Proceedings of the 2001 Winter Simulation Conference
B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, eds.

COMMANDER BEHAVIOR AND COURSE OF ACTION SELECTION IN JWARS

Deborah Vakas
John Prince
H. Ric Blacksten

CACI Inc. – Federal
1600 Wilson Blvd. Suite 1300
Arlington, VA 22209, U.S.A.

Chuck Burdick

Lockheed Martin Corp.
1555 Wilson Boulevard, Suite 620
Arlington, VA 22209, U.S.A.

ABSTRACT

The Joint Warfare System (JWARS) is being equipped with a Commander Model (CM) to perform situation assessment and Course of Action (COA) selection, and a Commander Behavior Model (CBM) to bias decisions with a commander's leadership style. The CM is a hybrid artificial intelligence system that models doctrine through the use of fuzzy rule sets, together with a tree-based lookahead algorithm for the strategy. The CBM employs behavior-based fuzzy rule sets to augment the CM in assessing the situation, and in biasing the COA selection criteria. Extending from Myers-Briggs personality traits, the CBM links personality traits to military attitudes, consequences and values. Employing the fuzzy rule sets, the resulting sets of values are combined to select a specific COA with an auditable trail. Users will have the ability to modify both the input parameters and the underlying rules. The CM/CBM is applicable to decisions at multiple echelons.

1 INTRODUCTION

The JWARS Land Commander Model (CM) and associated Commander Behavior Model (CBM) constitute a JWARS Course of Action (COA) Selector System designed to solve many of the problems associated with modeling doctrine. The system uses a modular design and fuzzy rule engine so that that doctrine may be expressed in JWARS nearly "as is." It uses data from well known personality tests in fuzzy rule sets to influence the interpretation of this doctrine, and a chess-like look-ahead engine to see the results of various applications of this doctrine. It then chooses the COA that gets to a goal while best satisfying other values such as "minimize attrition of friendly troops." There will be a CM available for every major decision-making agent in JWARS with an option to give each a unique personality. For the purposes of this paper, the term *agent* refers to headquarters units and, where applicable, specific commanders. The agents each know their own state and reason about the state of the environment and of opposing agents. They also each have a model of their rules (doctrine) and the rules (doctrine) of other agents along with an estimate of what the other agents likely perceive. This combination of capabilities allows the agents to reason concerning the likely course of future events. The better his mental model of the opposition is, the better a commander's confidence in the outcome can be. This confidence score puts quantifiable value on knowledge of the enemy and on denial of that same information to the enemy. This provides an opening to model information operations.

A simple stand-alone prototype of the CM/CBM COA system has been developed and applied to a simple scenario to demonstrate how psychological factors can influence the perception of the situation and the course of action chosen (JWARS Office 2001).

1.1 Statement of Problem

One of JWARS potential applications is as a test bed for doctrine evaluation. Doctrine evaluation can be problematic, however, since for any given situation there are several responses that may be doctrinally correct, but have very different outcomes. This variability presents a problem when modeling responses to doctrine. One way to deal with this problem is to obtain a statistical distribution of the responses to situations in the real world, and then roll the dice to draw from it. However, this would detach the decision-making process from specific current situations and the behavior of a particular commander would not be consistent over time. For example, a commander with a conservative response to one situation is likely to respond conservatively to a wide range of other situations as well, and the dice could not capture this. Nor is it possible to model every possible combination of situation, doctrine, and commander statistically, because there are just too many of them.

Another way to deal with this problem is the traditional expert system way of expressing doctrine. This approach converts doctrine to a form that is a blueprint for every possible case. However, because the result is a stylized form of the doctrine instead of the real doctrine, the method seldom takes into consideration all the time-dependent aspects of the upcoming battle, the impact of different personalities implementing the doctrine, or the perceived likely enemy responses. Furthermore, if we were to cover all cases, the resulting knowledge base would not be easily modified or modularized, because it is the doctrine itself that is being modified, not the agent that is implementing it.

A better way is to leave the doctrine the way it is, and to model the human being who interprets and applies it. Doctrine is not complete in itself. It is designed to be implemented by a cognitive being with perceptive states and a rational personality. Doctrine cannot cover every case, nor should it. Rather, it serves as a set of heuristics to achieve goals, which must be thought out strategically in each individual case. Modern social science has shown that personality has a considerable impact on both individual perceptions and on narrowing down (valuing) what is considered. Thus, the impact of personality may be thought of as a part of the overall heuristics for strategy, just as doctrine is.

Care must be taken in the modeling of strategy as a human endeavor. For example, in military history, there are numerous instances where a commander's personality shaped his perception of the situation to the great advantage or detriment of his force. Likewise, in dealing with risk, people do not always just maximize expected gain; they often limit their risks. The existence of the insurance industry is evidence of this. Just as JWARS takes into account the bounded rationality of perception, so it will also offer the opportunity to take into account human bias to accurately model the decision-maker.

In modeling the human decision-maker, it is also important to consider that human beings not only have a bounded rationality, but that they realize that their opponents do as well. This "knowing of knowledge" is called "second order cybernetics," and may be taken to even higher orders, where an agent knows what an opponent knows about what an agent knows, etc. Of course, the decision-maker's map of the opponent's bounds of knowledge is itself bounded by his knowledge of the opponent's sensor capabilities, the opponent's intelligence on friendly doctrine, etc. Thus, to more accurately model the impact of doctrine we strive to more accurately model the humans that use it.

It is important to understand that doctrine is not a blueprint that can be followed to the letter in every case, but is rather a guide to be used heuristically by strategic agents, those with a larger and longer term view of the battlespace. These strategic agents possess bounded rationality and some awareness of other agents' perceptions. Further, it is important to model these human effects in a modular way, so that psychological and social contexts and even theories of those contexts may be mixed and matched in a consistent manner with the doctrine to be implemented.

1.2 Overview of JWARS Solution

The JWARS CM is designed to implement doctrine accurately by keeping the doctrine nearly "as is" in modular rule sets within agents. Other rule sets on psychological, cultural, political, and economic constraints give different contexts in terms of different military attitudes, the importance of various values (e.g., own force casualties, collateral damage, risk avoidance), and the perceived consequences of specific actions. Doctrine may then be mixed and matched with different contexts, but it is always used in conjunction with a strategic look-ahead at the results of actions in particular scenarios. This strategic look-ahead includes probabilistic agent perception and a model of the opponent's perception of the agent.

The CM is able to implement doctrine well by virtue of its hybrid AI system design that puts together concepts from both the connectionist and the classic AI game playing paradigms in such a way that they help one another. Doctrine and personality are expressed in fuzzy rule sets, and are used as heuristic guides to strategic thinking about achieving an objective. For thinking strategically, agents use a look-ahead engine very similar to the type used to play chess.

The personality and doctrinal rules are implemented with fuzzy rule sets. CBM fuzzy rules describe behavior and personality, while CM fuzzy rules describe doctrinal reactions. Fuzzy rule sets themselves are a mixture of connectionism and classic AI and offer several advantages over traditional expert systems. Besides the fact that they can elicit behavior with fewer rules than traditional expert systems, their rules can have syntax very close to the language of doctrine, allowing doctrine to be kept nearly "as is." The rules in fuzzy systems are very additive, and it does not matter if they sometimes contradict each other. The rules ensure that there is always an interpretation of a situation, even if it is partial and even if there are multiple interpretations. Scalar ratings of situations based on all the information "fall out" from the membership functions. This characteristic makes them good for modular substitution and simple combination. It also opens the possibility for learning and adaptation, because fuzzy rule systems can be formulated to learn the same way that neural networks do, though this has not been done in the CM/CBM work to date.

The JWARS Land Commander Model is responsible for three primary functions. First, it has the responsibility for assessing the situation. This responsibility includes interpreting or setting the conditions of the assigned mission, estimating the enemy and his own troop strength (including external sources of firepower and supply status), the impact of the terrain on his and the enemy's force, and the feasibility of the assigned schedule given the foregoing conditions. Because JWARS operates on perception rather than ground truth for assessment, this assessment is a difficult task and the JWARS development team is still in the process of defining how personality biases perception, particularly of enemy strength. Second, the CM selects the course of action (COA) from those that are available. And third, the CM monitors the plan and makes the decision when to modify or abandon the plan and select a new one. Only the COA selector portion of the CM is discussed further in this paper, but the technique of applying individualistic behavior (the Commander's Behavior Model) to the decision-making process applies to all elements of the CM.

The COA selector is implemented with a chess-like look-ahead that takes perception into account by assuming that there may be different starting configurations for the game board. To some extent, a better analogy than chess may be a card game where the options to play, pass, bluff, or fold are available. The look-ahead capability allows simultaneous moves, players with incomplete knowledge of each other, and different rules for each player. Fuzzy rule sets influence this system in many ways. Fuzzy rule sets interpret a given situation and find the next available moves. They also serve as heuristics to rate the moves, so that not all moves are considered, and the size of the tree is constrained. Fuzzy rule sets rate the game boards corresponding to the leaves of the tree, and bias these ratings according to the agent's attitude towards risk, if the CBM is turned on. The fuzzy sets alone are not sufficient because they are not complete enough to be applied in every strategic context. The look-ahead alone is not enough because it needs a way to interpret ambiguous situations. It needs something to cut down the possible next moves, and it needs a way to rate situations in trees that do not extend to the final objective. Both AI methods help each other to provide an intelligence that uses doctrine in much the same manner that people use doctrine; i.e., interpreting the situation and applying the doctrine to the specific situation.

2 PHASE 1 SOLUTION

2.1 Prototype Description

The JWARS development team successfully implemented and demonstrated prototypes of the CM and CBM as stand-alone models that have JWARS-like data and software interfaces. The current goal is to refine the prototype models and integrate them into JWARS. The prototype is a valuable tool to "shake out" rules, algorithms, and personality traits and attitudes.

2.1.1 Overview of the Prototype

The prototype CM is a stand-alone Commander with a notional rule set, a simple look-ahead, and a simple scenario. When developed further, it will be integrated into JWARS and will also exist as a stand-alone that analysts may use to create and test rule sets. The associated prototype CBM has rules that convert Myers-Briggs Type Inventory scores to evaluations of consequences and ways of dealing with risk. The look-ahead treats perception as ground truth. It does not employ game theory, but simply scores the look-ahead simulation outcomes and picks them according to the notional CBM rules. The prototype scenario has a simple notional doctrine employed by both opposing forces and a single objective sought by each. Each aggregated force has a supply and consumption rate, an attrition rate tied to its strength, and a movement capability based on vehicle assets and terrain. The CM prototype selects a COA according to the outcome of its look-ahead simulation and its doctrinal rules. The CBM prototype biases the CM course of action selection by using differences in personality to modify the scoring weights. It should be noted that personality only counts for a portion of the selection score. A Commander with an overwhelming force who is ordered to attack would seldom decide not to attack. However, he might delay the attack because he was concerned with friendly supply status or with his lack of knowledge of the enemy.

2.1.2 Description of COA Selector

The CM uses the course of action (COA) selector to determine a series of actions and conditions (a plan) that will accomplish the mission and reach the objective. The COA selector starts with a simplified view of the actual battlefield situation including attributes of the commander's own situation and his perceived view of his opponent's situation. The attributes considered include: strength; initial supply level, supply consumption, resupply rates; mobility; orientation; position; and unit length and depth. These details are combined into a mental model of the situation, including a coarse simulation of the JWARS simulation, hereafter referred to as the game board. The commander starts with some information on himself and his perception of the enemy commander. This set of data includes scores for various personality characteristics, i.e.: introverted or extroverted; intuitive or sensing; thinking or feeling; and judging or perceiving. In addition, the commander has his own fuzzy rule set and a perception of his opponent's fuzzy rule set. These rule sets map personalities to evaluations of situations and preferred methods of choosing between alternative feasible courses of action. For the prototype, the data for the commander's rule set and the opponent's rule set are identical since no cultural or doctrinal differences are considered. However, simply changing the personality types engenders considerable change in the COA selected.

The COA selector takes the initial situation and selects what primary maneuvers can be done by each side, e.g. attack, defend, or withdraw. For each feasible combination of those moves, a new game board is created. The commander then plays out each of the battles via the game board for a set period of time or until one of the units on the game board has reached its objective. The commander then repeats the process of selecting maneuvers and playing out the resulting battles on the game board until a time limit has been reached. At that point, all of the existing game boards are stopped and are given a rating. In military terminology, this is analogous to the Commander's Wargame. The COA rating is determined by the fuzzy rule set based on such factors as strength loss for both the commander and the opponent, the remaining distance to the objective, and the amount of supplies available. The actual maneuver scoring highest depends on which type of evaluator the commander uses. The options are a maximum gain evaluator that picks the maneuver that leads to the best possible final result, a minimum loss evaluator that chooses the maneuver that avoids the worst result, a maximum robustness evaluator that chooses the maneuver that leads to a good result in the largest number of cases, and a maximum expected value evaluator that weights the final results by their relative likelihood and calculates the maneuver that yields the best result. The actual maneuver selected is determined by a weighted combination of all of the evaluator types. The weightings given to each evaluator are selected by the fuzzy rule set and can be biased by the personality of the commander. Maneuvers are also selected for the opponent based on an estimate of his capabilities and preferences. When each side has selected its maneuvers and those of its opponent, the appropriate child game board that represents that combination of maneuvers is promoted to parent and the process is repeated for a set length of time to determine each side's entire plan.

2.1.3 Description of Myers-Briggs

There are several psychological tests that can be used to summarize personalities numerically. The two that are most widely used are the MMPI (Minnesota Multiphasic Personality Inventory) and the MBTI (Myers-Briggs Type Indicator). Both may be used to find personality types. However, the Myers-Briggs, based on Jungian personality type theory, is solely used for personality typing while the MMPI is mainly used for diagnosis of pathological conditions. In any case, any personality descriptors could be substituted in our system, because we are using modular rule sets.

The Myers-Briggs Type Indicator is a self–report instrument that classifies people according to how they habitually perceive and act. They are described by values along four bipolar axes. The axes are *energy directedness*, *information processing*, *decision-making*, and *organization*. Along the energy directedness axis, if one is more directed to the outer world of activity and spoken words then he is *extroverted*, but if he directs his energy towards the inner world of thoughts and emotions then he is *introverted*. Along the information processing axes, if one processes information more in the form of known facts and familiar terms then she is *sensing*, while if she processes information in the form of possibilities or new potential then she is *intuitive*. Along the decision-making axis, if one prefers to make decisions on the basis of logic and objective considerations, then he is *thinking*, while if he prefers to make decisions on the basis of personal values then he is *feeling*.

Finally, along the organization axis, if one prefers to organize her life in a structured way, making decisions and knowing where she stands, then she is *judging*, while if she prefers to organize her life in a flexible way, discovering life as she goes along, then she is *perceptive*. These polar differences do not denote good or bad, just different. The four bipolar axes make sixteen personality types possible, one for each axis extreme. When fuzzy logic is applied, intermediate combinations are allowed permitting finer personality discriminations, e.g., 0.5 P and 0.5 J rather than 0.0 P and 1.0 J.

Individual scores along each of these four axes for both commanders are input to the prototype CBM. The emphasis in JWARS is to build an evaluation tool and make it available to the military analytical community. It will probably take several years to develop data that is widely credible.

2.1.4 Description of the Fuzzy Rule Sets Used

Personality is implemented through fuzzy rule sets, but there are also fuzzy rule sets in the COA selector even without the CBM. The CBM only puts a bias on which COA is selected, a bias that the analyst may turn off at will. When active, the CBM affects the prototype CM COA selector in two ways: in evaluating game boards and in picking a COA tree evaluation strategy.

The CM, without the CBM active, is designed to have three fuzzy rule sets: one that interprets situations, one that determines doctrinal reactions to situations, and one that tells how likely the objective is to be obtained in a given situation. The prototype only implements the last of these. The first rule set is not implemented in the prototype because the agents are given their knowledge of the situation (assuming it would be generated in JWARS). Neither is the rule set that determines specific doctrinal reactions to simulations implemented since a notional doctrine is being used in the prototype with only three reactions (attack, halt/defend, and withdraw), each of which always receives consideration in the COA selector.

Later (in the long-term JWARS implementation), the rule set that interprets situations will be aided by a rich role-based syntax that has convenient variables that can richly describe a situation. Here are some sample rules using this syntax to interpret situations:

- If attacker A is facing defender B's front and advancing, then A is frontal assaulting *B*.

- If attacker A is facing defender B's flank and advancing, then A is flank assaulting B.

Suppose that a game board with an ambiguous situation read into this fuzzy rule set that an attacker was coming at a defender diagonally. In that case, there would be partial memberships in frontal assaulting and in flank assaulting.

The list of memberships in situations is then sent through a rule set which determines doctrinally correct actions to perform in those situations. These rules are fuzzy, and so determine a preference score for each response to a situation. However, the membership value for a response depends not only on the amount of preference, but also the membership value for the situation that is responded to. Thus, ratings for several responses fall out. For the look-ahead, only the top *n* responses need be considered (*n* being analyst input).

Here are two sample rules in a the response rule set:

- If force A is flank assaulting force B, then B must face A.
- If force A's strength is much greater than B's strength and B's loss rate is very attriting, then B should withdraw.

Note that the stronger hedge "must" is used in the first example, while the weaker hedge "should" is used in the second example. This difference would give "facing" a higher rating in the example where an attacker comes at a defender diagonally, even at a 45-degree angle. However, the smaller the angle, the lower the rating for facing and the greater the advantage for staying in prepared defensive positions is scored.

The last COA selector rule set, the one that assesses the likelihood of achieving the objective for a game board, is used to rate the game boards corresponding to the leaves of the COA selector tree. An example of a rule from this rule set is:

- If attacker A is much stronger than defender B and A is close to objective C, then the objective is likely.

The CBM adds four additional fuzzy rule sets: one that converts personality scores to attitudes, two sets that work together to rate the game boards, and one that biases COA evaluation under uncertainty. All of these exist in the prototype. The fourth rule set will be modified later to incorporate confidence levels for intelligence data, such confidence levels being a feature of JWARS C4ISR-centric design. This input should lessen the impact of a commander's personality when confidence is high and increase the impact it when the confidence level is low.

The rule set that converts personality scores to attitudes uses personality profiles as input and has rules that are specific to a theory of personality. We use the Myers-Briggs for this rule set, but the MMPI or any rule set corresponding to a theory of personality may be used to generate this rule set because it interfaces to the other rule sets through an intermediate concept of attitudes. An example of a rule from this set is:

- A person whose decision-making axis is feeling has a benevolent attitude towards others

Two rule sets evaluate game boards. The first determines the consequences of a wargame's outcome and the second assesses Commander satisfaction with that outcome given his attitudes. Here is an example rule of the first rule set (linking game board situation to consequences):

- If friendly frontal attack is true, friendly attrition is very true.

Here is a sample rule from the second rule set that couples attitudes and consequences to value satisfaction:

- If attitude towards others is benevolent and friendly attrition is true, then value satisfaction is low.

Value satisfaction, the likelihood of reaching the objective, and commander personality each influence the overall score for a game board. For example:

- If evaluation valuation style is to have a broader view of winning and value satisfaction is low, then overall satisfaction is low.
- If evaluation valuation style is to value the objective, and the objective is likely, then overall satisfaction is high.

The next rule set of the CBM biases COA evaluation under uncertainty. An example for a rule for this rule set is:

- A commander whose attitude towards others is benevolent has a minimizing loss COA evaluation strategy.

This last rule set exists in both the prototype and the long-term CBM/COA selector design, but the score is used differently because the prototype has no uncertainty in it while the long-term version does.

2.2 Overall Design of CBM

2.2.1 COA Prototype to JWARS Long Term CBM/COA Implementation

The CM prototype contains a simplified version of the Commander's Behavior Model and COA selector that will be implemented in JWARS. Some of the key differences between the current prototype and the long-term JWARS version are described below in Table 1.

Table 1: Moving from Prototype to Long Term JWARS Implementation

Current Prototype	Long Term JWARS CBM/COA Selector
Agents treat perceptions as ground truth, i.e. have full confidence in them.	Agents recognize they only perceive situations (do not know truth), and take different fuzzy likelihoods into account in their decisions.
Agents all have the same knowledge of themselves that the opponent has.	Each agent has a model of what their opponent likely knows about that agent, based on opponent's position, sensor capability and intelligence, making deception and information operations possible.
Agents all have the same doctrinal rule sets.	Agents have individual rule sets based on their country's doctrine, and a model of opponent's rule sets based on their identification of the opponent and its doctrine.
Agents have a simple scoring function.	Agents act based on the assumption that opponents will do what is best for them given their doctrine and personality.
Agents have the same moves available to them at all times.	Agent's doctrinally correct responses are constrained by their forces and supplies, by the enemy situation, and by the environment, with only the top n options considered in COA tree creation.

2.2.2 How Attitude Towards Risk Affects Perception of Likelihood

Scoring of the COA tree according to risk is different in the JWARS CM/COA selector and the prototype because there is a recognition of the limits of perception in the long-term but not in the prototype. In the JWARS version, each agent has a perception, and the agent knows that it is operating based on perceived truth, so it assigns likelihoods to an interpretation of a situation based on its estimate of the truth of that interpretation. This likelihood is a function of the side's intelligence system and can be biased by the commander's personality. These likelihoods can be normalized to proper (subjective) probabilities by summing the initial estimates of interpretation likelihoods across all interpretations, and then dividing those estimates by that sum. For example, if an agent is informed of approaching enemy forces, he might also be informed that it is highly likely that an entire division is approaching and much less likely that the force is larger (a corps) or smaller (only a brigade). Each likely situation is taken into account in the COA tree evaluator, and the likelihood that a possible situation is true is calculated into its ultimate score.

When the CBM is active the likelihood estimator is biased so that the COA is not chosen by the standard "maximize expected value" alone. This is because people often exaggerate the likelihood of the occurrence of events, depending on their personality and attitudes. For example, a person with an "optimistic attitude" would probably maximize gain in response to risk, because he would be more certain that everything would come out all right than someone with a pessimistic attitude. If the CBM rule set that biases COA evaluation under uncertainty has determined that the commander has an attitude which maximizes gain in response to risk, then the likelihood of game board scores (from the other rule sets) that are higher is increased in proportion to how high they are and by how much this personality maximizes gain in response to risk. If a commander has an attitude of minimizing loss, then the likelihoods of low scores are increased by how low they are. Finally, if the commander seeks maximum robustness, then higher scores are given to courses of action with values above a threshold across the different possible scenarios, to yield an answer that will be satisfactory. Only after the commander's personality biases the likelihoods is the "maximize expected gain" evaluation done.

2.2.3 Creating a COA Tree with Perception

In the JWARS COA selector, each agent will have a model of what his opponents *think* about him. This "second order" knowledge facilitates the modeling of deception that is so important to war. Figure 1 illustrates this process. When an agent first assesses a situation by sending JWARS data through his situation assessment rule set, several game boards are created, one for each likely enemy position along in relation to its own position. We call these the *agent's awareness game boards*. These game boards are then sent through the agent's model of the opponent's situation assessment rule set. Which doctrinal rule set is used depends on who the agent has identified the opponent to be. The better the agent's knowledge of the enemy, the better this identification can be made. For each of the agent's awareness game boards there are several game boards created from the opponent's perceived situation awareness rules. What they represent is the possible "awarenesses" an opponent might have. We will call them the *opponent's awareness game boards*.

Then, each of the agent's awareness game boards is sent through the response rules, and each of the opponent's awareness game boards for that board is sent through the agent's model of the opponent's response rules. We now have a few responses for an agent's awareness board, and at least as many opponent's responses. Each combination

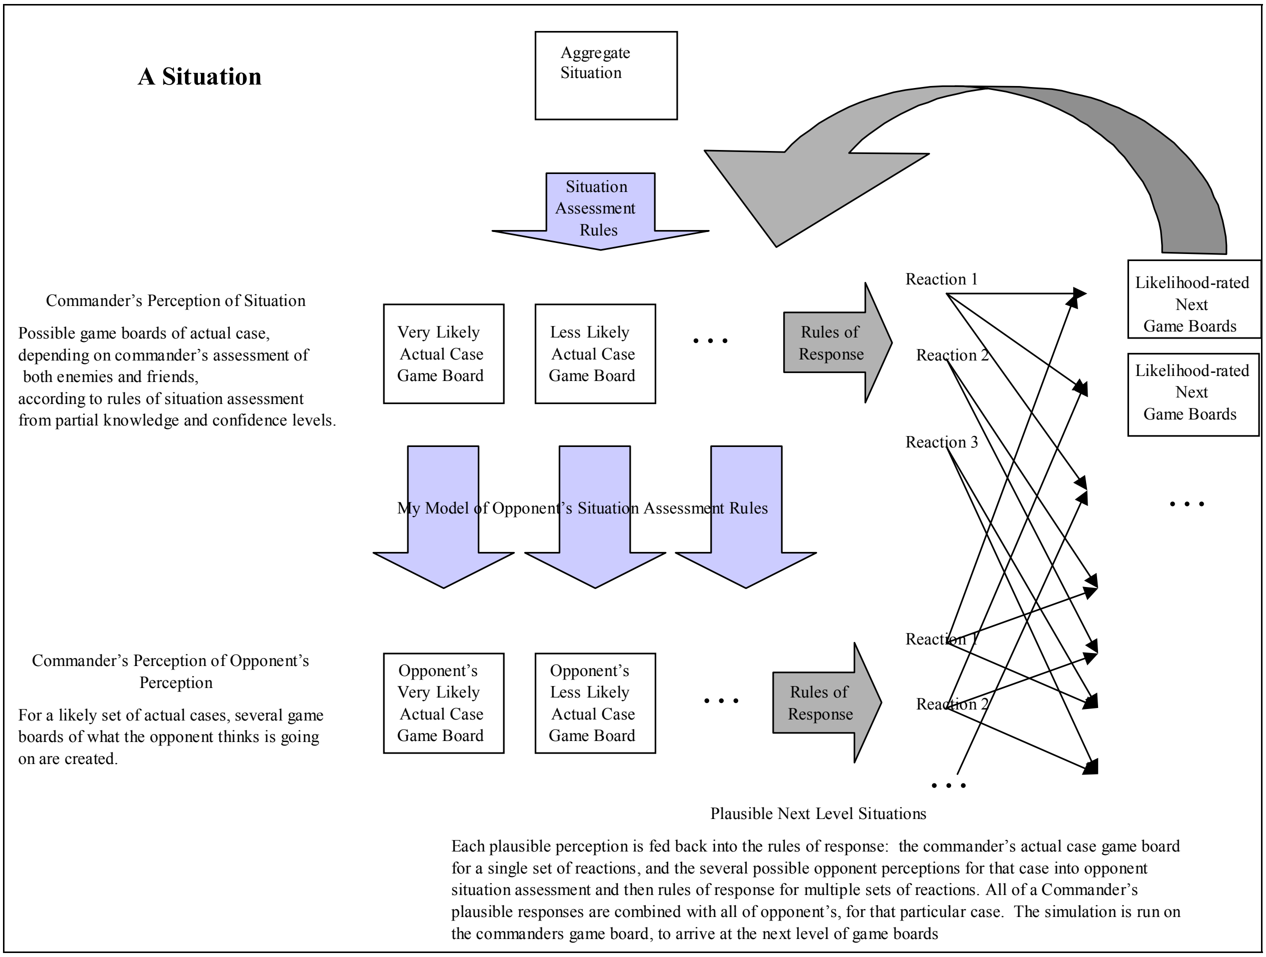

Figure 1: How a COA Tree is Created with Second Order Perception

is played out on a copy of that agent's parent awareness board, and these children become the next level of the tree. The resulting "successful" boards of one level become the agent's awareness boards for the start of the next level.

2.2.4 Strategic Scoring with Perception

Each leaf of the COA tree has a likelihood and a score from each commander. These scores are per board, and not per response. In order to send the scores back up the tree, we must convert the board scores to response scores. For example, suppose the bottom of the tree represented an agent's withdrawal and frontal attack responses, combined with an opponent's withdrawal, delay and frontal attack responses. To know what a specific response is worth, we compute its expected value from the likelihood of each of its boards, modified with CBM evaluation bias, times the score of those boards. The value for each side on the parent board is the value of the maximum response for that side, so that the response scores on the lower level become the board scores on the higher level. The process is repeated, and likelihoods of the higher level are given CBM bias, then used to compute response scores from board scores again, until the root is reached.

2.2.5 Modular Rule Sets

Doctrine, because it is written in human language, is too vague to be followed to the letter by a computer agent. There are many doctrinally correct responses that an agent can have to a complex situation. However, if the doctrine is artificially made more specific in response to this problem, then it cannot easily be mixed and matched with different social and behavioral contexts in a model. Our approach is to let the doctrine stay very much like it is in its own rule sets, and make it more specific to context by changing the agent's interpretation of it. In the JWARS COA selector, not only can a CBM personality be switched in and out to test its effectiveness in implementing doctrine, but its rules about the political, social, and cultural context can be switched in and out as well. These rule sets can be developed for individual cultures, societies, and political situations and can represent different theories of culture, politics and society just as the psychological rule set of the prototype represented a theory of personality to which test data was input.

As in the prototype, these modular rule sets will be interfaced with the rest of the system through concepts that are common to all human behavior theories in an area. For example, both the MMPI and the Myers-Briggs are based on personality theories that postulate a need for achievement. The concept of "need for achievement" could then become an interface with rules from a cultural rule set, which had something to say about cultural differences in what constitutes achievement. For example, some cultures might be represented by rules that say dying in battle is the highest achievement. It is through these "universal concepts" that the realms of the political, the cultural, and the economic may be treated modularly despite the fact that they are deeply connected.

The analyst may test the doctrine on a mix and match of different social theories, or she may combine two different theories by plugging them both in. Because the sets are fuzzy, they can be added together with no ill effects. It does not matter if the information within them is contradictory; the fuzzy engine can take all the information into account and still produce an outcome.

With modular rule sets, we could easily perform many kinds of tests of doctrine. For example, we could see how a specific CBM personality might act in another culture, or how our doctrinal responses would be seen in a different political climate. This capability will be especially important to the study of asymmetric warfare.

3 RESULTS

The JWARS commander model uses a course of action selector that examines many possible outcomes of maneuvers done by both the modeled commander and a postulated enemy commander, and determines what it considers the best COA for the commander. The simulation that represents the commander's wargame is done on a simplified and aggregated game board that represents a commander's internal thought processes. Factors such as greater military strength and mobility provide distinct advantages. A strong force can drive a weaker enemy force from an objective, but only if its supplies hold out; and a force with greater mobility may reach an objective faster and therefore hold onto it when another slower force attacks. Once a COA is selected, it is expanded into orders for subordinate units with individual missions, schedules, and guidance concerning required strength and external firepower, losses of personnel and equipment, and use of supplies. Intelligence resources are tasked to confirm that the enemy position and strength used in the COA selection continues to be true. Major deviation in any of the plan's major METT-T factors could result in a new situation assessment and trigger subsequent replanning.

The CBM model that has been developed allows a commander's personality to directly influence the selection of a course of action in the JWARS commander model. A fuzzy rule set provides a way for personality to have an effect on the course of action selected. The rule set affects how the selector chooses between different possible courses of action, and how the selector views his satisfaction with the end results of various courses of action. For example, in the case where a stronger force wishes to take an objective which a weaker force holds, an INTJ (Introverted-iNtuitive-Thinking-Judging) personality will typically attack and fight immediately while an ESFP (Extroverted-Sensing-Feeling-Perceiving) personality would typically spend more time in preparation and "shaping" the battlespace. Similarly, if the weaker force has a commander with an ESFP personality who perceives a major force ratio imbalance, he will withdraw when attacked to save the force, while a commander with an INTJ personality will stay and fight with the weaker force until actual casualties force him off the position. Note that these responses are based on the current rule set. Different rule sets might generate different responses and different perceptions almost always do.

The CM and CBM models demonstrate how different personalities can affect the selection of courses of action in a land battle. The models are flexible enough that with simple adaptations in their rule sets, they can be applied to a variety of other land situations, as well as to sea and air applications. Of particular interest might be applications to training simulations where a greater variety of enemy responses are desired without changing the enemy doctrine.

One challenge particular to JWARS is to generate COA without significantly slowing the simulation's run time. Without careful pruning of the game board tree, the computation time involved in the CM/CBM approach can quickly expand exponentially and potentially have a major impact on the JWARS' run time. Consequently, it is expected that the application of the CM/CBM will be tightly controlled and balanced against all of the other warfighting representations present in JWARS. Other applications may not have such stringent run time requirements and thus may freely expand both the scope and the depth of the CM/CBM COA selector searches.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the valuable feedback and assistance received on this paper from Dr. Jim Metzger, JWARS Director, LTC George Stone, III, USA, JWARS Technical Director, Dr. Jim Jones, CACI, JWARS Chief Systems Engineer, and Mr. George DeGovanni, CACI Vice President, Military Simulation and Analysis Division. Also contributing to this project were Maria Stropky and Judy Schandua, CACI object analysts, and James Huynh and Don MacQueen, GRCI software engineers. The authors further wish to express their appreciation to Dr. Ruth Willis, LtCol Eileen Bjorkman, USAF, and other members of the Defense Modeling and Simulation Office (DMSO) for their support on this effort. (Dr. Willis is now with the Office of Naval Research).

The opinions expressed in this document and its release for public consumption do not imply Department of Defense endorsement of either factual accuracy or opinion.

An earlier version of this paper has been published elsewhere (Vakas et al. 2001).

REFERENCES

JWARS Office. 2001. *Joint Warfare System Commander Behavior Model, Phase 1 Technical Report*, (2 volumes), Joint Warfare System Office, CACI under Contract No: DASW01-97-D-0041.

Vakas, D., J. Prince, H. R. Blacksten, and C. Burdick. 2001. Commander Behavior and Course of Action Selection in JWARS. *Proceedings of the Tenth Conference on Computer Generated Forces,* Orlando, Florida, Simulation Interoperability Standards Organization, 387-398.

AUTHOR BIOGRAPHIES

DEBORAH VAKAS is a senior cognitive scientist at CACI. She has 12 years of experience in the field of complex adaptive systems modeling of social and cognitive phenomena. She received her B.S. from Virginia Commonwealth University, her M.S. from the University of Alabama at Birmingham. She is currently completing her Ph.D. at George Mason University. Her email address is <dvakas@caci.com>.

JOHN D. PRINCE is a senior simulation engineer at CACI Inc – Federal. He received his Ph.D. from the University of Virginia in 1993. He is a member of the APS. His interests are in military simulations and artificial intelligence. His e-mail address is <jprince@caci.com>.

H. RIC BLACKSTEN is a mathematical operations researcher with CACI International, assigned to the JWARS Office. His research interests are mathematical modeling and stochastic simulation, military analysis. He received his B.S. from the University of Maryland and M.S. and Ap.Sc. degrees from The George Washington University. His email address is <rblacksten@caci.com>.

CHUCK BURDICK is a program manager with Lockheed Martin and is currently the leader of the JWARS Land Warfare Team. He has over 20 years of experience in simulating military systems and is a former Army MI officer and retired Colonel in the Army Reserves. He received his B.S. from Research Psychology from Rensselaer Polytechnic Institute (RPI) and his M.S.A. from George Washington University. His email address is <chuck.burdick@lmco.com>.