JUN 5 2000

SANDIA REPORT

SAND2000-0974
Unlimited Release
Printed May 2000

A Naturalistic Decision Making Model for Simulated Human Combatants

Keith O. Hunter, William E. Hart, and Chris Forsythe

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of
Energy under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

RECEIVED
JUN 20 2000
OSTI

Issued by Sandia National Laboratories, operated for the United States
Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.doe.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/ordering.htm>

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

SAND2000-0974
Unlimited Release
Printed May 2000

A Naturalistic Decision Making Model for Simulated Human Combatants

Keith O. Hunter and William E. Hart
Optimization and Uncertainty Estimation
kohunte@sandia.gov and wehart@sandia.gov

Sandia National Laboratories
PO Box 5800
Albuquerque, NM 87185-1110

Chris Forsythe
Statistics and Human Factors
jcforsy@sandia.gov

Sandia National Laboratories
PO Box 5800
Albuquerque, NM 87185-0829

Abstract

We describe a naturalistic behavioral model for the simulation of small unit combat. This model, Klein's recognition-primed decision making (RPD) model[1], is driven by situational awareness rather than a rational process of selecting from a set of action options. We argue that simulated combatants modeled with RPD will have more flexible and realistic responses to a broad range of small-scale combat scenarios. Furthermore, we note that the predictability of a simulation using an RPD framework can be easily controlled to provide multiple evaluations of a given combat scenario. Finally, we discuss computational issues for building an RPD-based behavior engine for fully automated combatants in small conflict scenarios, which are being investigated within Sandia's Next Generation Site Security project.

1. Introduction

Computer generated forces (CGF) are computer representations of combat entities that are used in simulation environments. The ultimate goal of CGF is to achieve realism while minimizing human intervention in the simulation process. Whether the simulation stands alone, links with other simulators, or interacts with humans, the accurate modeling of behavior is a significant challenge [1, 3, 4].

A number of approaches to modeling human behavior in combat simulation exist today, most of which can be criticized for lacking flexible, autonomous behavior [2, 3, 4]. Interaction and control on the part of subject matter experts (SMEs) plays a large role in some approaches to behavior simulation (eg. ModSAF). The human supervised approach has the obvious drawback of requiring a human to run the simulation, but it may also be criticized for its inherent limitation to that human's own knowledge, experience, and skill in incorporating desired factors into the simulation. Another approach is to rely upon the processing of prepackaged rules (eg. CCTT SAF). Like the human-supervised approach, rule based knowledge representation has had significant success in addressing the needs of the simulation community. However, systems that draw heavily upon a static set of preset rules lack any underlying model of human behavior [2]. Responses to situations are limited to what has been incorporated in the rule sets, and are difficult to make realistically flexible when novel situations are encountered. Rule based systems receiving rule sets from SME-run simulations have the additional drawback of being rather predictable. Any group of SMEs may have biases in terms of tactics or situations deemed to be important. These biases can easily be reflected in the simulation and enable an unacceptable level of

predictability in a training system or result in a scenario model that omits certain variations on known situations as well as novel situations.

New approaches to human behavior modeling for CGF need to do a better job of modeling the human cognitive functions for assessing and recognizing situations, drawing upon past experience, and considering novel solutions. Naturalistic Decision Making (NDM) models accommodate these. Klein [1] cites four characteristics of decision-making that NDM is designed to address:

1. Tasks and settings present poorly structured problems, uncertain and dynamic environments, competing or incompletely defined goals, action/feedback loops, time pressure, high stakes, multiple players, and organizational goals.
2. The entities being modeled possess experience.
3. Desirable features of the model include situation awareness, diagnosis, and plan generation rather than fixation on the moment of choice between options.
4. The purpose of the approach is to describe the strategies people use rather than prescribing the strategies they might be expected to use.

Modeling the way reasonably experienced people actually make judgements and decisions in realistic settings is the goal of NDM. Klein's four characteristics are all important to modeling human behavior in small-unit combat situations. From the perspective of the individual combatant being modeled, the environment is quite dynamic and requires quick recognition of local situations and

identification of actions that match those situations based upon experience or training. It is also desirable for these simulated combatants to have the capacity to generate short term local plans in support of the broader strategic plan being followed by all team members. The ability to realistically cope with novel situations is also important if the simulation is to remain accurate under all possible simulated conditions. This emphasizes the usefulness of the ability to mentally simulate actions under consideration when a situation is not recognized. We consider the NDM model known as Recognition-Primed Decision Making (RPD).

In this paper we describe the advantages offered by RPD for simulation projects that require the realistic modeling of human behavior in small-scale combat. First we define RPD in terms of its characteristics and briefly discuss options for its implementation. Finally, we present some of the implementation challenges associated with RPD and identify directions the implementation of RPD should take next.

2. Recognition-Primed Decision Making

Over the past decade, a relatively rich literature has developed describing the process by which decisions are made in naturalistic settings. A variety of domains have been studied (e.g., firefighters, commercial airline and military pilots, surgeons, military commanders, corporate managers, etc.) and similar findings have emerged for each domain. When faced with a novel problem, expert decision-makers focus their attention on comprehending the situation. Using critical cues, the objective is to attain a match between the current situation and past situations. Very little time is devoted to developing and evaluating alternative strategies. Instead, resources are committed to understanding the situation and once understood, the course of action is relatively clear. This tends to be a very robust finding and is supported by research suggesting that in long-

term memory, expert knowledge is pattern-indexed in relation to domain-specific tasks. The RPD model describes these kinds of situations, incorporating three levels of decision making.

2.1 Overview

Figure 1 illustrates the first level of decision making in RPD. A situation is experienced and the products of this recognition are expectations, knowledge of relevant cues, plausible goals and typical actions. Implicit within this recognition is an appropriate course of action for the current situation. Level 1 recognition is illustrated by the following example. Driving on a multi-lane expressway, a driver notices an illuminated indicator light on their dashboard. The initial response is to assess the situation. Does the indicator provide a word or icon communicating the nature of the problem? Are any of the gauges in the red zone? Are there noises or other abnormal cues emanating from the engine? How close is the nearest shoulder? How is traffic? Assume this assessment leads to recognition that the engine is overheating. Recognition of this situation implies certain expectations. The engine may catch on fire. It indicates relevant cues: the temperature gauge, the presence of steam or smoke. With the engine still running, moving to the shoulder becomes a plausible goal and a typical action might be to activate the hazard lights. In summary, once the driver recognizes that there could be a problem, attention is focused on situation assessment.

The second level of decision-making, Diagnose the Situation, is shown in Figure 2. This level captures an important characteristic distinguishing expert and novice decision-makers. During situation assessment, experts are adept at realizing when they do not have sufficient information to adequately assess a situation. The situation is diagnosed using techniques such as feature matching and story

building, each of which require that more information be extracted from the situation. Experts are also adept at recognizing anomalies between current and past situations.

The third level of decision-making, Evaluate Course of Action, occurs when there is uncertainty concerning a course of action. It

should not be inferred that the decision-maker compares alternative courses of action. Instead, as illustrated in Figure 3, a course of action is evaluated using mental simulation to confirm that it will work. If there is doubt, the course of action is modified, or when necessary, the situation is reassessed.

Figure 1 Simple Match (Level 1 Recognition) [9]

Figure 2. Diagnose Situation (Level 2 Recognition) [9]

Figure 3. Evaluate Course of Action (Level 3 Recognition) [9]

2.2 Evaluation

As we noted earlier, existing simulations place an excessive reliance on doctrine and consequently fail to exhibit sufficient flexibility or adaptability. Features of the RPD model enable development of a simulation that does not suffer these shortcomings. Instead of doctrine, the recognition of situations, as well as knowledge derived from this recognition, may be based on the experience of subject-matter experts with past situations. Furthermore, through machine learning approaches such as genetic programming [5], an RPD implementation can continually learn and adapt as experiences are accumulated. Thus RPD provides a highly flexible simulation that readily adapts in response to varying situational factors.

The RPD model provides an empirically valid foundation upon which representations of human behavior may be based [1]. Combat scenarios involving humans in tactical roles acting as a group working toward a strategic goal squarely fit within the themes Klein set out for the RPD model. Short distances and battle duration contribute to both high stakes and time pressure, and local decision imperatives can conflict with mission goals. Close combat of this nature represents a case in which real human combatants typically do not have time to perform extensive analysis of their circumstances. Instead experience is used both to recognize aspects of the situation that have been dealt with before as well as call forward the course of action that succeeded in that past instance.

Another central attribute of RPD is the mental simulation of various options in order to identify whether or not their outcome is desirable. Through this mechanism the behavioral model supports the development of solutions for only partially recognized problems as well as novel approaches to situations that are entirely new. This kind of situation-driven consideration of action consequences differs considerably from

merely selecting actions from a list of options or decomposing problems into basic elements that are then processed to produce an action.

Finally, it is desirable for simulations to account for psychological factors leading to variability in decision-making (e.g., stress, fatigue, fear) [2, 8]. RPD does not explicitly address such factors, but it does offer empirically sound mechanisms for representing their effects. For instance, factors causing decrements in performance, such as fatigue, could be simulated through the misinterpretation of cues, failure to detect cues, or mistaken recognition of situations. Each of these manifestations is consistent with reported observations.

3. RPD for Simulation

Figure 4 describes how the three essential elements of an RPD-based simulation fit together. There must be a representation of the environmental cues required for situation assessment, a pattern matching mechanism that allows situations to be recognized on the basis of environmental cues, and a collection of situations that embody the knowledge gained through accumulated experience. The following sections describe these aspects as they relate to agent-based combat simulation.

3.1 Environmental Cues. Agents simulating individual members of opposing teams may be modeled as continuously processing environmental data. At each time unit, they assess either the entire suite of environmental cues or a designated subset. Situations are recognized on the basis of patterns present within combinations of these cues. Environmental cues can include variables explicitly represented within the simulation (e.g., x, y, z coordinates of agents) and variables derived through computational processes. Agents continually monitor environmental cues

using pattern matching to recognize the situation, as well as changes to the situation.

3.2 Pattern Matching. Recognition of situations is a pattern matching process. Each known situation has a pattern of environmental cues associated with it. Pattern matching involves a continuous process of matching environmental cues with known situations. Where environmental cues form a pattern resembling that of a known situation, agents behave in accordance with that situation. Ambiguous patterns of environmental cues are anticipated. Consequently, it is necessary to develop mechanisms that enable situations to be inferred on the basis of ambiguous environmental cues. This process would be akin to undertaking Level 2 recognition wherein there is insufficient familiarity with the pattern of environmental cues to recognize the situation through Level 1 Recognition.

3.3 Collection of Situations. The recognition of situations described in the preceding sections implies formation of an association between immediate experience and one of a collection of cognitive schema, each the codification of accumulated knowledge gained from a set of similar past experiences. These cognitive schemas are

described here as "Situations." They are knowledge structures developed over time and continually refined in accordance with new experiences. The collection of situations does not represent a complete memory of previous experiences. Situations and their associated schemas are analogous to that portion of human experience that supports action based upon recognition of previously encountered situations. Close-combat tactics provide a basis for specifying a collection of situations. Tactics can be identified that are appropriate to our scenarios. These tactics are being associated with a pattern of environmental cues that specify their appropriate conditions. When agents recognize the pattern of environmental cues associated with a given tactic, it may be inferred that the situation is one in which that tactic offers an appropriate course of action.

4. Implementation Challenges We now turn our attention to discussing implementation issues. In order to realize the full benefit of a human behavioral model within an intelligent simulator, the software must be structured and the behavioral model should incorporate machine learning. These concerns lead to some significant challenges, which we discuss in the following sections.

Figure 4. The Three Elements Essential to an RPD-Based Simulation

4.1 Modularity, Flexibility, and Scalability

Modularity, flexibility, and scalability are important concerns in simulator implementation as they are in many software efforts. Modular components not only support software engineering but also contribute significantly to the capacity to verify successful implementation of key components of the RPD model. For example, different tasks associated with expected behaviors can be learned separately using different techniques. This modular independence facilitates the independent development and analysis of finer-grained aspects of the machine intelligence (e.g. holding navigation component constant while refining the shooting behavior, or testing the automatic recognition of situations). Flexibility in applying domain knowledge to the simulation is

also highly desirable. Using a layered design, any subtask module developed in software can be replaced by a hard-coded driver or even human interaction as automated components are developed or improved. Finally, scalability is a key concern because simulation fidelity needs may change or additional kinds of behavior may be desired. Hence, we advocate a design that does not assume small numbers of simulated combatants, strict numbers of environmental cues, or a set number of situations to be recognized rather than improvised. Expanding the suite of learned behaviors is made less difficult by the kind of architecture described here, since new modules can be developed independently and incorporated into the existing learning structure at the appropriate level. These factors are important to the implementation of RPD since some

environmental cues may best be processed differently or separately from others for use in situation recognition and decision-making. We also expect to benefit from the ability to isolate and combine environmental cues as we refine the models that drive pattern matching for situation recognition.

4.2 Machine Learning

Multiple learning approaches can be used to correlate environmental cues with realistic agent actions. Genetic programming [5], neural networks [6], and reinforcement learning [7] are among many possible approaches to learning agent behavior within the framework provided by RPD. For example, a dedicated simulator can be used to learn scenario recognition by using multiple iterations of training scenarios for the learning methods of choice. Hybrid learning approaches can be used along with combinatorial optimization algorithms as a part of an agent's overall toolkit. Even for opposing teams that have qualitatively different objectives within the simulation, the learning tasks involved may overlap significantly. Initially it may be appropriate to develop a single team's behaviors within manually controlled training scenarios. Subsequent training of the opposition could then utilize those behaviors in an alternating fashion.

4.3 Prediction

Development and exploitation of predictive capabilities will be very important to the implementation of RPD. An example of predictive capability is the projection of an opposing agent's location at some future point in simulated time based upon that agent's current speed and heading or based upon that agent's perceived intent. The modular development of this part of the architecture not only supports the independent development of predictive mechanisms but also enables the parallel development of agent capability to use predictive information. With respect to the RPD model, this capability will be of particular

importance in Level 3 recognition where a major feature is the mental simulation of actions under consideration in response to unrecognized situations.

4.4 Independence of Agents

While agents may need to act as members of groups and have their actions coordinated by tactical or strategic modules, agents act independently within the constraints of the roles they fill within the overall tactical implementation. Because roles can be reassigned, agent behavior may display varying degrees of independence. An isolated agent may take action based primarily on its own local situation and the requirements of its current tactical role. An agent that is assigned to a leading role within a group may also have its actions tightly coupled with the actions of other group members or use communications received from them. Supporting the full realistic range of agent independence will demand appropriate communication between agents and higher-level situational variables to accompany the low-level environmental cues described earlier.

4.5 Situational Variables

Deciding upon a set of situational variables and relevant cues is an important step in this and similar approaches to simulation [4]. The development and use of situational variables is key to the function of multiple components within RPD implementation. One area of concern is that of increasing problem complexity and, subsequently, computational cost. The machine learning is an off-line task, but the development of behaviors that rely upon costly real-time operations could result in unacceptable simulator performance. Similarly, the proliferation of decision criteria can also impede the success of learning. This proliferation may not only increase the amount of logic a learning algorithm must discover but may also negatively impact robustness of the learned behavior. These issues should not

discourage effort in developing and exploiting situational variables. Instead, these issues suggest that the careful consideration of variable count and variable interrelationships is merited as this model is implemented and tested.

4.6 Coordinating Levels of Decision Making

Providing training input that allows a learning algorithm to develop the capability of switching between recognition modes and their attendant activities is a nontrivial task. RPD's level 2 decision making involves further diagnosis of situational data when insufficient information is present. Given the vast range of situations possible within realistic situations, endless cycles of reevaluation can develop or inappropriate choices to deliberate may be made when immediate action is essential. Level 3 recognition requires the agent to evaluate the expected outcome of a course of action prior to carrying it out. It is unclear what the time penalty for such action should be or how accurately the agent should be able to project hypothetical outcomes. We expect that RPD's first level of decision-making, the relatively straightforward response to environmental cues, may be the least difficult level to enable. Pattern matching capability is critical in supporting transitions between levels of recognition.

4.7 Coordinating Tactical and Strategic Goals

Enabling the definition and pursuit of both tactical and strategic goals is another objective of RPD-based simulation. To accomplish this, agents need to learn how to recognize their local situation and take actions leading to the achievement of multiple goals that may be nontrivially related. For example, surviving the onslaught of an opponent may require a change in movement or activity that conflicts with the strict definition of an agent's role within a tactic. Also, agents may need to switch roles over the course of a simulation. Similarly,

tactics may need to shift temporarily or permanently during the implementation of a broader strategy due to attrition or some other change in a local situation or change in available tactical resources.

5. Conclusions

We have described RPD as a human behavioral model that supports key aspects of effective small combat simulation. Mechanisms built around the components of this model can provide behavior that more resembles that of real humans because they use situational cues the way humans do to drive decisions. RPD is a broadly endorsed cognitive model, and we have begun to identify an appropriate architecture for representing it in computer software.

Using current methodologies, we are working toward the goal of representing RPD's levels of recognition. In order to extend this work, several aspects of the approach must be addressed. Our consideration of the current effort in autonomous agent research and in CGF suggests that we investigate the following:

- Identification of relevant situation cues.
- Representation of situations for recognition.
- Automatic identification and representation of tactical and strategic goals for agents.
- Integration of RPD's levels of recognition within group and individual behaviors.

This list of obstacles is clearly not exhaustive, but it does illustrate the nature of the model's initial challenges.

6. Acknowledgements

We thank Ron Hightower, Nick Nicholson, Sabina Jordan and Mark Snell for their helpful input. This work was performed at Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed-Martin Company, for

the United States Department of Energy under Contract DE-AC04-94AL85000.

7. References

- [1] Klein Gary et. al.: *Naturalistic decision making and related research lines*. In Klein Gary and Zsombok C.E. (Eds), *Naturalistic Decision Making*, Lawrence Erlbaum Associates, Inc., Mahwah, New Jersey, pp. 29-35. 1997.
- [2] Klein Gary and Crandall Beth: *Recognition-Primed Decision Strategies*. U.S. Department of Commerce. Report ADA309570, March 1996.
- [3] Pew Richard, W. and Mavor, Anne S. (Eds): *Modeling Human and Organizational Behavior*. National Research Council, National Academy Press, Washington, D.C. 1998.
- [4] Booker, Lashon B.: *Motivational Control of Tactical Behaviors: Interim Results*. Proceedings of the Eighth Conference on Computer Generated Forces and Behavior Representation. 1999.
- [5] Koza, J.: *Genetic Programming: On Programming Computers by Means of Natural Selection and Genetics*. MIT Press, Cambridge, MA, 1992.
- [6] Rumelhart, David E., Hinton, Geoffrey E., and Williams, Ronald J.: *Learning Internal Representations by Error Propagation*. In Rumelhart, David E. (Ed) *Parallel Distributed Processing*, MIT Press, Cambridge, MA, pp. 318-362, 1986.
- [7] Sutton, R. S.: *Temporal Credit Assignment in Reinforcement Learning*, Ph.D. Thesis,

University of Massachusetts, Amherst, MA, 1984.

- [8] Zimm, Alan D.: *Modeling Maneuver Warfare: Incorporating Human Factors and Decision Making in Combat Modeling*. Proceedings of the Eighth Conference on Computer Generated Forces and Behavioral Representation, 1999.
- [9] Klein, Gary: *The Recognition-Primed Decision (RPD) Model: Looking Back, Looking Forward*. In Klein Gary and Zsombok C.E. (Eds), *Naturalistic Decision Making*, Lawrence Erlbaum Associates, Inc., Mahwah, New Jersey, pp. 285-292. 1997.

Author Biographies

KEITH HUNTER is a technical staff member at Sandia National Laboratories working on global optimization and agent-based modeling as a part of the Optimization and Uncertainty Estimation Department. He received his Master of Science in Computer Science from the University of Central Florida in 1997.

WILLIAM HART is a Principal Member of Technical Staff at Sandia National Laboratories in the Optimization and Uncertainty Estimation Department. He develops and applies global optimization in support of a broad range of problems including agent-based learning. Dr. Hart received his Ph.D. in Computer Science from University of California at San Diego in 1994.

CHRIS FORSYTHE is a technical staff member working with the Statistics and Human Factors Department at Sandia National Laboratories. He holds a Ph.D. in Psychology from Memphis State University.

Distribution:

1 MS 9201 Howard Hirano, 16000
1 0839 Elaine Raybourn, 16000
1 1109 Rich Pryor, 09201
1 1153 Stewart Cameron, 15336
1 9202 Rob Rinne, 08100
1 0839 Gerry Yonas, 16000
1 1221 Russ Skocypec, 15002
1 0185 Dorothy Stermer, 01317
1 1188 John Wagner, 15334
1 0428 Dave Carlson, 12300
1 0747 Allen Camp, 06410
1 6413 John Forester, 0748
1 0746 Dwight Miller, 06411
1 9011 Nina Berry, 08920
1 1137 John Ganter, 06535
1 1110 Ron Hightower, 09201
10 0780 Sabina Jordan, 05838
1 1110 William Hart, 09211
5 0829 Chris Forsythe, 12323
1 0759 Mark Snell, 05845
1 1110 David Womble, 09222
1 0318 George Davidson, 09201
1 0780 Stephen Ortiz, 05838

1 9018 Central Technical Files, 8940-2
2 0899 Technical Library, 9616
1 0612 Review & Approval Desk, 9612
For DOE/OSTI