Event Prediction for Modeling Mental Simulation in Naturalistic Decision Making

Dietmar Kunde
LTC, German Army
US Naval Postgraduate School
Monterey, CA
831- 656-2248
dkunde at nps.edu

Christian J. Darken, PhD
US Naval Postgraduate School
Monterey, CA
831- 656-2095
cjdarken at nps.edu

Keywords:
Mental Simulation, Naturalistic Decision Making, Event Prediction, Combat Simulation, Combat XXI.

ABSTRACT*: Nearly all armies of the Western Hemisphere use modeling and simulation tools as an essential part for analysis and training their leaders and war fighters. Tremendous resources have been applied to increase the level of fidelity and detail with which real combat units are represented in computer simulations. Current models digress from Lanchester equations used for modeling the big Cold War scenarios towards modeling of individual soldier capabilities and behavior in the post Cold War environment and increasingly important asymmetric warfare scenarios. Although improvements in computer technology support more and more detailed representations, human decision making is still far away from being automated in a realistic way. Many "decisions" within a simulation are based on rules and/or stochastic processes (qualified coin tossing) and hardly at all on cognitive processes. One cognitive model in naturalistic decision making is the Recognition Primed Decision Model developed by Klein and Associates. It describes how the actual process humans use to come up with decisions in certain situations is radically different from the traditional model of rational decision making. Mental Simulation is an essential part of this model in order to picture possible outcomes in the future for given courses of actions. This paper describes the current development of a computational model for mental simulation and the initial results of experiments conducted with a prototype in a combat simulation environment.*

1. Introduction

Running combat simulation models for training and analysis purposes is very time and personnel intensive. The low degree of artificial intelligence possessed by the constructed units in the simulation requires both extensive manual input of initial orders and human monitoring during the simulation run. The capabilities of autonomously acting units are very limited. Models of military decision making range from sophisticated methodologies, e.g., by comparing scored values of possible actions and taking the highest or the lowest value depending on circumstances (Norling et al, 2000), to less sophisticated methods, where units execute their initial orders according to an internal script – these are mainly "movement orders" – and react through simple logic to opposing fire or properties of the terrain or movement data. Their perception of the environment is restricted to that which is directly relevant to the application domain—for example, a simulated tank commander knows only some knowledge about tank combat, and nothing else. By contrast, human tank commanders have life experiences that may sometimes influence their decisions more strongly than domain knowledge (Forsythe, 2002). In many research prototypes of agents the learning capability has been addressed. However, in combat simulation models these issues have not been implemented to a satisfying degree. Therefore, the learning component lacking in simulated commanders precludes them from making complex decisions of human scale, and many decision points in the simulation have to be resolved externally and then put into the system, requiring much skilled assistance.

Logically, increasing the degree of AI should increase the cost-effectiveness of a simulation system during use. Fewer staff are needed for scenario input and system setup. During a run, the greater autonomy of the system leads to longer decision cycles before the units reach unreasonable or unacceptable conclusions. With enhanced AI, an assistant can control or monitor more units, which is especially valuable for the analytical application domain of modeling and simulation, for which there is usually a paucity of personnel. However, there is a drawback. Decisions made within the system by simulated commanders are not normally as good or of high quality as those from human commanders. This observation is valid not only vis-à-vis the ingenious decisions made by great generals or admirals, but to conventional and small-scale decisions as well. One of the differences between the performance of artificial and human commanders lies in the ability of humans to mentally simulate possible outcomes of their actions.

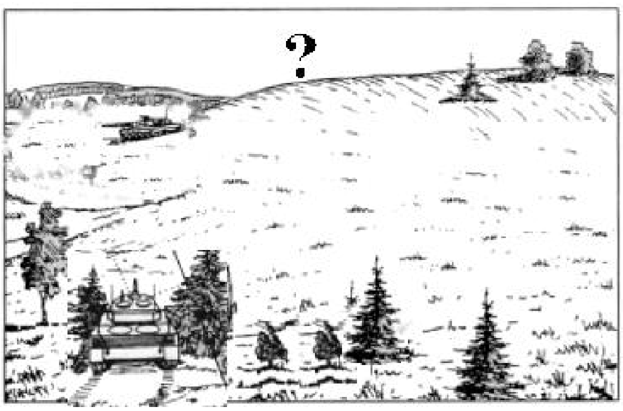

Figure 1. The simulated defender's perception: One tank is visible; therefore, he concludes that there is one tank in front of him. He can not deal with the question mark!

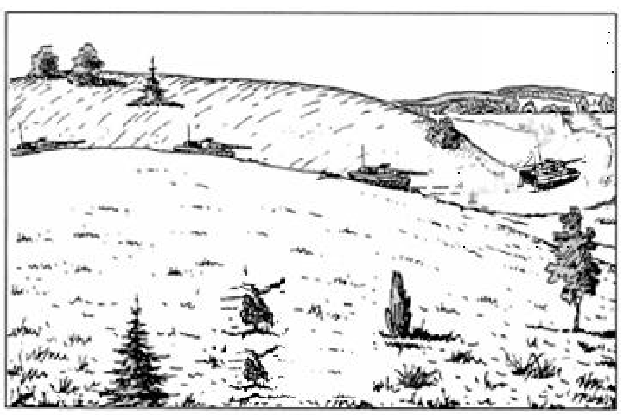

Figure 2. The same scene from the attacker's perspective. Actually there are four tanks coming around the corner. When the first tank gets shot the remaining tanks would behave differently in reality. In a simulation model normally they follow their scripted path.

The above example in Figure 1 and 2 illustrates this capacity. A platoon is defending a position with tanks.
Enemy tanks are expected to come around the hillside within firing range. A human platoon commander, having seen one enemy tank, would expect additional tanks and would therefore probably wait longer to begin surprise fire than would a simulated commander. If he fires before the other tanks round the corner, they will be warned and may try to outflank him, seek cover, use artillery fire, choose a different path, etc.

So the human platoon commander projects, or simulates, forward in time the possible consequences of his actions. Since mental simulation is beyond the present capability of simulated commanders, they might choose a different tactic, leading to a different outcome, unless overridden at particular decision points.

2. Why Naturalistic Decision Making?

Embedding mental-simulation capability in constructive units will contribute to the enhancement of AI and to overall economy and quality. Why can this be expected? Because, that is how humans think.

In nearly 20 years of empirical research, Klein has investigated the decision making of firefighters, pilots, nurses, military leaders, nuclear-power-plant operators and experts in a range of other domains (Klein, 1999). He developed a model that focuses on human strengths and capabilities that have not been modeled in classical decision theory. He described how commanders and leaders (or experts in general) are sometimes required to make urgent decisions under conditions of uncertainty. Thunholm also stated: "The study of military tactical planning and decision-making has shown that experienced commanders, quite contrary to what is prescribed by traditional military prescriptive planning models, make intuitive decisions based on recognition and mental simulation" (Thunholm 2000). In these situations they use recognition-based reasoning instead of the classical rational approach (Hutchins, 1996). That does not necessarily mean they decide irrationally in the original sense of the word; rather, they arrive at good decisions on a different path. The above references claim that we are not discussing a prescriptive theory about decision making, but a descriptive theory; that is, how humans actually perform decision making processes.

Recognition-primed decision making (RPD) is an established subfield in the domain of psychology. In the annual conferences since 1998 a large number of applications and advances in the field have been described (http://www.ndm7.org/). The attractiveness of the approach and degree of adaptation possible within the military is quite enormous. Klein has conducted approximately fifteen studies funded by the U.S. Army Research Institute and investigated decision making in a military environment (Klein, 2003). The Committee on Technology for Future Naval Forces stated in 1997 that the Navy should pursue an approach to joint-model development with a long-term view and an associated emphasis on flexibility. Especially with respect to technical attributes needed in joint models, decision models should represent the reasoning and behavior of commanders at different levels, naturally reflecting the actions, plans, and adaptations that commanders make in the course of operations (Committee on Technology for Future Naval Forces 1997). It is thus appropriate that the Naval Studies Board of the National Research Council at Washington DC has foreseen the advantage of mental simulation, as recommended in 2000: "The Department of the Navy may need to train commanders in recognizing patterns in typical cases and anomalies encountered in operations to improve their mental simulation skills and enable quicker and better decisions" (National Research Council, 2000).

Mental simulation is still a new field. Though several approaches have been undertaken (Sokolowski, 2002; Warwick et al, 2001), they all focused on issues of RPD in general, not explicitly on mental simulation. By contrast, this paper will focus on the mental simulation component of RPD. Mental simulation in this context can enhance decision making in three major areas: It improves the decision maker's diagnosis capability, it generates expectations to look for in the current situation and by evaluating the possible course(s) of action it helps in finding pitfalls (Klein, 1999).

3. The Modeling Design

3.1 Purpose

We build a decision model with the purpose of enhancing the cognitive capability of certain entities (intelligent software agents) used in combat simulation models. In this work we will use Combat XXI (see Chapter 4). The agents shall be enabled to

- predict next events based on statistical estimates,
- be sensitive to context of decisions,
- have an improved situational awareness,
- determine potential actions and to
- provide an explanatory component for the reasoning.

Figure 3 displays the role of Mental Simulation within situation awareness and decision. Since the combat simulation model is currently under development the decision model is external. Therefore, the analysis is done in a post-processing mode.

3.2 Components

The decision model consists basically of three major components.

- Situational Awareness Component
- Mental Simulation Component
- Decision Component

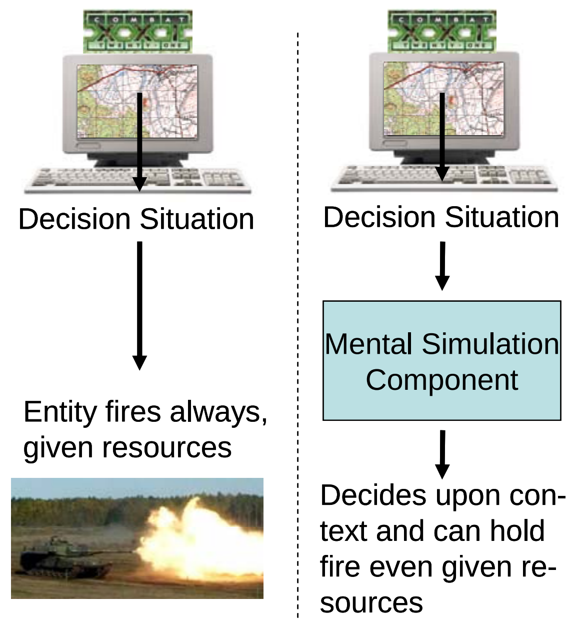

Figure 3. The role of Mental Simulation in the current work. At left, an entity fires when resources are available and always fires unless told otherwise. At right, the entity considers context, predicts the next event and fires accordingly.

3.2.1. Situational Awareness Component

Situational Awareness is a critical component in a decision environment. The better the awareness the more accurately all the parameters can be assessed that influence a decision situation. Good situational awareness is a prerequisite for a "good" and successful decision. The Situational Awareness Component represents the growing knowledge of the entity commander. It resembles a human cognitive picture of the battlefield situation. There is no data retrieval from ground truth. The output files from the combat simulation model yield data about detections and engagements on the battlefield plus associated data like time, location, shooter, targets, etc. The observations are ordered chronologically. Currently there are sensors "scattered" over the battlefield that detect red entities. The sensors are stationary and have no operational impact, which means they do not engage and they do not get engaged. The red movement is also not influenced by the sensors. These sensors yield the operational picture for the platoon commander that would be given by him via various enemy situation reports.

The model starts from scratch with respect to knowledge about the enemy formations. It has no pre-assumptions about the enemy behavior and formation. In a later step this could be supplemented with a “knowledge or experience database”. This receptive status allows the model to be flexible, since the commander cannot count on meeting with strict formations aligned according to old Warsaw Pact rules. Formations detected and categorized carry information about size, type, direction and speed. So far the formations are homogeneous. That means a forward artillery observer accompanying a combat unit is a distinct formation even if they operate together. This feature must be enhanced later. In the current model, the size of a formation is taken to be the number of distinct entities per formation that have been detected. Possible enemy objectives with respect to terrain such as seizing key terrain have not yet been represented. The situational awareness component yields the data used in the prediction model explained in the next paragraph.

3.2.2. Mental Simulation Component

The task of the Mental Simulation Component is to give an estimate of the next probable observation and the average (typically we use the median) time before this event happens. This is currently accomplished by utilizing a Markov Chain. This stochastic state machine assigns probabilities to the state transitions from state i to state j. These probabilities reflect the frequency of state transitions in the observations analyzed up to the current observation and normalized so that all the probabilities of emitting arcs in a particular state add up to 1. There are other possible models for this data; however, considering the current status of the combat model used, a finite state machine suited the available data the best. A state is defined as the number of entities detected at an observation time i. Each agent “tracks” the observations according to the state diagram in Figure 4.

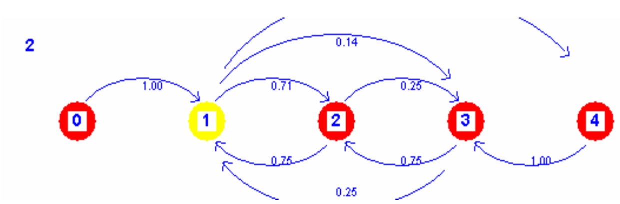

Figure 4 shows one of the state machines for a defending platoon that is currently in state “1”. A state indicates how many entities have been seen with the current observation. The arcs carry the transition probability to the next state. The median dwell times are listed separately.

State “1” means that the agent currently sees 1 entity and he stays in this state as long as he does not see another observation. If he then sees two tanks then he moves into the state “2” and the transition probabilities and mean dwell times (duration of staying in a state) are updated. Since the combat model does not currently provide data when an entity goes out of sight, the state “0” is never revisited. With the availability of “undetection” events the mean or median dwell times would be more realistic. So far the model is trained by having various sensors along the main approaches. Another setup could be to use comparable scenarios and then initialize the state machine with the probability and dwell time values.

An easy prediction criterion for the next transition could be to choose the arc with the highest probability. However, using this approach would not exercise transitions to states of lower likelihood. This would also be a very simple model of mental simulation that has the flavor of RPD. The degree to which humans take less likely outcomes into account when mentally simulating is, we believe, still a research question. In order to ensure that events with a lower probability can also be predicted the authors use a Monte Carlo simulation for sampling the values from the probability distribution as estimates. With Monte Carlo Simulation all kinds of questions can be addressed. This method is widely used when an analytically computation is very hard, even though the mathematical model is completely determined (Axtell, 2000). One class of problems of this type is referred to in the mathematical and statistical literature as boundary crossing (Giraudo, Sacerdote, Zucca; 2001). Many simulation runs can be conducted and then the frequency of occurrence, in the case of boundary crossing when the curve hit the line, can be taken as an estimate. Considering the current decision context there are mainly two questions that seem promising for the model. The first is related to the estimated time to expect a transition and the second addresses the multiple state sequences. The precise questions are:

- What is the most likely state sequence with the next x observations?
- When will a state (or set of states) of interest next be entered?

These mathematical formulations correspond to the human questions “What will happen next?” and “How long until (some anticipated event) occurs?”. We are still exploring how various questions of military interest can be best cast in these terms.

3.2.3. Decision Component

The Decision Component takes the prediction as the driving input for the decision. In this basic initial example of an agent for a tank platoon commander the tank will fire when it is likely that the next state will either have less targets or the mean transition time exceeds a threshold. That enables the platoon commander to make arguably realistic decisions. In existing models the immediate firing remains unpunished because the attacker also behaves inappropriately by ignoring the first shot or even a resulting kill and keeps following the scripted path. There is a maximum amount of time the agent will wait for more targets after the first tank is in view. If the enemy comes within a certain threshold distance then the tanks will fire anyway in order to avoid being shot first. This assumes that the blue tanks are in turret- or hull-down positions where the probability of detection is relatively small.

3.4. Process Flow

Figure 5 shows how the information or process flow works. Starting with a new observation the entities are assigned to the most probable formations (in this case the formations are platoons). Note that these formations are hypothetical; the agent does not have access to the ground truth.

The entities are detected and assigned to formations depending on the time and location they were spotted. The main assumption is that entities belonging to the same formation are not spread apart more than some threshold distance, currently taken to be 500 m. If an entity has not been seen for a while, then it is assigned to a new formation or reassigned to an old formation according to a space-time calculation. This ensures that for each detected entity, the model knows after a certain number of processed observations what other entities belong to that formation and where and when they have been detected the last time. The situational picture is updated with each new observation until a decision is required.

A decision situation is invoked when a blue tank receives a red tank detection message. Then the blue tank commander has to consider whether he will engage immediately or wait. All the observations before this event were made by sensors that have no operational impact on the movement of the red forces. They are used exclusively for situational awareness. Given a decision situation, all related information is revealed by the situational awareness component. That means, e.g., how many distinct entities of the formation have been detected in total, and where any entities that are currently not observed were reported for the last time.

The prediction is made based on knowledge from the state diagram. Given the chosen prediction method it is either a look-up for the most likely transition or a Monte Carlo simulation. Independent of the method chosen, the prediction will give an estimate of what state can be expected next and with what probability and expected wait time.

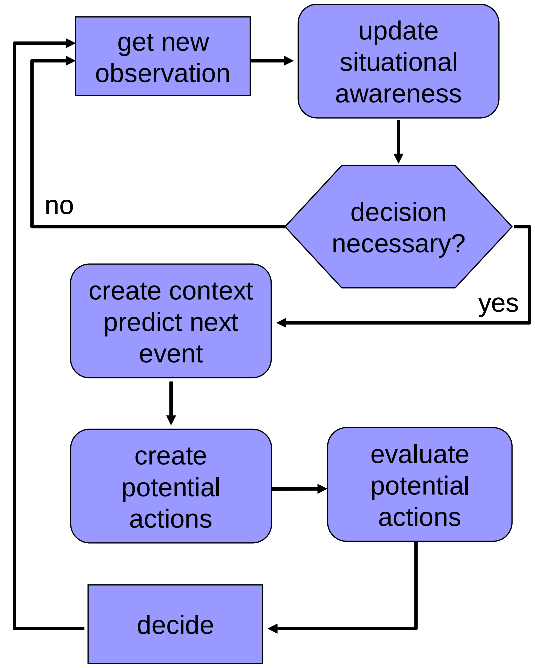

Figure 5. Process Flow. The Situation Awareness Component gets updated one observation at a time until a decision situation is encountered. Then the mental simulation component predicts the next event and evaluates the potential action. Based on the prediction and evaluation a decision is made.

The potential actions that have been implemented so far are firing right away or after the red tank has taken on a certain combination of distance, speed and direction. This combination is a threshold whose crossing will cause the blue tank to decide to fire regardless of the predicted event. The decision based on the prediction has not been implemented yet. Currently the decision does not impact the simulation during a run since the combat model and decision model have not been coupled together.

4. The Combat XXI Model

Combined Arms Analysis Tool for the 21st Century (CXXI) is a high-resolution, closed-form, stochastic, analytical combat simulation. CXXI is being developed by the TRADOC Analysis Center – White Sands Missile Range (TRAC-WSMR) and the Marine Corps Combat Development Command (MCCDC). CXXI will be used for the analysis of land and amphibious warfare in the Research, Development and Acquisition (RDA**)** and Advanced Concepts and Requirements (ACR) Modeling and Simulation (M&S) domains. It is the next-generation analytic tool for modeling theater-level ground warfare for the Army and Marine Corps.

CXXI is an information-driven model. Physical algorithms, primitive behaviors and tactical behaviors are represented by the model's computer code. However, it is the processing of user-defined information, from the study scenario and the information in the model's database, which drives the interactions between entities in the simulation.

The model's unit of resolution is a platform-level entity (tank, aircraft, dismounted soldier, etc.). CXXI models primitive behaviors such as movement, search, and engagement in separate modules. Tactical behaviors (bounding over-watch, close-air support, etc.) are represented in decision-making modules. Each entity is assigned a set of modules to describe its physical behavior and decision-making capabilities (TRAC-WSMR, 2004).

The current perceptual model used consists of event messages that are native to CXXI. This information is stored in several output log-files. They yield information about detections with associated data regarding observer and target IDs, coordinates, distances, sensors detected and time. Other log files deliver engagements, kills, and damages and detonations.

For the sake of completeness it should be mentioned that Combat XXI has several built-in engagement levels. Depending on the accuracy of detection with respect to target type (detect, aim point, recognition and identification) the model fires after a certain level of target identification is achieved. This also leads to a delay in firing in the case that the respective accuracy of detection level has not been met. However, this is unrelated to this work, because this delay is not the result of the assessment of the agent's own behavior. If the target identification level increases and meets the threshold of engagement then firing will commence immediately.

5. First Experiments and Initial Results

The scenario modeled in this paper is shown in Figure 6. The initial scenario demonstrates that the allocation of observed entities to formations works correctly and displays the decision context based on the observations so far. It also shows the distribution of the number of tanks per observation until the tank encounter occurs. Currently three red tank formations are following scripted paths southbound. They represent three platoons of a red tank company. A decision point is reached when a blue tank entity detects one or more red tank entities. The context for the decision point is shown in Figure 7. It displays the entities involved in the current observation in dark colors ❶. The faded entities indicate the other current blue tank positions ❷ and the red positions of the known tanks with the last observed location ❸. The positions are scaled and the green circle indicates a 500 m distance. ❹ indicates that so far 4 tanks have been identified that belong to that

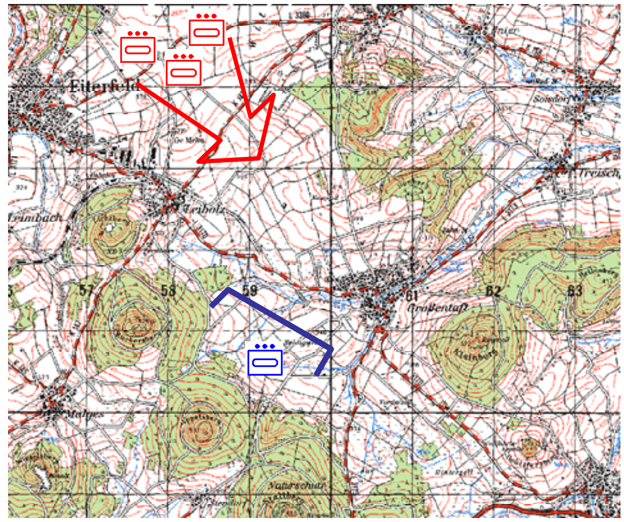

Figure 6 shows the initial scenario. A red tank company is represented by three tank platoons. Their objective is somewhere south east of the blue defending platoon.

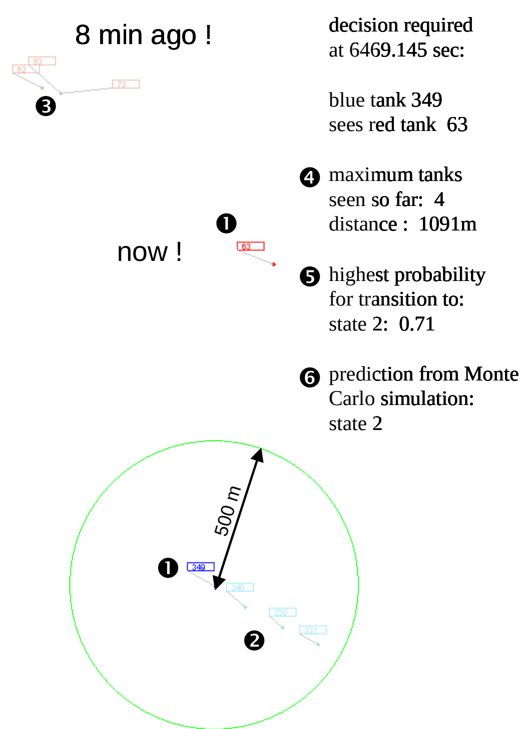

Figure 7 shows the first decision situation in the scenario. The blue tank (349) sees one red tank (63). In the top left corner are the other tanks with their last reported location. The data field on the right displays the current information set for this decision situation.

formation . The current distance is 1091m. Applying the highest probability from the state machine as estimate for the transition to the next state would predict state 2 ❺. Using the Monte Carlo simulation method predicts as the probable transition also state 2 ❻. The data display indicates that the tank should not fire. There is an 71% chance that within the next observation cycle there may be one more tank in sight. Therefore the decision should be to wait. On the other hand if the tank waits too long the risk increases of being shot first. Currently a combination of speed, distance and direction defines a threshold whose crossing will lead to fire regardless of further predictions.

Figure 8 shows the actual state sequence after the first decision situation. It cycles 4 times between state 1 and 2 (which have the also the highest transition probabilities) and moves then to state 3 and back to 2. After this there are no more observations therefore it does not go back to state 1 or 0. In this experiment the engagement option in Combat XXI has been turned off. Otherwise there would be only one or two observations after the first detection due to the destruction of the red tanks.

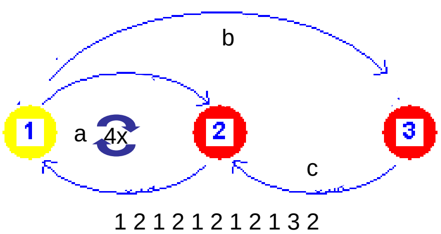

Figure 8 shows the actual state transitions. The numbers below the state machine indicate the state sequence numerically after the first decision situation.

Table 1 shows the initial results obtained from the scenario in Figure 6. The time indicates seconds after simulation start. Column 2 displays the state the blue platoon goes to when observing the number of tanks visible at that point in the simulation. Column three indicates the next three actual transitions from the current observation time. At each decision point the Monte Carlo simulation runs three transitions ahead 100 times. The mode (most common single outcome) of the 100 runs and the associated probability are denoted in column 4 and 5. The probability of the continuation of the real state sequence is calculated based on the transition probabilities starting from the current time, and therefore they vary depending on the observations so far. They are displayed in column 6. The last column displays the probability for the real state sequence according to the Monte Carlo simulation.

Time at Decision Point	Current-State	Real Sequence Ahead	Mode from the 100 Simulations	Prob of Mode	Prob of real Sequence from State Machine	Prob of real Sequence from Simulation
6469.145	1	1212	1212	0.43	0.38	0.43
6827.905	2	2121	2121	0.52	0.42	0.52
6840.526	1	1212	1212	0.42	0.44	0.42
7051.431	2	2121	2121	0.47	0.55	0.47
7076.572	1	1212	1212	0.54	0.57	0.54
7095.882	2	2121	2121	0.52	0.57	0.52
7113.815	1	1213	1212	0.72	0.06	0.02
7134.462	2	2132	2121	0.58	0.05	0.07
7155.963	1	1321	1212	0.63	0.05	0.03
7165.961	3	321	3212	0.52	0.65	-
7176.971	2	21	2121	0.58	0.86	-

Table 1 shows the analyzed data for the decision points. In case of the first row at 6469.145 one blue tank sees one red tank. The real sequence (hidden from the agent) shows what the next real states will be. Running the Monte Carlo simulation yields an estimate for the next state the tank commander can expect. In this case the prediction matches the actual outcome. The last two columns display the probability values for the real sequence obtained from the state machine and from the Monte Carlo simulation respectively.

In the first 6 cases the predicted sequences match the actual observations. The probability values estimated with the Monte Carlo simulation indicate a good match. The probability values become very small after 7113.815 sec. This looks bad on the first glance, but actually reflects what we should expect the model to provide. The transition from state 1 to state 3 has a very low probability ($p_{13} = 0.07$). Due to the low frequency of occurrence the simulation does not hit this arc very often. Therefore, the small probability values show that the model is working correctly. In the last two rows the real sequence values shrinks due to encountering the end of the observation stream.

6. Conclusion

In many combat models insufficient representation of human behavior still cancels out because it occurs on the blue and on the red side as well. However, with the need of more sophisticated representation of human behavior, in order to represent the more sophisticated combat situations, it is mandatory to base the decisions in the system on more accurate entity representation. This first approach to the computational modeling of mental simulation is far from being perfect or comprehensive. However, we believe that the initial scenario already depicts a promising way of improving the cognitive capabilities of constructive forces. The conducted experiment still lacks sufficient accuracy in the prediction with respect to time mainly due to data that is not currently provided by the simulation platform, but it also shows the benefits by enabling the entities to anticipate certain behavior of the enemy. This will enable pro-active behavior not seen so far in combat models. With further research and implementation the representation of constructive forces by mental simulation will certainly improve and meet the requirements for simulation models of the next generation.

Acknowledgment

The authors like to thank the people at TRAC White Sands Missile Range and TRAC Monterey for their support in using Combat XXI.

7. References

Axtell, Robert., (2000). Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences. Working paper. The Brookings Institution. Available at www.brookings.edu/es/dynamics/papers.

Committee on Technology for Future Naval Forces, National Research Council Technology for the United States Navy and Marine Corps, 2000-2035 Becoming a 21st-Century Force: Volume 9: Modeling and Simulation. (1997). National Academy Press Washington, D.C.

Forsythe, C. & Xavier, P. (2002). Human emulation: Progress toward realistic synthetic human agents. Proceedings of the 11th Conference on Computer-Generated Forces and Behavior Representation, Orlando, FL. 257-266.

Giraudo M.T.; Sacerdote L.; Zucca C. (2001). A Monte Carlo Method for the Simulation of First Passage Times of Diffusion Processes - Methodology And Computing. In Applied Probability, June 2001, vol. 3, iss. 2, pp. 215-231 (17) Kluwer Academic Publishers, Boston, U.S.A

Hutchins, S. G. (1996). "Principles for Intelligent Decision Aiding." Naval Command, Control, and Ocean Surveillance Center, RDT&E Division. Technical Report 1718. San Diego, CA.

Klein Associates Inc. (2003). KA PUBS 2003. URL= http://www.decisionmaking.com/approach/publications.html

Klein, G. (1999). Sources of Power: How People Make Decisions. Cambridge MA, MIT Press.

National Research Council Washington DC Naval Studies Board (2000). Network-Centric Naval Forces: A Transition Strategy for Enhancing Operational Capabilities. Commission on Physical Sciences, Mathematics, and Applications (CPSMA): p. 56.

Norling, E., L. Sonenberg, and R. RÄonnquist. (2000). Enhancing Multi-Agent Based Simulation with Human-Like Decision Making Strategies. In Multi-Agent Based Simulation: Proceedings of the Second International Workshop, MABS 2000, Boston, MA, Springer, pp 214 - 228.

Sokolowski, J.A. (2002). "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, Orlando, FL., May 7-9 2002. pp. 473-478.

Thunholm, P., (2000). ETT STEG MOT EN FÖRESKRIVANDE MODELL FÖR MILITÄRT – TAKTISKT BESLUTSFATTANDE. 5th Conference on Naturalistic Decision - Making, Tammsvik, Sweden May 26 28, 2000.

TRAC-WSMR. (2004) COMBAT XXI User's Guide v4.0. White Sands Missile Range, NM

Warwick, W., S. McIlwaine, R. Hutton, and P. McDermott. (2001). Developing Computational Models of Recognition-Primed Decision Making. In Proceedings of the Tenth Conference on Computer Generated Forces, May 15-17, Norfolk, VA, pp 323-331.

Author Biographies

Dietmar Kunde, LTC (GE) is a PhD-Candidate at the US Naval Postgraduate School in Monterey, California. He received a Master of Science in Operations Research from the Naval Postgraduate School at Monterey, CA and a Diploma in Surveying from the University of Federal Armed Forces in Munich, Germany. Before he started his Doctorate he worked for five years in the German Army Office responsible for the conceptual development of Decision Support Tools for the German Army.

Christian J. Darken, PhD, is an Associate Professor of Computer Science at the Naval Postgraduate School, where he also collaborates intensively with the MOVES Institute. Previously he was Project Manager of the Decision Support Systems project and Senior Member of Technical Staff at Siemens Corporate Research in Princeton, NJ, where he was variously associated with the Learning Systems, Adaptive Information and Signal Processing, and Software Engineering Departments. He was also a programmer of one of the first commercial first-person perspective massively-multiplayer games. He received his Ph.D. in Electrical Engineering from Yale University in 1993, and previously received the M.S. and M. Phil. in Physics from the same institution.