Behaviors that Emerge from Emotion and Cognition: A First Evaluation

Amy E. Henninger, Ph.D.
Soar Technology, Inc.
Ann Arbor, MI.

Randolph M. Jones, Ph.D.
Colby College
Waterville, ME.

Eric Chown, Ph.D.
Bowdoin College
Brunswick, ME.

ABSTRACT

This paper presents an initial evaluation of an emotions model developed for a sophisticated synthetic forces model. Sponsored by ARI, the objective of this research is to make the decision-making process of complex agents less predictable and more realistic, by incorporating emotional factors that affect humans. To this end, we have adopted an approach that promotes the emergence of behavior as a result of complex interactions between factors affecting emotions, integrated in a connectionist-style model, and factors affecting decision making, represented in a symbolic model.

This paper explains how we used the concept of emerging behaviors to test our framework. This includes a description of the behaviors we used in our prototype, the design of our experiments, a representative set of behavior patterns that emerged as a result of exercising our model over the design space, and our project's lessons learned.

ABOUT THE AUTHORS

Dr. Amy E. Henninger is Senior Scientist at Soar Technology, Inc. Prior to joining Soar Technology, Henninger worked in the modeling and simulation community as a Staff Scientist at SAIC, a Research Assistant at the Institute for Simulation and Training, and a Research Fellow for the U.S. Army Research Institute at U.S. Army STRICOM. She holds three BS degrees in Psychology, Industrial Engineering and Mathematics from Southern Illinois University; two MS degrees in Engineering Management and Computer Engineering, from Florida Institute of Technology and University of Central Florida (UCF), respectively; and a Ph.D. in Computer Engineering from UCF. Dr. Henninger currently teaches as an Adjunct Faculty Member in the Modeling and Simulation Graduate Program at UCF.

Dr. Randolph M. Jones received his B.S. in Mathematics and Computer Science from the University of California, Los Angeles, in 1984. In 1987 and 1989, respectively, he received M.S. and Ph.D. degrees in Information and Computer Science from the University of California, Irvine. He has held research positions at Carnegie Mellon University, the University of Pittsburgh, and the University of Michigan. He is currently Chief Scientist and Vice President of Soar Technology, Inc., and an assistant professor in the Computer Science Department at Colby College, where he is also department chair. His research interests lie in the areas of interactive intelligent systems, executable psychological models, and human and machine learning.

Dr. Eric Chown is an Assistant Professor of Computer Science at Bowdoin College. Previously he was a post-doctoral fellow at Oregon State University where he applied machine learning techniques to the problem of developing ecosystem models such as are involved in global climate change. His general areas of research include computational models of human learning and human cognitive mapping. He is one of the primary contributors to the SESAME cognitive architecture originally developed at the University of Michigan. He earned M.A. (1985) and M.S. (1987) degrees from Northwestern University, and a Ph.D. (1994) from the University of Michigan. Chown was awarded a National Science Foundation Career Grant in 2001 for his work on modeling the human spatial system and pioneering applications of that work to robotics and other domains.

Behaviors that Emerge from Emotion and Cognition: A First Evaluation

Amy E. Henninger, Ph.D.
Soar Technology, Inc.
Ann Arbor, MI.

Randolph M. Jones, Ph.D.
Colby College
Waterville, ME.

Eric Chown, Ph.D.
Bowdoin College
Brunswick, ME.

INTRODUCTION

A recent panel report sponsored by the National Research Council (Pew and Mavor, 1998) has called for the use of personality factors, behavior moderators and emotions to develop more realistic computer generated forces (CGFs). These recommendations have spawned a number of studies incorporating fatigue representations (French et al., 2001; Jones et al., 1997), defeat mechanisms (Heeringa & Cohen, 2000), personality paradigms (Hudlicka & Billingsley, 1999; McKenzie et al., 2001), and emotion models (Fransechini et al., 2001; Gratch & Marsella, 2001; Hudlicka & Billingsley, 1999) in CGF systems. While we currently know of no studies that have investigated whether military training is indeed improved by the use of CGFs with these capabilities, Army instructional courseware designers have recognized the significance of emotions in learning and training (Abell, 2000).

Classical Emotions Literature

The study of emotion has had a fickle history in psychology, appearing to correlate to prominent psychological theories of the day (Schultz, 1981). As such, there is currently no universally accepted, comprehensive theory of emotions. Instead, there exist a host of “mini-theories” that emphasize cognitive, motivational, physiological, and behavioral dimensions of emotion. For example, cognitive theorists tend to focus on thoughts and evaluations when defining emotions, motivational-theorists (Ferguson, 1982) tend to focus on goals and drive-reduction, physiologists on physiological reactions, and behaviorists (Lindsay and Norman, 1977) on emotional behavior, etc.

Common Emotional Constructs

While individual camps exist, there is a growing list of researchers (Lazarus, 1984; Ortony, 1988; Levine and Leven, 1992) who generally support the concept that emotional states can be manipulated by a combination of different factors. At a minimum, these factors include cognitive processes (expectations) and physiological states (usually interpreted as arousal or emotional intensity). Other factors have included environmental influences and behavioral expressions. These notions have led some psychologists, such as Lazarus (1984), to argue that emotion is multifaceted, and that all facets must be present in order to label something as an emotional state.

A second concept that is common to many emotion theories is the existence of a central evaluative mechanism that determines whether a given situation is potentially harmful or beneficial to the individual. For example, LeDoux and Fellous (1995) have discovered neural circuitry that processes stimuli according to whether they threaten or enhance the survival of the organism or its species. Also, a related discovery of an emotional memory system that works in concert with this circuitry has further added to the recent thrust of emotion research. Emotional memory has been associated with the amygdala and appears to add an “emotional flavor” to a declarative memory, which is thought to primarily originate in the hippocampus. This theory, exercised at its most primitive level, suggests that emotions are strong, “hard-wired” responses to stimuli that have a positive or negative survival value. The accompanying work on emotional memory suggests that these responses are mostly learned through classical conditioning (LeDoux, 1992) and performed as unconscious processes (Damasio, 1995).

EMOTIONAL IFOR FRAMEWORK

The models of emotions proposed in the psychological community are still in their formative stages. This gives rise to a system that is difficult to express in computational terms. After all, the integration of an emotion model in a computer agent ultimately requires the expression of that model in a formal and executable language, and making the translation from an imprecise psychological model to a formal computational language is an onerous task. However, as indicated earlier, there are some consistencies among the theories and their constructs, and it is our strategy to use these generally accepted common themes to the extent possible. In those cases where no one theme prevails, we adopt a more functional, physiologically based approach, as it tends to be more readily expressed in computational terms.

Our model of emotion and its interactions with cognition is based on a symbolic-connectionist hybrid architecture. A number of researchers have investigated and advocate the use of this type of architecture for a variety of cognitive modeling tasks (Sun & Alexander, 1997; Tan, 1997). In our system, cognition is represented within the Soar[1] software architecture. Soar has been under continuous development for over 18 years as a symbolic model of natural intelligence (Rosenbloom et al, 1993; Newell, 1990). It combines the abilities to react immediately to situations, use knowledge in deliberative decision making, step back from the immediate situation to perform various forms of problem solving and planning, and learn from experience. As an indicator of the maturity and utility of Soar-based entities, the system has been used successfully as the production model in a number of large-scale military exercises (Jones et al., 1999; Nielsen et al., 2000).

The connectionist model (Chown, 1993) is used to represent "emotional intensity", also called "arousal". In tune with modern theories of emotions (Damasio, 1987; LeDoux, 1992), we regard emotions essentially as subconscious signals and evaluations that inform, modify, and receive feedback from higher cognitive processes. Thus, our work explicitly distinguishes the subconscious processes (in a connectionist implementation) and the decision making that is subject to emotional influences (in a symbolic cognitive architecture). The following two sub-sections present an overview of the connectionist component representing emotion and the symbolic component representing cognition, respectively.

[1] The Soar architecture is in the public domain, with source code available at: http://ai.eecs.umich.edu/soar/

Detailed descriptions and computational representations of these components and their interactions may be found in Jones et al. (2002) and Henninger et al. (2001). After describing the model's components, we present a personality framework for exercising the model and a narrative of the scenario used to prototype the model.

Emotion: The Connectionist Component

As shown in Figure 1, the connectionist model consists of several interacting components: arousal, pleasure/pain and clarity/confusion.

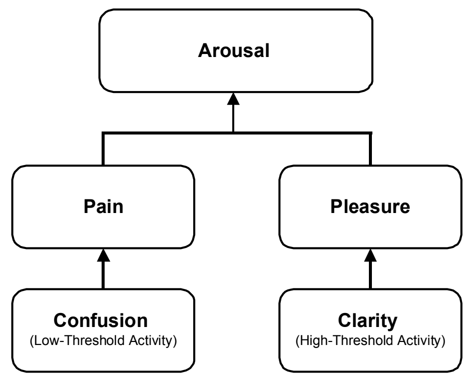

Figure 1. Block Diagram of Computational Arousal Mechanism

The clarity and confusion system, based on Kaplan et al. (1991), represents important correlates of pleasure and pain in forms of higher intelligence. Kaplan suggests that members of each species possess particular characteristics (e.g., agility, camouflage, etc) that facilitate their survival. Since humans are not particularly fast, fierce, or camouflaged, like other species are, Kaplan asserts that we rely on our ability to organize, store, and use information to enhance our survival. Consequently, for humans, confusion is a potentially dangerous attribute and clarity is a desirable attribute.

The pleasure/pain system interprets the level to which a stimulus represents a threat or enhancement to survival. In other words, stimuli that impede one's chances of survival would be tagged as painful and stimuli that would help one survive or reproduce would be tagged as pleasurable. This applies both to immediate sensations of physical pain as well as to deliberate predictions of situations and outcomes.

Whereas pleasure, pain, confusion and clarity all work to detect events of importance to an agent, the arousal system functions as a kind of interface between the emotional and higher cognitive systems. This relationship has been demonstrated by a number of researchers (e.g., D'Ydewalle, et al. 1985; Milner, 1991) documenting the effects of arousal on a variety of cognitive factors such as learning, memory, and attention.

In this framework emotions can be viewed as arising from a combination of pleasure/pain, clarity/confusion, and arousal. While other researchers opt to assign specific symbolic labels (Gratch, 1999; Ortony, Clore, and Collins, 1988), our emotion model does not make such high-level explicit assignments (e.g., "joy", "fear", etc). Instead we deliberately avoid use of these labels, opting to use continuously valued variables ranging between 0 and 1. However, "fuzzy" mappings between these attribute values and emotional labels can be made. For example, primitive mappings for use in facial gestures of animated characters (Wray et al., 2002) could include:

Clarity && Pleasure ➔ Joy
Clarity && Pain ➔ Anger
Confusion && Pleasure ➔ Surprise
Confusion && Pain ➔ Fear

Incorporating dimension of arousal, could yield:

Confusion && Pain && High Arousal ➔ Panic
Confusion && Pain && Low Arousal ➔ Anxiety

Or, incorporating dimension of time could yield:

Clarity && Pain && Past ➔ Regret

Cognition: The Symbolic Component

The two key parts of the symbolic component are an appraisal system and a response system. The appraisal system provides information to the connectionist emotions model, and the response system accepts information from the connectionist model. The theory behind these sub-systems is described in more detail below.

Response. The output of the emotional system includes numeric levels of pleasure, pain, and arousal. Arousal also creates a focusing of attention in the agent's higher-level cognitive processes. Soar, the architecture used to model the agent's higher-level cognitive processes, includes two types of memory (working memory and long-term memory), and arousal influences attention to elements in both of these memories. In working memory, arousal has similar consequences as it has for the agent's perceptual systems. That is, arousal tags particular working memory elements that are associated with the arousing stimulus. Since the "emotional" version of an agent has extra preference knowledge (for guiding the temporal order of decision-making actions) that prefers attending to such tagged elements, the agent tends to focus its attention on the active concepts that are perceived to be relevant to the current level of arousal. Accordingly, in low-arousal situations, the agent's normal preference knowledge makes decisions with less regard for any particular focus of attention.

In long-term memory, arousal focuses attention on retrieving well-rehearsed chunks of knowledge with which the agent is very familiar and experienced. Since humans prefer the comfort of using well-rehearsed knowledge in high-arousal situations, the agent will tend to revert to ingrained or default patterns of behavior when arousal is high. In general, these default patterns are more restricted than what the agent would exhibit in more calm and deliberative situations and they vary across agents, depending upon personality (discussed in the following sub-section) as well as background knowledge.

Appraisal. Following Gratch (1999), appraisal in our system is based around goals and expectations. The most straightforward types of appraisal require monitoring whether goals have been achieved, have become likely or unlikely to be achieved, or have been deemed unachievable. As the agent monitors its progress, the appraisal system signals events that feed into the connectionist system, such as goal failure and achievement, and interpretations of the environment that cohere (suggesting clarity) or confound (suggesting confusion). The appraisal system emphasizes relevance to the system's current goals, but it also reflects basic information processing needs, such as the dangers of not having enough knowledge. For example, typically, the agent should only be concerned about knowing the precise location of a particular enemy tank if that piece of knowledge is germane to the agent's current context. However, since the presence of enemy tanks is arousing and stimulates a need for information, an emotional agent will pay higher attention to the tanks. Thus, sometimes appraisal is cognitive (e.g., monitoring the progress towards an agent's mission goals) and at other times it is more reactive (e.g., monitoring progress of unexpected events).

One important form of appraisal revolves around an agent's expectations. An agent generates expectations as a result of past actions and situational interpretations. These expectations reside in working memory. If the expectations are violated, this suggests that the agent has a "confused" internal model, and the pertinent confusion inputs are candidates for processing. On the other hand, if the expectations are confirmed, this verifies the agent's internal model, and the pertinent clarity inputs would become candidates for processing. Since the clarity is preferred to confusion, the agent will seek situations where its expectations are confirmed.

Personality

Since different people have different reactions to the same situations (i.e., emotions and emotional responses are unique to individuals), we use the body of research in emotions and temperament to develop the bounds of an experimental design region for testing our model. Such differences can be thought of as an "emotional style" or temperament. Using this framework, we are able to account for individual differences in temperament by changing the connection strengths in the emotional subsystem. For example the psychological literature has long theorized that the critical factor that distinguishes introverts and extroverts is the relative susceptibility to becoming aroused (D'Ydewalle, et al., 1985; Eysenck & Eysenck, 1985).

Figure 2 illustrates the personality framework used to exercise the emotions model. Fundamentally, we model the introversion/extraversion dimension of personality by incorporating a susceptibility to arousal parameter in the emotions model, the neuroticism/stability dimension with susceptibility to pain, and the explorer/preserver dimension with susceptibility to confusion. Thus, by adopting this mapping, we are able to model an individual's emotional style such that it can lead to distinct decision making profiles in a variety of emotionally charged scenarios. Again, more information on these mappings can be found in Jones et al. (2002) and Henninger et al. (2001).

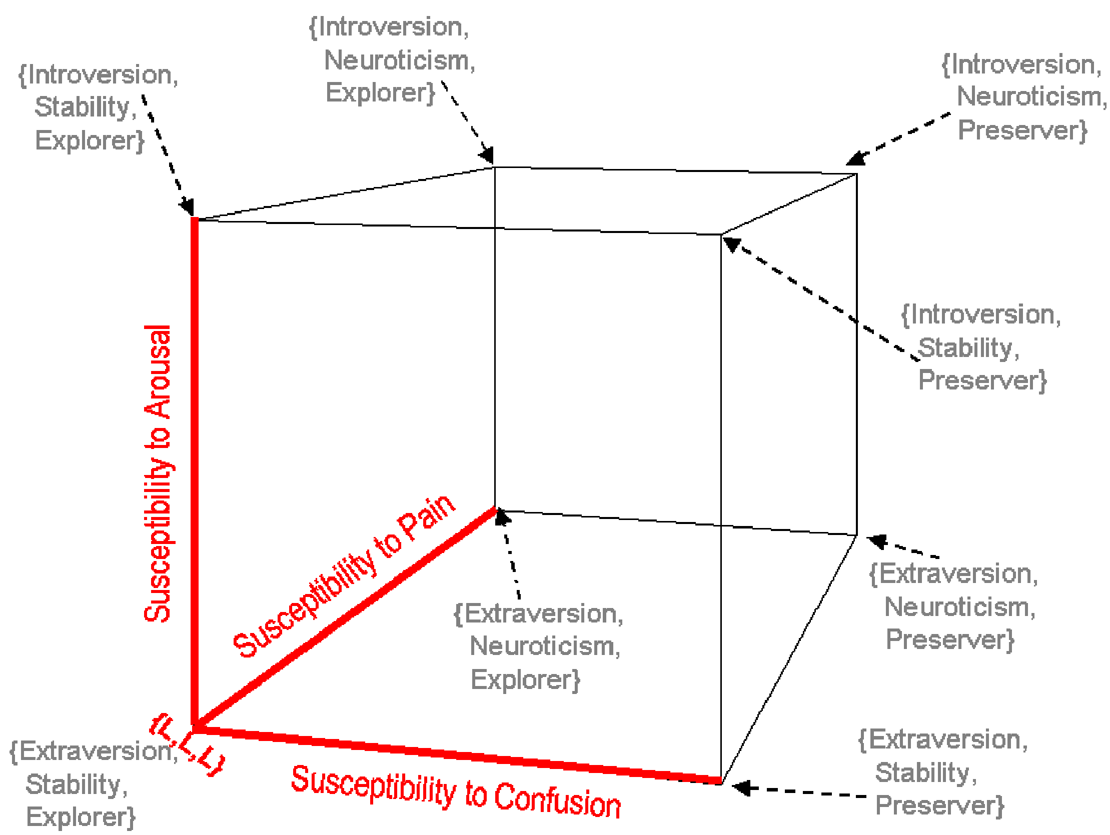

Figure 2. Personality Space used to Exercise Emotions Model

Prototype Scenario

The scenario used to prototype the emotions model was the Special Operations Forces (SOF)Soar Long Range Reconnaissance Mission. This task involves a 6-man team inserted deep within enemy territory. Once inserted, they travel anywhere from 20-50 km to the Objective Rally Point and split into three 2-man teams (i.e., two 2-man observation teams, and one 2-man radio team).

Figure 3 shows an example of how the physical mission parameters may be arranged. In this

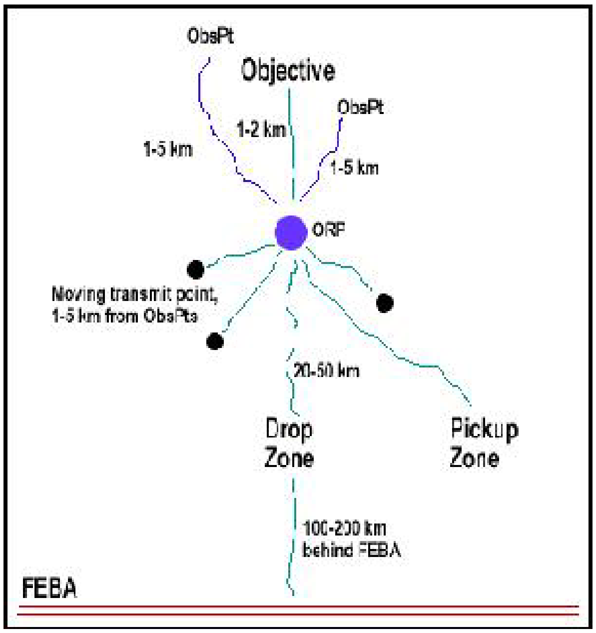

Figure 3. Example Map of Long Range Reconnaissance Mission Parameters

mission, there are 5 types of critical points (Drop, Rally, Observation, Transmit, and Pickup points) and one critical area, the Objective area. Currently, the prototype focuses on the behavior of a two-man team at an Observation point.

After splitting into 2-man teams, the observation teams seek cover and set up at the Observation point near the designated Objective area. Once an appropriate objective has been sighted, the team will report back to the radio team via wireline radios. At the conclusion of the mission, the teams will make their way to the Pickup Zone for exfiltration.

A run-time screen shot of our prototype may be seen in Figure 4, which shows the SOF Agents running in JSAF, an Agent's Soar Interface Panel, and an Agent's Emotions Interface Panel. The Soar Interface Panel enables operator control of an agent and communicates agent's decisions and actions. The Emotions Panel Interface is used to monitor the agent's emotional sub-systems and responses. To evaluate our system, we enhanced the SOF Agents with emotional responses, given some range of triggers. For example, we could alter the scenario by allowing detections, engagements, injuries, etc. The specific cases used in this prototype are presented in the next section.

EVALUATION FRAMEWORK

As described in the previous sections, there are a number of important interactions in our system: 1. the interaction amongst the emotional subsystems (i.e., arousal, pain/pleasure, clarity/confusion), 2. the interactions between the emotional system and cognition, and 3. the interactions between emotional entities. Taken together, the whole creates a fairly complex feedback system, in which the resulting external behavior would be very difficult to predict analytically (Riekel, 1995). This justifies the approach of building these models within executable intelligent agents, so that the resulting behaviors and emotional effects can be characterized empirically.

The multi-level, embedded nature of these interactions forms the foundation of a model that exhibits emerging behavior. Emerging behavior, broadly defined, is a modeling approach that attempts to generate high-level behavior of a system from the interaction of low-level rules and properties (Ilachinski, 1996). One form of an emerging behavior model that relates specifically to simulated agents is the individual-based model or IBM (Riekel, 1995).

Validating Models of Emerging Behavior

Because behaviors of IBMs exist only in software and are based on nonlinear and multi-threaded process controls, they are difficult to verify through conventional, analytical methods (Riekel, 1995; Ropella, 2002). For example, the emotional SOF Agents discussed in this paper exist only in software. And, like other models of emerging behavior, because this representation is theoretically based and lacks data for conventional forms of validation, we adopt principles of the IBM community (Riekel, 1995; Ropella, 2002) in validating the model. Essentially, this approach considers validation an issue of deciding whether the model output meets the required performance

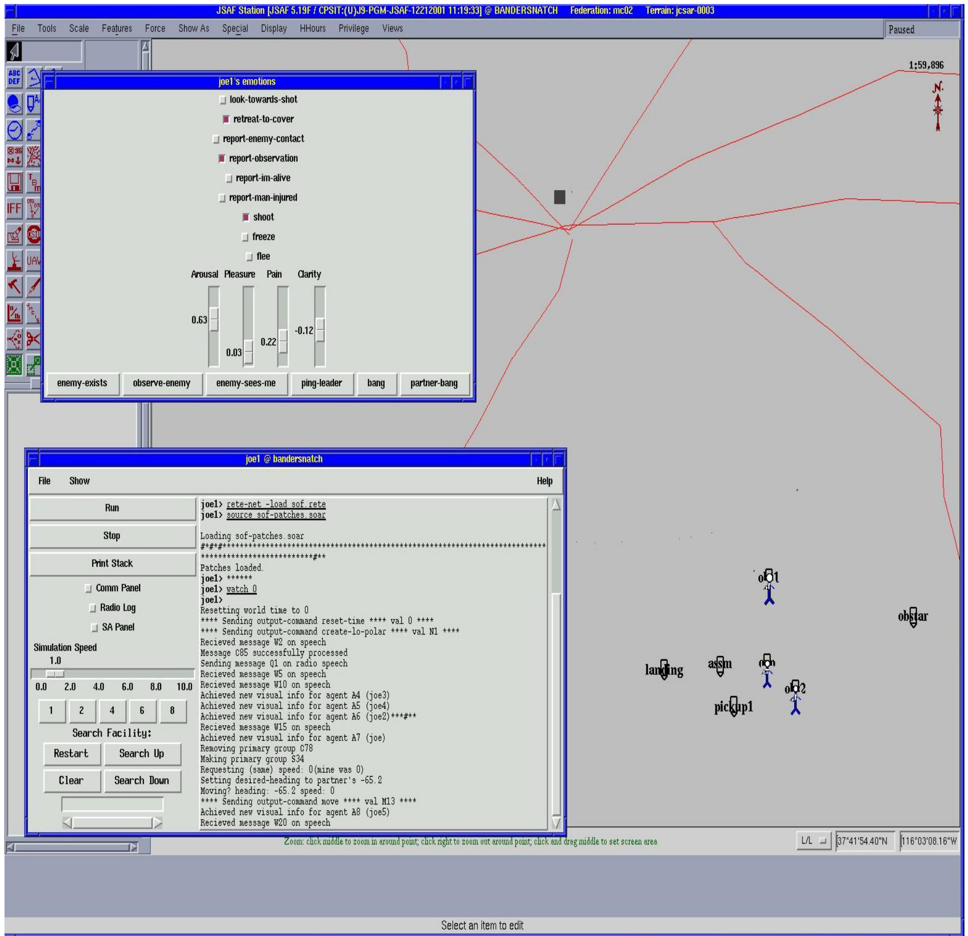

Figure 4. Runtime Screen Shot of Emotional SOF Soar Agent

standards according to the model's purpose. In the IBM community, this approach is known as "Operational Validation". Of course, another important part of this process is verifying the code, that is, checking that model does what it is intended to do.

Riekel offers a number of strategies for verifying and validating the software accuracy of IBMs or other similar types of models based on principles of emerging behaviors. Central to these strategies is a clear experimental procedure, where testing is treated as a scientific enquiry with testers designing experiments, predicting the outcomes, and then running the code to compare the actual outcomes to the predicted outcomes. Verification and validation strategies offered by Riekel include: code reviews, spot checks of key model subcomponents, pattern tests, and systematic tests against an independent implementation.

Validation Methodology

To test the model presented in this research, we adopt a dual approach, measuring the within-agent emerging behavior pattern (i.e., the interaction of the emotional model substructures and cognition for an individual agent, given a personality type) and the between-agent emerging behavior pattern (i.e., the interaction between agents, both of whom can have unique emotional states and temperaments). These tests are performed over a number of increasingly complex test cases. Initially, as a means of verifying model code, we consider static test cases to measure the within-agent patterns. These measures include verifying the emotional system's numeric output according to personality types, and verifying that the correct response is selected, given some state vector. All of these measures are compared against a manual simulation of each test case to determine accuracy. Follow-on tests evaluate the more complex cases measuring the dynamic behavior emerging from between-agent interaction. As in the simple case, these tests also systematically increase in complexity. The test cases are presented in the next section.

Test Case Scenarios

The general form to our testing procedure is to implement the simplest cases first and then allow the more complex cases to emerge as a result of the behaviors that have already been established and evaluated. For example, for the within-agent tests, each of the scenarios 1-5 (see Table 1) is

Scenario Title/Variation	Description
1. At Observation Point.	Two-man Observation team is stationed at the Observation Point and there is no enemy in sight.
2. At Observation Point and Enemy Detected.	Two-man Observation team is stationed at the Observation Point and a high number of enemy have been sighted.
3. At Observation Point, Enemy Detected, and Shooting.	Two-man Observation team is stationed at the Observation Point, spotted enemy, and detected shooting.
4. At Observation Point, Enemy Detected, Shooting, and Teammate Hit.	Two-man Observation team is stationed at the Observation Point, detected enemy, and agent's teammate has been shot.
5. At Observation Point, Enemy Detected, Shooting, and I'm Hit.	Two-man Observation team is stationed at the Observation Point, detected enemy, and the agent has been shot.

Table 1. Progression of Test Scenarios

used to statically evaluate the emotional system, in isolation. Thus, even in this simple block of static scenarios, a total of 160 tests are performed (emotional levels of each personality type for each scenario), assuming some baseline arousal level. Still focused on within-agent behavior, the next more complicated round of tests evaluates the first-order case where emotion and cognition interact. Again, this is accomplished by verifying model results with independently implemented manual simulations.

Once these initial results are verified through static tests, dynamic cases starting with scenario 1 progressing to variations of scenarios 4/5 are executed to record patterns in between-agent behavior. Differences in behaviors over these scenarios will be due to differences in agent's arousal level and how that impacts cognition, where these components of the model were previously verified in within-agent tests. Additional differences, still to be isolated through this latter set of tests, are the result of how the behaviors of one agent can impact the emotional intensity, and hence response, of another agent.

RESULTS

As stated in the previous section, the primary focus of the validation effort is to develop a sense of the model's utility, given its purpose. The objective of this research, as communicated in abstract, was to make the decision-making process of complex agents less predictable and more realistic, by incorporating an emotions model. The following sub-sections present results pertaining to both of these cases. In the first sub-section, we demonstrate the model's utility by presenting scripted output of simple test scenario. Next, in the second sub-section, we report on means of determining how this system reduces the predictability of an agent's behavior.

Model Output

Output behavior of one example of a simple dynamic case is scripted in Table 2. This test case contrasts the behavior of Agent assigned emotional style of <Extravert, Stability, Explorer> with an Agent assigned emotional style of <Introvert, Neurotic, Preserver>. In this scenario, the SOF Agents are at the Observation Point and detect enemy, the objective on which Agents should report. Up to this point, even though the Agents exhibit differences in the values of their emotional parameters, the Agents propose and select the same reactions based on the same world events. The next event, "Enemy-Sees-Me" causes Agent2 to propose one more action ("flee"). However, both Agents choose to "Retreat-to-Cover". During the next event, "Partner-Shot", behavior of the two Agents starts to diverge. That is, Agent1 chooses to report the injury, whereas Agent2 continues to seek cover. Lastly, as the "Shooting" (final event) continues, Agent1 remains active in seeking cover, whereas Agent2 "freezes", in essence rendering him useless in the rest of the scenario.

Reducing Agent Predictability

To measure the first objective we compared the range of the agent's response space using a classic, deterministic state-transition approach with the range of the agent's response space using our emotional model, which is also deterministic. Thus in a classical state-transition construct based on change in world state, as seen in equation 1, there

Event	Agent1(ESE)	Agent2(INP)
At ObsPt	Arousal = .54 Pleasure = .2 Pain = .0 Clar/Conf = .15 **Proposed action(s):** no change **Selected action:** no change	Arousal = .77 Pleasure = .59 Pain = .0 Clar/Conf = .35 **Proposed action(s):** no change **Selected action**: no change
Observe Enemy	Arousal = .54 Pleasure = .15 Pain = .06 Clar/Conf = .09 **Proposed action(s):** report-observation **Selected action:** report-observation	Arousal = .77 Pleasure = .43 Pain = .14 Clar/Conf = .21 **Proposed action(s):** report-observation **Selected action:** report-observation
Enemy-Sees-Me	Arousal = .54 Pleasure = .03 Pain = .19 Clar/Conf = -.03 **Proposed action(s):** retreat-to-cover, report-observation, shoot **Selected action:** retreat-to-cover	Arousal = .76 Pleasure = .07 Pain = .47 Clar/Conf = -.07 **Proposed action(s):** retreat-to-cover, report-observation, shoot, flee **Selected action:** retreat-to-cover
Partner-Shot	Arousal = .62 Pleasure = .00 Pain = .22 Clar/Conf = -.12 **Proposed action(s):** retreat-to-cover, shoot, report-man-injured **Selected action:** report-man-injured	Arousal = .93 Pleasure = .00 Pain = .62 Clar/Conf = -.28 **Proposed action(s):** retreat-to-cover, shoot **Selected action**: retreat-to-cover
Shooting	Arousal = .87 Pleasure = .0 Pain = .27 Clar/Conf = -.39 **Proposed action(s):** flee, retreat-to-cover **Selected action:** retreat-to-cover	Arousal = .98 Pleasure = .0 Pain = .63 Clar/Conf = -1.0 **Proposed action(s):** flee, freeze **Selected action:** freeze

Table 2. Example of Simple Dynamic Test

is some fixed number of outputs, given a current world state and an input.

$$\lambda_i = \lambda(q, E_i) \qquad (1)$$

where λ_i = set of outputs,

for any external input, E_i and any state, q

On the other hand, our approach, still viewed from perspective of state-transition construct, also selects output as a function of the world state and an input. In this case, however, that input is augmented by another state variable internal to the agent (e.g., arousal).

$$\lambda_j = \lambda(q, E_j, I_j) \qquad (2)$$

where I_j = is an input internally generated by Agent

To compare the two methods, we manually calculated the number of behaviors possible for each state in the prototype system for the classic approach and the emotions approach. Comparison of these numbers reveals an increased size in response space by average of 3.1, as shown in equation 3.

$$3.1(\lambda_i) \cong \lambda_j \qquad (3)$$

What makes this approach useful for generating less predictable behavior is the fact that the additional input is internal to the Agent and thus, not detectable by humans interacting with the scenario. So, for example, while it might be easy for a human participant to learn that "when X happens in the world, the agent will do Z", it is more difficult to learn that "when X happens in the world and the agent's emotional state is Y, the agent will do Z". Primarily, the reason this is difficult to predict is because "Y", the agent's emotional state, is not obvious to the human participant.

FUTURE WORK AND DISCUSSION

From a practical perspective, future work for this model includes tailoring it for use in phase II of the Virtual Technologies and Environments (VIRTE) project sponsored by the Office of Naval Research. In this capacity, our emotions model will be incorporated in to fully autonomous OPFOR agents populating a building clearing scenario (Wray et al., 2002). From a research perspective, future work includes continued experimentation to assist in understanding and improving the theory underlying this model. As we have attempted to make clear throughout this paper, our research is the implementation of one theory of emotions. Many theories exist. Also, other implementations of other theories exist (Gratch and Marsella, 2001; Franceschini et al, 2001; Hudlicka and Billingsley, 1999). We offer high-level comparisons of our approach with other approaches in Henninger et al. (2001) and Jones et al. (2002).

Ultimately, in the training community, the worth of these models must be measured in terms of improved training. That is, a student interacting with CGFs whose behavior is moderated by emotions should have better performance scores on some training task than a student interacting with "vanilla" CGFs. However, we know of no studies that have investigated whether military training is indeed improved by the use of CGFs with these capabilities. A less ambitious test for model performance has been to determine whether behavior moderators make CGFs appear more realistic (i.e., more humanlike). Of course, since field data to adequately perform such tests do not exist, this becomes a highly subjective measure. Moreover, adopting this measure of performance tacitly and baselessly advances the assumption that CGFs with "more humanlike" behavior (e.g., emotions) will improve training. Again, there is no evidence to support this assumption.

Based on our research, we believe that the incorporation of an emotions model into CGFs can make the behavior of the CGFs less predictable. But, we have not demonstrated that emotional CGF behavior is more realistic, nor have we demonstrated that the use of emotional CGFs will improve training. It is our opinion and recommendation that somewhere along this vein of research, funding agencies and sponsors of behavior moderator research formally investigate the assumed benefits of incorporating these models into CGFs.

ACKNOWLEDGEMENTS

This work was sponsored by the U.S. Army Research Institute's Emotional Synthetics Forces STTR – Ph II, contract number: DASW01-99-C-0037. That support is gratefully acknowledged.

REFERENCES

Abell, M. (2000). Soldiers as Distance Learners: What Army Trainers Need to Know. *Proceedings of the 2000 Interservice/Industry Training, Simulation, and Education Conference.* Orlando, FL.

Chown, E. (1993). The Adaptive Power of Affect: Learning in the SESAME Architecture. In J.A. Meyer, H.L. Roitblat, and S. Wilson (eds.), *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation Adaptive Behavior.* Cambridge, MA: MIT Press.

Damasio, A. R. (1995). *Descartes' error: Emotion, reason, and the human brain*. New York, NY: Avon.

D'Ydewalle, G., Ferson, R., and Swerts, A. (1985). Expectancy, Arousal, and Individual Differences in Free Recall. *Journal of Memory and Language*. pp. 519-525.

Eysenck, H.J. and Eysenck, M. (1985). *Personality and Individual Differences: A Natural Science Approach.* New York, NY: Plenum Press.

Ferguson, E.D. (1982). *Motivation: An Experimental Approach*. Malabar, FL: Krueger Publishing Company.

Franceschini, R., McBride, D., and Sheldon, E. (2001). Modeling the Vincennes Incident Using Affective Computer Generated Forces. *Proceedings of the 10th Conference on Computer Generated Forces and Behavioral Representation.* May 15-17, Norfolk, VA. pp. 65-75.

French, J., Wetteland, C., Bower, S., and Edgin, T. (2001). Estimating Crew Fatigue in Discrete Event Simulations. *Proceedings of the Summer Conference on Computer Simulations.* Orlando, FL.

Gratch, J. (1999). Why You Should Buy an Emotional Planner. *Proceedings of the Agents '99 Workshop on Emotion-based Agent Architectures (EBAA'99)* and *ISI Research Report ISI/RR-99-465*.

Gratch, J., and Marsella, S. (2001). Modeling Emotions in the Mission Rehearsal Exercise. *Proceedings of the 10th Conference on Computer Generated Forces and Behavioral Representation*. May 15-17, Norfolk, VA. pp. 457-465.

Heeringa, B. and Cohen, P. (2000). An Underlying Model for Defeat Mechanisms. *Proceedings of the 2000 Winter Simulation Conference*. Orlando, FL, pp. 933-939.

Henninger A.E., Jones, R.M., and Chown, E. (2001). A Symbolic-Connectionist Framework for Representing Emotions in Computer Generated Forces. *Proceedings of the 2001 Interservice/Industry Training, Simulation, and Education Conference*. Orlando, FL.

Hudlicka, E. and Billingsley, J. (1999). Representing Behavior Moderators in Military Human Performance Models. *Proceedings of 8th Conference on Computer Generated Forces and Behavioral Representation.* Orlando, FL.

Ilachinski, A. (1996). Land Warfare and Complexity, Part II: An Assessment of the Applicability of Nonlinear Dynamics and Complex Systems Theory to the Study of Land Warfare. Alexandria, Virginia, Center for Naval Analysis.

Jones, R. M., Henninger, A.E., and Chown, E. (2002). Interfacing Emotional Behavior Moderators with Intelligent Synthetic Forces. *Proceedings of the 11th Conference on Computer Generated Forces and Behavioral Representation*. Orlando, FL.

Jones, R. M., Laird, J. E., Nielsen, P. E., Coulter, K. J., Kenny, P., and Koss, F. V. (1999). Automated intelligent pilots for combat flight simulation. *AI Magazine.* Vol(20), no(1), pp. 27-41.

Jones, R.M., Laird, J.E., and Neville, K. (1997). Modeling pilot fatigue with a synthetic behavior model. *Proceedings of the 7th Conference on Computer Generated Forces and Behavioral Representation. Orlando, FL.* pp. 349-357.

Kaplan, S. (1991). Beyond rationality: Clarity-based decision making. In T. Garling and G. Evans (eds.), *Environment, Cognition, and Action: An Integrative Multidisciplinary Approach*. pp. 171-190. New York, NY: Oxford University Press.

Lazarus (1984). On the primacy of cognition. *American Psychologist*. Vol(39), pp. 124-129.

LeDoux, J. E. and Fellous, J. (1995). Emotion and Computational Neuroscience. *The Handbook of Brain Theory and Neural Networks.* Michael A. Arbib, editor. Cambridge, MA: The MIT Press, pp. 356-360.

LeDoux, J.E. (1992). Brain Mechanisms of Emotion and Emotional Learning. *Current Opinions in Neurobiology*. Vol(2), pp. 191-197.

Levine, D.S. and Leven, S. J., (1992). *Motivation, Emotion, and Goal Direction in Neural Networks.* Hillsdale, NJ: Erlbaum.

Lindsay, P.H. and Norman, D.A. (1977). *Human Information Processing.* New York, NY: Academic Press, Inc.

McKenzie, F., Catanzaro, J., and Petty, M.D. (2001). A Personality-Based Command Decision-Maker: Results and Recommendations. *Proceedings of the 10th Conference on Computer Generated Forces and Behavioral Representation.* May 15-17, Norfolk, VA. pp. 351-359.

Milner, P.M. (1991). Brain Stimulation Reward: A Review. *Canadian Journal of Psychology*, vol(45), pp. 1-36.

Newell, A. (1990). *Unified Theories of Cognition.* Cambridge, MA: Harvard University Press.

Nielsen, P., Smoot, D., Martinez, R., and Dennison, J. (2000). Participation of TacAir-Soar in Road Runner and Coyote Exercises at Air Force Research Lab, Mesa, AZ. *Proceedings of the 9th Conference on Computer Generated Forces and Behavioral Representation*. Orlando, FL. pp. 173-180.

Ortony, A., Clore, G.L., and Collins, A. (1988). *The Cognitive Structure of Emotions*. New York, NY: Cambridge University Press.

Pew, R.W. and Mavor, A.S., eds. (1998). *Modeling Human and Organizational Behavior: Application to Military Simulations.* Washington, DC: National Academy Press.

Rosenbloom, P., Laird, J., and Newell, A. (1993). *The Soar Papers: Research on Integrated Intelligence.* Cambridge, MA: MIT Press.

Ropella, G.E.P., Railsback, S.F., and Jackson, S. K. (2002). Software Engineering Considerations for Individual-based Models. *Natural Resource Modeling.* Vol(15), no(1).

Rykiel, E. (1995). Testing ecological models: the meaning of validation. *Ecological Modeing. Vol(*90), pp. 229 – 244.

Schultz, D. (1981*). A History of Modern Psychology.* New York, NY: Academic Press.

Sun, R. and Alexandre, F. (1997). *Connectionist-Symbolic Integration: From Unified to Hybrid Approaches.* Mahwah, NJ: Lawrence Erlbaum Associates.

Tan, A. H. (1997). Cascade ARTMAP: Integrating Neural Computation and Symbolic Knowledge Processing. *IEEE Transactions on Neural Networks.* Vol(8), no(2), pp. 237-250.

Wray, B., Laird, J., Nuxoll, A., and Jones, R. (2002). Intelligent Opponents for Virtual Reality Trainers. To appear in *Proceedings of the 2002 Interservice/Industry Training, Simulation, and Education Conference.* Orlando, FL.