

Representing Knowledge and Experience in RPDAgent

John A. Sokolowski, Ph.D.
Virginia Modeling, Analysis & Simulation Center
College of Engineering and Technology
Old Dominion University
Norfolk, VA 23529
757-686-6215
jsokolow@odu.edu

Keywords:

Knowledge Representation, Recognition-Primed Decision Making, Multiagent System Simulation

ABSTRACT: *Military simulations lack algorithms that adequately describe the cognitive decision process employed by military commanders. This is especially true at the operational level of warfare. Research was undertaken to improve upon these algorithms. What resulted was a model named RPDAgent that was able to mimic the human decision process and produce decisions that were equivalent to those made by humans for a given operational decision scenario. Key to the functioning of this model was its method for representing knowledge and experience. RPDAgent captured important concepts defined by Recognition-Primed Decision making. By modeling these concepts, it produced a unique and effective methodology for representing the experience needed to define complex decision-making.*

1. Introduction

The U.S. military is relying more and more on constructive simulations for force training, for war plans analysis, and for experimentation with new warfighting concepts. To support these domains, the military requires models of human decision-making that incorporate realistic human behavior [1]. However, most of these decision models are focused on the tactical level of warfare. Very few decision models exist that mimic the cognitive decision processes of senior military commanders at the operational level of warfare. In the training and experimentation domains, role players are required to make most of the operational decisions and to input them into the models. In the analysis domain, one runs the risk of erroneous results because of the stereotypical and homogeneous decision algorithms that currently populate the models.

To address these shortcomings, an effort was undertaken to develop a decision model that more accurately mimicked the decision process of senior military commanders. What resulted is a model called RPDAgent. RPDAgent is based on the *Recognition-Primed Decision (RPD)* model of cognitive decision-making by Klein [2]. RPD has been shown to accurately describe the cognitive decision process used by military commanders in the field [3,4]. RPDAgent was implemented using multiagent

system simulation techniques including the concept of a *composite agent* set forth by Hiles, et al., [5]. RPDAgent's performance was validated by comparing its decisions against the decisions of a group of military officers playing the role of a Joint Task Force Commander. Analysis showed that RPDAgent produced decisions equivalent to those of the role players for a typical operational military decision scenario. See the companion paper [6] for a discussion of the model validation results.

Under the RPD concept, decisions are influenced, to a significant degree, by a person's past experience. RPD is meant to describe the decision process used by a person who possesses significant experience in a decision area. In other words, RPD describes how experts (e.g. senior military commanders) make decisions. Capturing this expertise in a computational form was critical to producing a model that could mimic this decision process. The remainder of this article describes basic RPD principles and the techniques used in RPDAgent to represent human experience and its application to selecting a decision that will satisfy a given situation. Of note, RPDAgent does not produce optimal decisions. It was meant to simulate the variable and imperfect decisions that military commanders often are forced to make because of incomplete or erroneous information, time pressure, and their own personality characteristics and past experiences.

2. Pertinent RPD Concepts

Understanding how RPD_{Agent} represents experience will require some knowledge of RPD concepts. A decision-maker's ability to recognize the context of a particular decision situation and his or her ability to identify an appropriate action to take is based on past experience. The broader a person's experience base in a particular domain, the more likely he or she is to match that experience to a current situation. When situational recognition occurs, four byproducts result. They are: cues, goals, actions, and expectancies. These four elements come from experience and describe the cognitive concepts on which a decision maker operates. Cues represent those physical and mental elements on which a decision maker keys to understand and to monitor a situation. Knowing what small set of cues to monitor, out of all the possible pieces of information reaching the decision maker, is the mark of an expert. Cues are often made up of aggregated pieces of information that a decision maker assembles in his or her mind.

Goals are a key part of the recognition process. They represent an end state that the decision maker is trying to reach. He or she may be trying to satisfy several goals and must pick the action that can best accomplish them.

Actions represent the set of potential decisions from which a decision maker can select. A decision maker, experienced in the decision domain, intuitively knows which action is likely to be most favorable. He or she will use mental simulation to evaluate this action to see if it is appropriate for the specific context of the current situation. If so, it will become the decision. If not, the next most favorable action is evaluated until one is found that will produce satisfactory results. Note that the decision is not necessarily optimal.

Expectancies act as a control mechanism in the decision process. They represent criteria against which the decision maker can gauge the progress of the situation and determine if any adjustments are required because of a changing context or ineffective results. See Klein [2] for a complete description of RPD.

3. A Structure for Modeling Experience

To understand how RPD_{Agent} represents experience from the ideas discussed above, one must understand the concept of a frame. Minsky [7] was the first person to identify a frame as a data structure to hold information about a person's environment. In his words:

“When one encounters a new situation (or makes a substantial change in one's view of the present problem), one selects from memory a structure

called a frame. This is a remembered framework to be adapted to fit reality by changing details as necessary. A frame is a data structure for representing a stereotyped situation...Attached to each frame are several kinds of information. Some of this information is about how to use the frame. Some is about what one can expect to happen next. Some is about what to do if these expectations are not confirmed.”

Because of Minsky's work, a frame was chosen as the data structure best suited to hold the major portion of RPD_{Agent}'s experience. Each frame within RPD_{Agent} corresponds to a single experience that holds all the cues, goals, and actions describing that experience. It is defined by the following structure:

$$F = (C^*, G^*, A^*)$$

where F is a frame, C^* is a structure containing all cues for an experience, G^* is a structure containing all goals for an experience, and A^* is a structure containing all actions for an experience. *The number and type of goals, actions, and cues associated with a particular frame is, itself, an embodiment of experience.*

When presented with a decision request, RPD_{Agent} searches its table of frames, looking for a match. If no match is found, the model does not possess the experience to render a decision for the situation. If a match is found, the matching frame, with its associated cues, goals, and actions is retrieved. This action represents the recognition process within RPD.

Once a frame is identified, RPD_{Agent} must develop an internal representation of the decision situation, much like how humans develop an internal interpretation of their external environment based on their past experience. To accomplish this internalization, RPD_{Agent} aggregates the environmental variables that describe its external environment into cues that represent its internal interpretation of that environment. Cues are higher-level abstractions of lower-level environmental variables that describe elementary physical or mental parameters. Military commanders, at the operational level of warfare, do not concern themselves with these elementary parameters. Instead, they, or their staffs, aggregate them into this higher-level abstraction that represents one or more of these parameters. For example, in deciding where to conduct an amphibious landing, a military commander may consider landing zone hydrography as a cue. The commander would want to know if the hydrography of each potential landing zone (each landing zone corresponds to a potential action or decision) satisfactorily supports the amphibious landing. Making up the evaluation of hydrography may be many

Table 1. Hydrography cue structure

Cue	Environmental variables	Description	Value ($e_{i,j}$)
Hydrography	Reef	None	2
		Partial	1*
		Full	0
	Water depth	Shallow	2
		Moderate	1
		Deep	0*
	Anchorage	None	0*
		Yes	2
	Tides	Small	2
		Moderate	1*
		Large	0
	Currents	Light	2
		Moderate	1
		Severe	0*

environmental variables such as water depth, tides, and currents. The environmental variables associated with the hydrography cue must somehow be combined to produce a representation of hydrography for a particular landing location. In RPDAGENT, environmental variables are assigned numeric values that represent a variable's characteristics. The higher the value, the more favorably the variable influences the cue. The cue value is calculated by summing the values of its associated environmental variables. The following equation represents this calculation:

$$cv_j = \sum_{i=1}^n e_{i,j} \quad (1)$$

where $e_{i,j}$ is the i th environmental variable value associated with the j th cue, cv_j is the cue value associated with the j th cue, and n is the number of environmental variables associated with the j th cue. This calculation is repeated for all j cues.

Once the cue values for all cues are calculated, they are further refined into fuzzy values. Humans tend to interpret their environment in fuzzy terms rather than discrete values. A military commander may categorize landing zone hydrography as unsatisfactory, marginal, or satisfactory. To achieve the same type of cue interpretation, RPDAGENT generates fuzzy values for each cue based on the cue value. Each cue has three fuzzy sets associated with it, an unsatisfactory fuzzy set, a marginal set, and a satisfactory fuzzy set. The higher the cue value, the more likely it is to fall in the satisfactory range. Triangular-shaped fuzzy sets were used since they best represent an optimum value for the set and a trailing off of that value as one moves to either side of it. See Zadeh for a discussion on fuzzy variables and fuzzy sets [8]. *The*

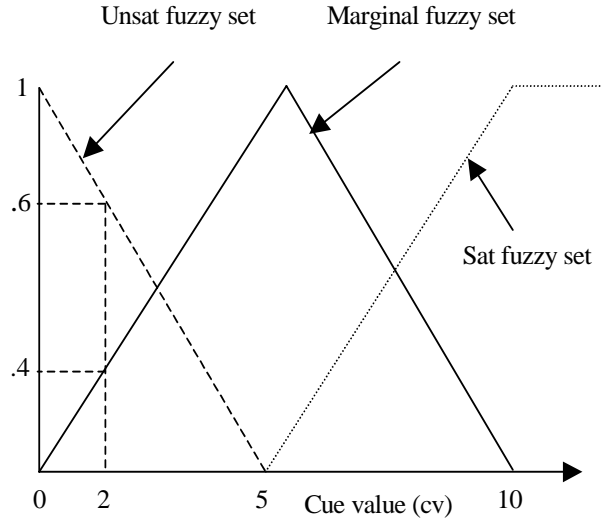


Figure 1. Hydrography fuzzy sets

shape and range of these fuzzy sets is a key element that defines experience within RPDAGENT. By manipulating their shape and range, one can adjust RPDAGENT's internal interpretation of a cue based on the type of experience the model user wants RPDAGENT to possess. One can generate different fuzzy sets if one desires different fuzzy interpretations of the cue based on past experience.

The following example serves to illustrate cue computation. It is based on the hydrography cue of the amphibious landing location decision mentioned earlier. Table 1 depicts one possible structure for the hydrography cue. The cue has five environmental variables associated with it. Each variable has two or three descriptive values (experience provides this characterization) and corresponding numeric values. The values with the asterisks identify the variable values for the particular landing location under consideration. Using Equation (1), hydrography would have a cue value of 2 out of a possible 10 for this location (10 if all environmental variables were at their most favorable values). This value is then *fuzzified* as illustrated in Figure 1. For a cue value of 2, the unsat fuzzy set height is 0.6 and the marginal set height is 0.4. The sum of these heights provides a normalized value on which to base the percentage probability of membership to the sets. The subjective probability of being unsat is therefore 0.6/1.0 and the subjective probability of being marginal is 0.4/1.0. A random number is generated to make the selection.

Each action within a frame represents either a past or a current decision option. Actions can be thought of as choices that are characterized by their environmental variables. Thus, RPDAGENT can determine, in part, how satisfactory an action is, by summing the cue values for

each action. This computation is given by the following equation:

$$AV_i = \sum_{j=1}^n cv_{j,i}$$

where $\sum cv_{j,i}$ is the sum of all cue values associated with the i th action, n is the total number of cues associated with action i , and AV_i is the action value for the i th action. The action with the largest action value is considered the most favorable. RPDAGent uses this method to mimic the decision maker's intuitive identification of the most favorable action to further evaluate. Cue values are used for this computation rather than cue fuzzy values because this calculation is meant only as an intuitive indicator of the most favorable action. RPDAGent must carry out further evaluation before this action is chosen as the most suitable for the situation.

Once the most favorable action is identified, RPDAGent must use its experience to determine how well the selected action meets the goals of the situation. In RPD, this is analogous to a decision maker using mental simulation to determine if the selected action satisfies the current situation. Associated with each goal is a set of cues that are indicators of how well the goal is being satisfied. *The association of particular cues to particular goals is another element of experience over which the model user has control.* Goals also have fuzzy sets associated with them. These fuzzy sets define how well a goal is being satisfied. Again, unsatisfactory, marginal, and satisfactory sets were used. To generate a goal assessment, a goal value is calculated from the cues associated with the goal. This calculation is given by:

$$GV = \sum_{i=1}^n (GV_c^i * cw_i)$$

where GV_c^i is a numeric representation of the i th cue fuzzy value associated with goal j , cw_i is the cue weight associated with the i th cue, n is the number of cues associated with goal j , and GV is the goal value for the j th goal. Cue weighting is included because decision makers often consider some cues more important than others. This calculation is repeated for all j goals.

Goal fuzzy values are then generated by applying a *goalfuzzyvalue* function to the goal value as previously described for cue fuzzy values. *The definition of the goalfuzzyvalue function is another key part of experience definition within RPDAGent.* Goal fuzzy values represent RPDAGent's assessment of how well each goal is being met by the action under consideration.

To illustrate this computation, the landing location decision will be examined once again. Suppose one goal

for the landing location decision is to accomplish the mission. This goal may have five cues associated with it.

Table 2. Goal evaluation example

Goal	Cues	Fuzzy value	Numeric value
Accomplish mission	Beach topography	Marginal	1
	Beach hydrography	Sat	2
	Beach obstructions	Sat	2
	Beach staging area	Marginal	1
	Route to objective	Sat	2
Goal value (GV)			8

These cues, their fuzzy values, and the numeric representation of the fuzzy values are given in Table 2. The goal value has been computed to be 8 out of a possible value of 10 (10 if all fuzzy values were sat). The goal fuzzy value is then calculated using fuzzy sets similar to those in Figure 1.

If all goals were evaluated as satisfactory, the action under evaluation becomes the model's decision. If RPDAGent evaluated one or more goals as other than satisfactory, it must see if it can negotiate a compromise. Compromise is carried out by the final concept that defines experience, a *negotiation function*.

Very rarely can a decision maker find an action that will completely satisfy all his or her goals. Experience tells a decision maker how far he or she can compromise on a specific goal and still arrive at a satisfactory decision. RPDAGent performs in the same way. It instantiates a *reactive agent* for each goal in the frame. See Wooldridge [9] for a discussion of reactive agent characteristics. Each reactive agent is responsible for evaluating the attainment of its assigned goal. When one or more goals are not evaluated as satisfactory, the reactive agents try to negotiate a compromise by lowering the standards by which they evaluate their goals. Since reactive agents are autonomous, they are not influenced by other agents' evaluations. There is a threshold set, below which a reactive agent will not go to achieve a compromise. *The negotiation methodology and the compromise threshold are parameters that further define experience within RPDAGent since they represent how a person resolves goal conflicts.* Reactive agents implement negotiation by mapping their previously calculated goal value to a revised goal value through a multiplication factor. The following equation represents this calculation:

$$GV^n = GV * RF$$

where GV is the previously calculated goal value for a specified goal, RF is a risk factor, and GV^n is the compromise goal value. The risk factor is a real value that represents RPD Agent's tolerance for risk. Each reactive agent performs this calculation for its respective goal.

Once the new goal value is computed, a new goal fuzzy value is determined. If all goals were evaluated as being satisfactorily met, then a compromise has been reached. The action under evaluation is chosen as the decision. If no compromise is possible, the next most favorable action is selected and evaluated in the same manner as described above. If RPD Agent can not find a satisfactory action, it renders a default decision that is relevant to the situation.

4. Summary

RPD Agent was designed as a model to improve upon the decision-making capability of military simulation systems especially at the operational level of warfare. Key to having an effective model was the ability to represent human knowledge and experience in a computational form. RPD Agent accomplished this by defining several concepts that captured a person's cognitive experience as defined by the RPD model. These concepts included:

1. The number and types of actions, cues, and goals that were used to define an experience.
2. The use of cues and fuzzy sets to capture how a person internalizes his or her external environment.
3. The association of specific cues to goal evaluation.
4. The use of fuzzy sets to evaluate goal accomplishment.
5. The negotiation methodology and the compromise threshold parameters that define RPD Agent's ability to arrive at a decision.

Combining these concepts helped produce a computational model that mimicked the human decision process characterized by RPD and practiced by operational military commanders.

The author is grateful for funding support for this research provided by Joint Warfighting Center, U.S. Joint Forces Command and the Defense Modeling and Simulation Office.

5. References

[1] Pew, R.W. and Mavor, A.S. *Modeling Human and Organizational Behavior: Application to Military Simulations*. National Academy Press, Washington, D. C., 1998.

[2] Klein, G. *Sources of Power: How People Make Decisions*. The MIT Press, Cambridge, MA, 1998.

[3] Pascual, R. and Henderson, S. "Evidence of Naturalistic Decision Making in Military Command and Control," In *Naturalistic Decision Making*, Zsombok, C.E. and Klein, G. Eds., pp. 217-226, Lawrence Erlbaum Associates, Publishers, Mahwah, NJ, 1997.

[4] Klein, G. "Strategies of Decision Making," *Military Review*, Vol. 69, No. 5, pp. 56-64, 1989.

[5] Hiles, J. et al. "Innovations in Computer Generated Autonomy," Tech Report NPS-MV-02-002, Naval Postgraduate School, Monterey, CA, 2001.

[6] Sokolowski, J.A. "Enhanced Military Decision Modeling Using a MultiAgent System Approach," In *Proceedings of the 12th Conference on Behavior Representation in Modeling and Simulation (BRIMS)*, Scottsdale, AZ, May 12-15 2003.

[7] Minsky, M.L. "A Framework for Representing Knowledge," In *Computation and Intelligence*, Luger, G.F. Ed., pp. 163-189, AAAI Press, Menlo Park, 1995.

[8] Zadeh, L.A. *Fuzzy sets and their applications to cognitive and decision processes*. Academic Press, New York, 1975.

[9] Wooldridge, M. *An Introduction to MultiAgent Systems*. John Wiley & Sons, Ltd., West Sussex, England, 2002.

Author Biography

JOHN A. SOKOLOWSKI is a research professor and Director of Research at the Virginia Modeling, Analysis & Simulation Center. Prior to this, he spent twenty-seven years in the Navy as a submarine officer. His last assignment was Head, Modeling and Simulation Division, Joint Warfighting Center, U. S. Joint Forces Command. He holds a Bachelor of Science degree in Computer Science from Purdue University, a Masters of Engineering Management from Old Dominion University (ODU), and a Ph.D. in Engineering with a Concentration in Modeling and Simulation from ODU. His research interests include human behavior and decision modeling and multiagent system simulation.