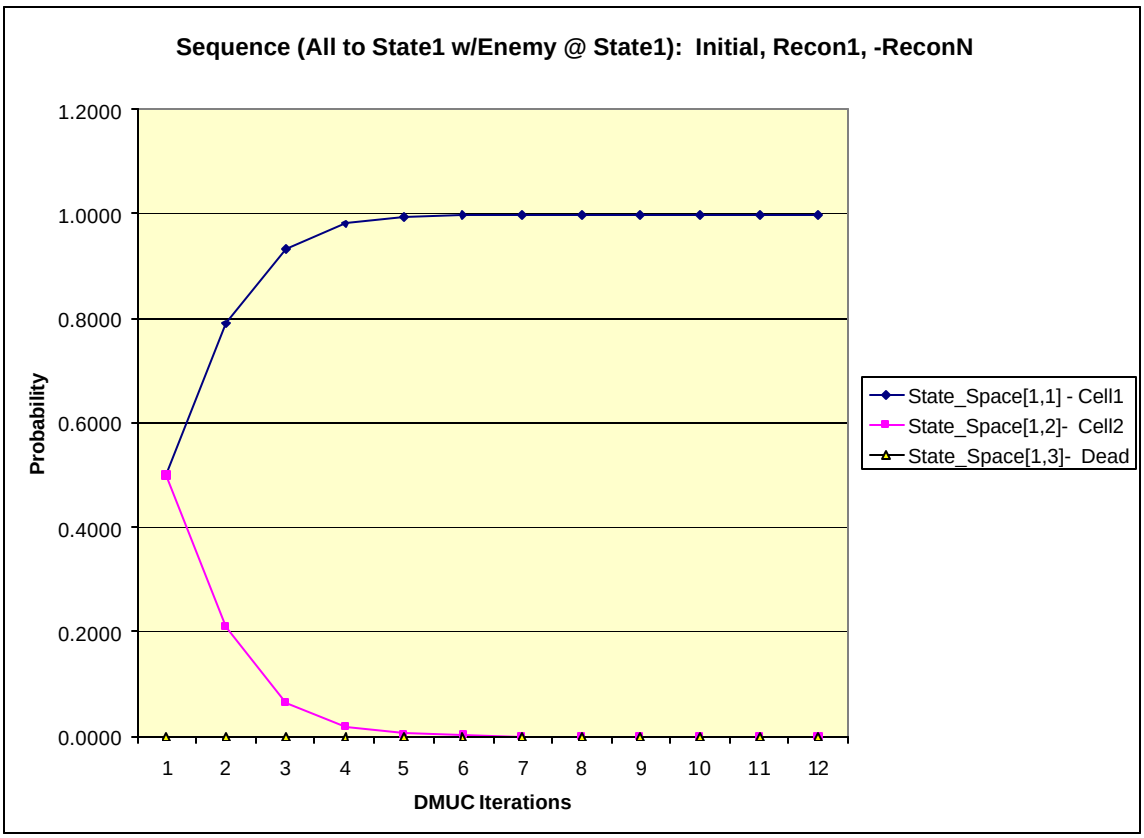

Figure 7 – Belief Graph for Action Sequence Recon1-N

Here, as would be predicted by common sense logic, the dead state remains at zero as recon missions have no ability to damage the enemy. The belief probabilities for the enemy at $State_1$ and $State_2$ complement each other with the $State_1$ belief that asymptotes at 100% while the $State_2$ belief asymptotes at 0%.

3.4. Action Sequence Assessment: Strike1-N

Finally, Table 7 and Figure 8 show the results for the case of only firing artillery for multiple strikes at $State_1$. Initial conditions remain the same as before. In this case, while the artillery has the ability to kill the enemy, it does not have the ability to report back its effects. Thus, the belief probability of the commander for both where the enemy is located and what his dead state is, is non existent as artillery strikes only report back that the mission was fired. As a result, only firing artillery will provide no information to the commander as to the status of the enemy even though the artillery might have actually destroyed the enemy. Thus, successive artillery strikes to $State_1$ result in no change to the belief that the enemy is at $State_1$ while causing only a maximum belief of 50% that the enemy might have been destroyed by the artillery missions. These models represent the total belief possible in the commander with an indication of 100% summed from all the individual beliefs. As a result of this 100% possible belief, this causes the belief that the enemy is at $State_2$ to asymptote at 0%.

Table 7 – Probabilities for Sequence Strike1-N

NumDMUC_ Iterations	Action	State_Space[1,1] - Cell1	State_Space[1,2]- Cell2	State_Space_ Dead
1	Strike1	0.5000	0.5000	0.0000
2	Strike2	0.1250	0.5000	0.3750
3	Strike3	0.0313	0.5000	0.4688
4	Strike4	0.0078	0.5000	0.4922
5	Strike5	0.0020	0.5000	0.4980
6	Strike6	0.0005	0.5000	0.4995
7	Strike7	0.0001	0.5000	0.4999
8	Strike8	0.0000	0.5000	0.5000
9	Strike9	0.0000	0.5000	0.5000
10	Strike10	0.0000	0.5000	0.5000
11	Strike11	0.0000	0.5000	0.5000
12	Strike12	0.0000	0.5000	0.5000

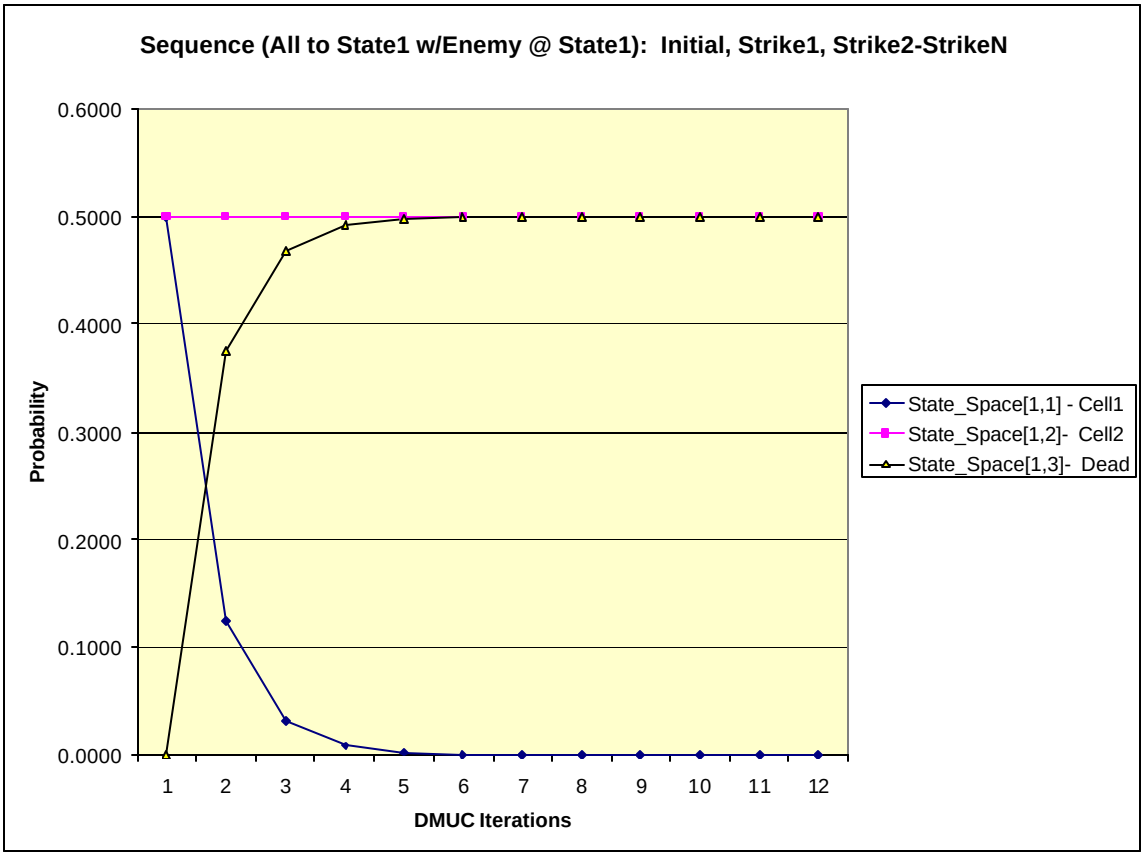

Figure 8 – Belief Graph for Action Sequence Strike1-N

4. Conclusions / Continuing Work.

This project is at the initial stage of developing a quantitative predictive model of optimal decision making performance in a form usable in simulations of command and control activities. The Bayesian statistical approach to the modeling is not intended to be construed as an attempt at presenting a human cognitive model of decision making, but rather an attempt to understand what optimal performance could be in a given set of circumstances. Once this optimal performance is understood then the Bayesian models can be used to compare predicted optimal performance against actual observed human performance data to gain an understanding of how human cognitive limitations can be affected by changes to different components of the human-computer interfaces in the work system. Some of the human performance areas where this technique might be applied include memory, decision strategy, and perception, to name a few. This technique has the promise to allow investigations of how modifying technology can effect decision strategies as they are represented by battlefield success and the cost of achieving that success.

The work to date has been and continues to be a partnership between basic empirical research that is investigating optimal performance through the modification of belief presentations and the accuracy of belief vectors, and applied research working to develop simulations of highly dynamic battlefield performance that are moderated by predicted optimal performance conditions existent within the workgroup by employing the C3TRACE modeling environment.

These initial efforts have developed the framework for a more complex series of models by developing performance algorithms based on a simple 3 element state space of 2 location states and one status state (Dead). The results presented here are not intended to provide a basis for actual investigative work. These data are included only for the purpose of illustrating the potential for this research. Future work will expand the effort to a 5 state space environment of a 2 x 2 location grid + Dead and on to a 26 state space environment of a 5x5 location grid + Dead. The current empirical work is based upon the 5x5 location grid + Dead, or 26 state space environment, condition. Also, the current work is based upon an enemy statically located in the grid. Future work will incorporate a dynamic enemy moving within the location grid according to some predetermined algorithm.

5. Acknowledgements.

This work has been supported by a grant from the congressionally funded University XXI program in a partnership between the faculty and staff of The University of Texas at Austin through the Institute for Advanced Technology and the Human Research and Engineering Directorate of the U.S. Army Research Laboratory.

Thanks go to colleagues who provided technical reviews of this manuscript. From the U.S. Army Research Laboratory's Human Research and Engineering Directorate these are Dr. Bruce Sterling at Fort Knox, Kentucky, Dr. Jim Ainsworth at Fort Hood, Texas, and Dr. Paul Rose and Dr. Dallas Johnson at Aberdeen Proving Ground, Maryland. From The University of Texas at Austin's Department of Psychology is Associate Professor Dr. Bradley Love.

The authors would also like to thank the following members of the faculty at the U.S. Military Academy at West Point, New York for their support during this project:

- LTC Garry Lambert, Ph.D., ARL/USMA R&D Coordinator, Department of Mathematics.

- Dr. Michael D. Matthews, Professor of Engineering Psychology, Department of Behavioral Sciences and Leadership.

Special thanks also go to West Point Cadet Edward Klein for his contributions to the project during his Advanced Individual Academic Development (AIAD) tour during the summer of 2005 while at Fort Hood, Texas and the UT Austin campus.

6. References.

Belknap, M. (1996). Military Decision Making in the Information Age. Unpublished Monograph, Naval War College, Newport, RI. Joint Military Operations Department. Ad-a307-648-6-xab.

Cassandra, A. R. (1998). Exact and Approximate Algorithms for Partially Observable Markov Decision Processes. Unpublished Dissertation, Brown University.

Cassandra, A. R., Kaelbling, L. P., & Kurien, J. A. (1996). Acting under Uncertainty: Discrete Bayesian Models for Mobile-Robot Navigation. In Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems.

Cassandra, A. R., Kaelbling, L. P., & Littman, M. L. (1994). Acting Optimally in Partially Observable Stochastic Domains. In Proceedings of the Twelfth National Conference on Artificial Intelligence. (AAAI), 1023--1028. http://citeseer.ist.psu.edu/cassandra94acting.html Seattle, Washington. MIT Press.

Geisler, W. S. (1989). Ideal-Observer Theory in Psychophysics and Physiology. Physica Scripta, 39, 153-160.

Hecht, S., & Shlaer, S. (1942). Energy, Quanta, and Vision. Journal of General Physiology, 25, 819-840.

Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and Acting in Partially Observable Stochastic Domains. Artificial Intelligence, 101, 99-134.

Kilduff, P. W., Swoboda, J. C., & Barnette, B. D. (2005). Command, Control, and Communications: Techniques for the Reliable Assessment of Concept Execution (C3TRACE) Modeling Environment: The Tool. U.S. Army Research Laboratory, Human Research and Engineering Directorate. Aberdeen Proving Ground, MD. ARL Technical Report ARL-MR-0617.

Legge, G. E., & Hooven, T. A. (2002). Mr. Chips 2002: New Insights from an Ideal-Observer Model of Reading. Vision Research, 42, 2219-2234.

Legge, G. E., & Klitz, T. S. (1997). Mr. Chips: An Ideal-Observer Model of Reading. Psychological Review, 104, 524-553.

Liu, Z., & Knill, D. C. (1995). Object Classification for Human and Ideal Observers. Vision Research, 35(4), 549-568.

Middlebrooks, S. E. (2001). Experimental Interrogation of Network Simulation Models of Human Task and Workload Performance in a U.S. Army Tactical Operations Center. Unpublished Thesis, Virginia Polytechnic Institute and State University, Blacksburg, VA.

Middlebrooks, S. E. (2003). The Compass Paradigm for the Systematic Evaluation of U.S. Army Command and Control Systems Using Neural Network and Discrete Event Computer Simulation. Unpublished Dissertation, Department of Industrial and Systems Engineering. Virginia Polytechnic Institute and State University, Blacksburg, VA.

Middlebrooks, S. E. (2004). Neural Net and Discrete Event Simulations of C2 Performance in Army Tocs. In Proceedings of the Human Factors and Ergonomics Society's 48th Annual Meeting, September 20-24, 2004 (pp. 513-517). New Orleans, LA: HFES.

Middlebrooks, S. E., Knapp, B. G., Barnette, B. D., Bird, C. A., Johnson, J. M., Kilduff, P. W., Schipani, S. P., Swoboda, J. C., Wojciechowski, J. Q., Tillman, B. W., Ensing, A. R., Archer, S. G., Archer, R. D., & Plott, B. M. (1999a). Cohost - Computer Simulation Models to Investigate Human Performance Task and Workload Conditions in a U.S. Army Heavy Maneuver Battalion Tactical Operations Center. In Proceedings of the HFES 43rd Annual Meeting (pp. 242-246). Houston, TX: Human Factors and Ergonomics Society.

Middlebrooks, S. E., Knapp, B. G., Barnette, B. D., Bird, C. A., Johnson, J. M., Kilduff, P. W., Schipani, S. P., Swoboda, J. C., Wojciechowski, J. Q., Tillman, B. W., Ensing, A. R., Archer, S. G., Archer, R. D., & Plott, B. M. (1999b). Cohost (Computer Modeling of Human Operator System Tasks) Computer Simulation Models to Investigate Human Performance Task and Workload Conditions in a U.S. Army Heavy Maneuver Battalion Tactical Operations Center. Human Research and Engineering Directorate, U.S. Army Research Laboratory, Aberdeen Proving Ground, MD 21005. ARL Technical Report ARL-TR-1994; DTIC ADA368587.

Middlebrooks, S. E., & Williges, R. C. (2002). Experimental Design Interrogation of Network Simulation Models of U.S. Army Command and Control Centers. In Proceedings of the 46th Annual Meeting of the Human Factors and Ergonomics Society (pp. 458-462). Baltimore, MD: HFES.

Najemnik, J., & Geisler, W. S. (2005). Optimal Eye Movement Strategies in Visual Search. Nature, 434, 387-391.

Plott, B. (2002). Software User's Manual For: C3TRACE (Command, Control and Communication - Techniques for Reliable Assessment of Concept Execution). U.S. Army Research Laboratory, Human Research and Engineering Directorate, Aberdeen Proving Ground, MD.

Plott, B., Quesada, S., Kilduff, P., Swoboda, J., & Allender, L. (2004). Using an Information-Driven Decision Making Human Performance Tool to Assess U.S. Army Command, Control, and Communication Issues. In Proceedings of the Human Factors and Ergonomics Society's 48th Annual Meeting, September 20-24, 2004 (pp. 2396-2400). New Orleans, LA: HFES.

Salas, E., Morgan, B. B., Jr., Glickman, A. S., Woodard, E. A., & Blaiwes, A. S. (1986). Measurement of Team Behaviors in a Navy Training Environment. Human Factors Division, Department of the Navy. Technical Report NTSC TR-86-014. DTIC ADA185237.

Schunk, D., & Plott, B. (2004). Use of Discrete Event Simulation to Model Human Performance. In Proceedings of the Human Factors and Ergonomics Society's 48th Annual Meeting, September 20-24, 2004 (pp. 2126-2130). New Orleans, LA: HFES.

Sondik, E. J. (1971). The Optimal Control of Partially Observable Markov Processes. Unpublished Dissertation, Stanford University, Stanford, California.

Tjan, B. S., & Braje, W. L. (1995). Human Efficiency for Recognizing 3-D Objects in Luminance Noise. Vision Research, 35(21), 3053-3069.

Tjan, B. S., & Legge, G. E. (1998). The Viewpoint Complexity of an Object Recognition Task. Vision Research, 15(16), 2335-2350.

Trommershäuser, J., & Gepshtein, S. (2004). Optimal Compensation for Changes in Effective Movement Variability. Nature Neuroscience Under Review.

Wojciechowski, J. Q., Plott, B., & Kilduff, P. (2005). Human Performance Model Development of a Battalion Tactical Operations Center. U.S. Army Research Laboratory, Human Research and Engineering Directorate. Aberdeen Proving Ground, MD. ARL Technical Report ARL-TR-2635.