

Department of Computer & Information Science
Center for Human Modeling and Simulation

University of Pennsylvania

Year 2002

Constructing Virtual Asymmetric
Opponents from Data and Models in the
Literature: Case of Crowd Rioting

Barry G. Silverman* Michael Johns† Kevin O'Brien‡
Ransom Weaver** Jason Cornwell††

*University of Pennsylvania, basil@seas.upenn.edu

†University of Pennsylvania

‡University of Pennsylvania

**University of Pennsylvania

††University of Pennsylvania

Postprint version. Published in *Proceedings of the 11th Conference on Computer Generated Forces and Behavioral Representation, SISO 2000*, 10 pages.

This paper is posted at ScholarlyCommons.

<http://repository.upenn.edu/hms/1>

Constructing Virtual Asymmetric Opponents from Data and Models in the Literature: Case of Crowd Rioting

Barry G. Silverman^{1,2,3}, Michael Johns², Kevin O'Brien³, Ransom Weaver⁴, Jason Cornwell³

1 – Systems Engineering Department and Institute for Research in Cognitive Science

2- Computer Science Dept. and Human Modeling & Simulation Lab

3 - Ackoff Center for Advancement of Systems Approaches (ACASA)

4 – Department of Asian and Middle Eastern Studies

University of Pennsylvania
Philadelphia, PA 19104-6315
barryg@seas.upenn.edu

KEYWORDS:

asymmetric adversary agents, stress, emotion, utility

ABSTRACT: *This paper describes an effort to integrate human behavior models from a range of ability, stress, emotion, decision theoretic, and motivation literatures into a game-theoretic framework appropriate for representing synthetic asymmetric agents and scenarios. Our goal is to create a common mathematical framework (CMF) and an open agent architecture that allows one to research and explore alternative behavior models to add realism to software agents - e.g., physiology and stress, personal values and emotive states, and cultural influences. Our CMF is based on a dynamical, game-theoretic approach to evolution and equilibria in Markov chains representing states of the world that the agents can act upon. In these worlds the agents' utilities (payoffs) are derived by a deep model of cognitive appraisal of intention achievement including assessment of emotional activation/decay relative to value hierarchies, and subject to (integrated) stress and related constraints. We present the progress to date on the mathematical framework, and on an environment for quickly editing opponents in terms of the various elements of the cognitive appraiser, utility generators, value hierarchies, and Markov chains. We illustrate the approach via an example training game for counter-terrorism and crowd management. Future research needs are elaborated including validity issues and ways to overcome the gaps in the behavioral literatures that confront developers of asymmetric forces.*

1. Introduction

A common concern amongst agent developers is to increase the realism of the agents' behavior and cognition. In training, wargaming, and operations rehearsal simulators there is a growing realization that greater cognitive subtlety and behavioral sensitivity in the agents leads to both (1) a greater ability to explore alternative strategies and tactics when playing against them and (2) higher levels of skill attainment for the human trainees: e.g., see [1] and [2]. For this to happen, the tactics, performance, and behavior of agents must change as one alters an array of behavioral and cognitive variables. As a few examples, one would like agent behavior to realistically change as a function of: the culture they come from (vital for mission rehearsal against forces from different countries); their level of fatigue and stress over time and in different situations; and/or the group effectivity in, say, the loss of an opposing force's leader. At present, however, this does not happen, and in most of the available combat simulators the agents conduct operations endlessly without tiring, never make mistakes of judgment, and uniformly (and predictably) carry out

the doctrines of symmetric, sometimes vanquished opponents, such as the Warsaw Pact, among others.

Closely related to the topic of emulating human behavior is that of "believability" of agents. The basic premise is that characters should appear to be alive, to think broadly, to react emotionally and with personality to appropriate circumstances. There is a growing graphics and animated agent literature on the believability topic (e.g., see [3], [4] and [5]), and much of this work focuses on using great personality to mask the lack of deeper reasoning ability. However, in this paper we are less interested in the kinesthetics, media and broadly appealing personalities, than we are in the planning, judging, and choosing types of behavior -- the reacting and deliberating that goes on "under the hood" of embodied agents. Finally, and perhaps most importantly the human behavior literature is fragmented and it is difficult for agent developers to find and integrate published models of deeper behavior. Our research involves developing an integrative framework for emulating human behavior in order to make use of published behavioral results to construct agent models. We are not attempting basic research on how humans think but on how well existing models might work

together in agent settings. That is, the framework presented here is intended for experiments on how to integrate and best exploit published behavioral models, so as to improve the realism of agent behaviors when one seeks to model individual differences such as stress, emotion, and culture.

In particular, we are interested in *emergent macro-behavior due to micro-decisions of bounded-rational agents* and with developing a framework that promotes the study of specific phenomena (i.e., emotions, stress, and cultural values) that lead to limits of rationality. What motivates agents to select actions that sub-optimize their own utility as well as that of groups whose causes they seek to advance? To explore this question, we have been researching a framework that allows one to investigate the duality of mind-body interaction in terms of the impact of environment and physiology on stress and, in turn, of stressors on rationality. Our framework also attempts to integrate value systems and emotion-based appraisals of decision options along with the stress constraints. That is, we have been working towards a framework that permits one to examine the *impacts of stress, culture, and emotion upon decisionmaking*. With such a framework, one should, as an example, be able to readily model and visually render what makes one protesting crowd throw stones while another peacefully demonstrates.

As soon as one opens the door to modeling the impact of stress, culture, and emotion on rationality, one must be amenable to the idea that competing views, results, models, and approaches have to be examined and potentially integrated. The point of such a research program should not be to argue for one approach or theory over another, but to provide ways to readily study alternative models of whatever contributes to the phenomena of interest.

1.1 Role of Emotion and Concern Ontologies in Agent Behavior

“Emotive computing” is often taken to mean the linking of the agent state to facial and body expressions, vocal intonation, and humorous or quirky animation effects: e.g., see [6], [7] and [4]. However, recent theories identify emotions as vital to the decision-making process and to manage competing motivations [8]. According to these theories, integrating emotion models into our agents will yield not only more believable decision-makers, but also more realistic behavior by providing a deep model of utility. These agents will delicately balance, for example, threat elimination versus self-preservation, in much the same way it is believed that people do. These theories suggest that without adding emotional construal of events, the agents won’t know what to focus upon and what to ignore, and won’t know how to balance the set of next-

step alternative actions against larger concerns, as in the case of Damasio’s pre-frontal cortex damaged patients who spend the entire day mired in highly logical decision analyses of banalities, even at the cost of their own self-interest and survival.

Important implementations of these ideas and theories were attempted in the “believable agents” movement such as [4] and [5] which seek to improve the believability of characters’ behavior in fictional settings with the help of an emotion model. The OCC model is probably the most widely implemented of the emotion models (e.g., [9], [10] and [11]) and it explains the mechanisms by which events, actions, and objects in the world around us activate emotional construals. In both Oz [4] and the Affective Reasoner [5] projects, emotion was largely modeled as a reactive capability that helped characters to recognize situations and to reflect broad and believable personality characteristics. Later versions of Oz include a behavior planner, but the link between emotion construals and behavioral choice is never well articulated in their published accounts. On the other hand, [12] and [13] concretely extend the OCC model via the use of an event planner into a deeper, deliberative reasoning mode where agents were able to construe the value of plans and plan elements (events that haven’t happened yet). In the current paper, we extend this still further so that agents can construe the value not only of plan elements (future events), but so they also can construe the impact of objects and behavior standards both on themselves and on those they like/dislike. We go beyond this too to the area of what is probably unconscious construals of stressors such as fatigue, time pressure, and physiological pressures. This means we attempt a fairly full implementation of the OCC model for reactions and deliberations of all types of events, actions, and objects.

This approach provides a generalizable solution to another issue in the OCC model. The OCC model indicates what emotions arise when events, actions, or objects in the world are construed, but not what causes those emotions or what actions an agent is likely to take as a result. There is no connection between emotion and world values, even though other theories suggest such a link [8], [10] and [11]. In contrast, concern or value ontologies are readily available in the open literature (e.g., the ten commandments or the Koran for a moral code, military doctrine for action guidance, etc.) and may readily be utilized to implement an agent of a given type in the framework we present here. Ideally, one would like to tie such concern ontologies indirectly to the emotional processes of the agent, so that situation recognition as well as utilities for next actions are derived from emotions about ontologies and so that both reacting and deliberating (judging, planning, choosing, etc.) are affected by emotion.

2. Cognitive Architecture and Framework

The research described here is not to propose the best cognitive architecture or agent algorithms but to propose a reasonable framework within which the many contributions from the literature can be integrated, investigated, and extended as needed. That framework includes four somewhat arbitrarily separated subsystems plus a memory that form the stimulus-response capability of an agent as shown in Figure 2. There are a large number of similar frameworks in the literature: e.g. a useful comparison of 60 such models may be found in Crumley & Sherman [14]. The model we depict here shows an agent that receives stimuli and formulates responses that act as stimuli and/or limits for subsequent systems. The flow of processing in a purely reactive system would be counter-clockwise starting at the “stimuli” label, however, we are also interested in a deliberative system, one that can ponder its responses and run clockwise from the “cognitive system” to seek stimuli to support alternative response testing.




Figure 1 – Top Level of the Integrative Architecture for Researching Alternative Human Behavior Models for Generic Agents

The agent model of interest to us is that of a modified Markov Decision Process (MDP). That is, the agent seeks to traverse a hierarchical and multi-stage Markov chain which is the set of nested games such as the one depicted partially in the case study (Sect. 3). In order for the agent to be aware of this chain one would need to place it into the agent’s working memory as $G(A,C)$, a set of possible goals and tasks that the agent might wish to work its way through as the game unfolds. More broadly, working memory should store and process beliefs, desires, and intentions. In keeping with the BDI agent model, the

beliefs are those processed in the game theoretic sense of observing the world and of forming and remembering simple statistical models of the actions of those near us in the situation of interest. Desires are not well-defined in the BDI model, so here we define them as the future-focused affective states of hope and fear as generated by the emotion system (Section 3.3). Intentions are the planned actions and sets of orders that the agent is seeking to carry out (a_{mn} in A).

2.1 Stress and the Physiological Subsystem

The physiological subsystem of Figure 1 initially reacts to a set of stimuli that are perceived from and/or experienced in the environment. This subsystem includes all sensory apparatus, but also grouped into here are a number of physical processes that may be thought of as reservoirs that can be depleted and replenished up to a capacity. At present we model eight physiological reservoirs or stressors, including: energy, sleep, nutrients, noise and light impacts, and other physical capacities: [15] provides more detail. For each of these there are a large number of stressors that moderate an agent’s ability to perform up to capacity, and that in some cases send out alarms, for example when pain occurs or when other thresholds are exceeded (e.g., hunger, fatigue, panic, etc.). An important criterion for such a module is that it supports study of common questions about performance moderators: e.g., the easy addition or deletion of reservoirs of interest to a given study or training world (e.g., pain from virtual injuries, stress from proximity to land mines, etc.), individual differences in reacting to the same stressors, and/or how to model reservoir behaviors either linearly (our present approach) or non-linearly such as with bio-rhythms. Another vital criterion for such a module is that it should support studying alternative mechanisms for combining the many low level stressors and performance moderator functions into a single stress level. It is the overall stress that effects each of the other subsystems, and one would like a framework that shows how to compute an integrated level and then each of the subsequent modules need capabilities to reflect how their functioning is effected – emotions about stress, judgments under stress, and stressed motor/expressive acts.

In particular, we model integrated stress or iSTRESS as a result of three prime determinants – (1) event stress (ES) which tracks agents’ adverse and positive events, (2) time pressure (TP) which is a normalized ratio of available vs. required time for the tasks at hand, and (3) effective fatigue (EF) which integrates a normalized metric based on current level of many of the physiological reservoirs. Each of these is quantitatively derived and then emotionally filtered since a stoic will construe the same facts differently than a nervous type. The next section describes the emotional filtering. The quantitative factors

that go into these modifiers are then summarized via the following where $f\{.\}$ is currently a linear additivity model:

$$iSTRESS(t) = f\{ES(t), TP(t), EF(t)\} \quad [1.0]$$




Figure 2 - The Classic Performance Moderator Function is an Inverted-U

It is one thing to quantitatively derive an integrated metric called *iSTRESS*, but it is another to interpret its meaning and to translate that meaning into overall agent coping style. The approach we've adopted for accomplishing this translation is derived from Janis & Mann [16] who provide what is probably the most widely cited taxonomy of decision strategies for coping under stress, time pressure, and risk. We interpret this taxonomy as the steps of the inverted U-curve of Figure 2 and define it below. The taxonomy includes a decisional balance sheet that indicates how stress, time pressure, and risk drive the decision maker from one coping strategy to another and we depict these items across the X-axis of Figure 2.

In particular, we use the framework without further elaboration here to label the cutoff points for the integrated stress, or the *iSTRESS* variable and to constrain the decision making since a given stress level dictates the agent's ability to collect and process both information and action alternatives ($a \in A$) when in a given state, s .

All but the third of the coping patterns vigilance regarded by Janis & Mann [16] as "defective." The first two, while occasionally adaptive in routine or minor decisions, often lead to poor decision-making if a vital choice must be made. Similarly, the last two patterns may occasionally be adaptive but generally reduce the DM's chances of averting serious loss. The authors note, vigilance,

although occasionally maladaptive if danger is imminent and a split-second response is required, generally leads to decisions of the best quality". Some authors have since refined these ideas as with Klein et al. [17] who shows that experts work effectively in the "near panic" mode where they immediately recognize a best or near best alternative without vigilant scanning of other alternatives.

Unfortunately, Janis & Mann [16] do not provide either (1) precise threshold values (Ω_i) that indicate when decision makers trigger a change in coping style, or (2) any insight into how to integrate the many diverse stimuli, factors, or PMFs that determine stress and time pressure or risk. For these purposes, at present we use logic rules to combine these three factors. For example, such rules must account for facts such as a Very High value of any one of the factors could push the agent to panic. However, panic is more likely if at least one factor is very high and another is high. Or alternatively, if one factor is very high and both of the others are moderately high, panic might also result.

The results of physiology and stress are thus a bounding on the parameters that guide the agent's decision or cognitive subsystem and that dictate the coping style it is able to select. These parameters and decision style constraints do not in themselves provide any guidance on how to construe the situation, on the sense-making that needs to go on. For that we turn to the emotion subsystem.

2.2 Emotion Appraisal as a Deep Model of Utility

In particular, the emotion subsystem receives stimuli from the sensors as adjusted and moderated by the physiological system. It includes a long term associative or connectionist memory of its concern ontologies that are activated by the situational stimuli as well as any internally recalled stimuli. These stimuli and their impact on the concern ontologies act as releasers of alternative emotional construals and intensity levels. These emotional activations in turn provide the somatic markers that serve as situation recognition and that help us to recognize a problem that needs action, potential decisions to act on, and so on. In order to support research on alternative emotional construal theories this subsystem must include an easily alterable set of activation/decay equations and parameters for a variable number of emotions. Further, since construals are based on concern ontologies, this module must serve as a concerns ontology processor and editor. Simply by authoring alternative concern ontologies, one should be able to capture the behaviors of alternative "types" of people and organizations and how differently they would assess the same events, actions, and artifacts in the world. This requires the emotion module to derive the elements of

utility and payoff that the cognitive system will use to access alternative actions.

In the next section we will examine how to combine multiple emotions into a utility estimate for a given state. For now we will only examine how our different emotions arise when confronted by a new state, s , of the world, or in reaction to thinking about being in that state. In general, we propose that any of a number of ξ diverse emotions could arise with intensity, I , and that this intensity would be somehow correlated to importance of one's values or concern set (C) and whether those concerns succeed or fail for the state in question. We express this as

$$I_x(s_k) = \sum_{j \in J_x} \sum_{c \in C_{ijkl}} [W_{ijl}(c) * f1(r_j) * f2(O,N)] \quad [2.0]$$

Where,

$I_{\xi}(s_k)$ = Intensity of emotion, ξ due to the k th state of the world

J_{ξ} = The set of all agents relevant to x . J_1 is the set consisting only of the self, and J_2 is the set consisting of everyone but the self, and J is the union of J_1 and J_2 .

$W_{ij}(C_{ijkl})$ = Weighted importance of the values of agent j that succeed and fail in one's i th concern set.

C_{ijkl} = A list of paths through the i th ontology of agent j triggered to condition l (0 =success or 1 =failure) by state k .

$f1(r_{jk})$ = A function that captures the strength of positive and negative relationships one has with the j agents and objects that are effected or spared in state k

$f2(O,N)$ = A function that captures temporal factors of the state and how to discount and merge one's emotions from the past, in the present, and for the future

This expression captures the major dimensions of concern in any emotional construal – values, relationships, and temporal aspects. For the sake of simplicity, we assume linear additivity of multiple arousals of the same emotion from the $i=1, I$ different sets of values that the state may precipitate.

There are several emotion models from the psychology literature that can help to provide greater degrees of detail for such a model, particularly a class of models known as cognitive appraisal theories. These include the models mentioned earlier [9], [10] and [11] that take as input a set of things that the agent is concerned about and how they were effected recently, and determine which emotions result. Most of them fit into the structure of equation 2.0 but they have different strengths to bring to bear. At present we have decided to pursue the OCC model [9] to see how it helps out. In the OCC model, there are 11 pairs

of oppositely valenced emotions (ξ). One pair we use here as an example is pride-shame. Another pair we mentioned earlier was hope-fear for future events. One can experience both emotions of a given pair at the same time and if their intensities are equal, they cancel out from a utility perspective.

The OCC model assumes a decision making agent has 3 types of concern trees about the world: goals for action, standards that people should follow, and preferences for objects. Let us suppose as in Figures 3a & b that we have a terrorist agent who has two concern trees (let $|C| = 2$): one for standards ($i=1$) about how agents should act and one for preferences about objects or artifacts in the world ($i=2$). Of course any such agent would have many more concern trees and each might be more richly filled in, but these will suffice for the sake of the example. And in fact, the stopping rule on filling in concern trees for any agent is the limit of what behavior is needed from them in the scenario or micro-world in question. One can see from Figure 3 that concern trees bottom out in leaf nodes that can be tested against elements (events, actions, nearby objects, etc.) of the current state, k . Further, concern trees hold an agent's previously learned values or importance weights. Each link of a concern tree is labeled with a weight, w , and the sum of child weights always sums to 1.0 for the sake of convenience. The children can be either strictly or non-exclusively conjunctive or disjunctive.

Thus far in our research we have derived the structure and weights on these trees manually as part of the process of building agents for a given micro-world, though one could in principle derive these trees via machine learning and knowledge discovery when interacting with a news event dataset about a given terrorist group. The way we use these trees in Equation 2.0 is as an evaluation function for W_i . That is, when a given state of the world causes a leaf node to fail or succeed, that leads to the w_i being multiplied together up the branch of the tree from leaf node to root, and the overall W_i of that concern tree is computed. We gave details of how this works in Johns & Silverman [18].

Consider how the use of the trees of Figure 3a&b result in the weighting on a strategy resulting in being dead. Upon the agent contemplating his death (k ="dead"), no preferences are caused to succeed or fail by being dead. Consequently, no preference-based emotions would be generated from this agent's object preference ontology. However, k ='dead' does effect the agent's standards tree and one standard ($i = 1$) directly succeeds and one fails. He feels pride at having attempted his mission (c ="attempt current mission") for two reasons: he has fulfilled his commitment to the organization, and has attempted something to correct a perceived injustice.




Figure 3 – Concern Ontologies Showing Part of the Standards and Preferences of a Sample Terrorist

However, his mission involved returning home safely, which is clearly thwarted by failing to survive. Consequently, he will feel shame at his incompetence as well.

On balance, in the current state, pride slightly outweighs shame at being a martyr. Whether an agent's decision subsystem would choose death, however, is also a function of its iSTRESS or Ω level and of its current goal tree construals, a topic we omitted from this example due to space considerations, though we illustrate a goal tree construal in Sec.3. Also omitted from this discussion are several other dimensions of the agent's reasoning in social situations, a few examples of which are: (1) construing relationships to others in the scenario that the agent likes, dislikes, etc.; (2) explicit modeling of partial knowledge of the emotions of those others to further guide his own actions; (3) assigning credit/blame to others for various actions and events; and (4) managing likelihood and temporal factors. The OCC model provides a number of

inroads into how to handle these and we address them rather fully, along with a number of open research questions, in Silverman [15].

2.3 Game Theory and the Cognitive Subsystem

The cognitive subsystem serves in our model as the point where the diverse emotions, stressors, memories, and other factors are all integrated into a decision for action (or inaction) to transition to a next state (or return to the same state) in the Markov decision process sense. In essence, at each node of the Markov chain (and at each tick of the simulator's clock) each agent must be able to process the following information: the state name (or ID); the allowable transitions and what action might cause those state transitions (a_{nm} in $A(iSTRESS)$); current intentions as provided in a task list or plan and the intentions of their prior actions; expectations of what other agents are going to do in this state based on recent history and other memories/beliefs $G(A, U, C)$; desires for actions based on the 11 pairs of emotional scales ($I_{\xi}(s_k)$ where $\xi = 1,2,2$); stress-based coping level Ω_i where $i = 1,5$); and a mood, μ , that we discuss below. Using all this information as stimuli, the agent must select a decision style, Φ , also defined below, and process the stimuli to produce a best response (BR) that maximizes expected, discounted rewards or utilities in the current iteration of the game. The cognitive subsystem is thus governed by the following equation:

$$BEST\ REPLY\ (BR_t) = \Phi_{\mu, iSTRESS, \Omega}\{U_{mn}(s_t, a_{mnt}), p_{mn}\},$$

subject to $a_{mnt} \in A(iSTRESS)$ [3.3]

Where,

$\Phi_{\mu, iSTRESS, \Omega}\{.\}$ = as defined below for the alternative values of μ , iSTRESS, and Ω

p_{mn} = perceived probability = $(1 - \Delta) e_m + \Delta_{m\tau} p_{m\tau}$

u_{mn} = $(1 - \delta)x(U$ from equation 3.1)

Δ = memory coefficient (discounting the past)

τ = number periods to look back

e_m = $\begin{cases} 0 & \text{action } m \text{ not} \\ 1.0 & \text{action } m \text{ is situationally relevant} \end{cases}$

δ = expectation coefficient (discounting the future)

$A(iSTRESS)$ = action set available after integrated stress appraisal (see Section 2.1)

We assume utilities for next states are released from the emotional activations. The previous section used the OCC model to help generate up to 11 pairs of emotions with intensities (I_{ξ}) for the current and/or next state of iterative play. Utility may be thought of as the simple summation of all positive and negative emotions for an action leading to a state. Since there will be 11 pairs of oppositely

valenced emotions in the OCC model, we normalize the sum as follows so that utility varies between -1 and +1:

$$U = \sum_{\xi} I_{\xi}(s_k) / 11 \quad [3.1]$$

While one can argue against the idea of aggregating individual emotions, this summation is consistent with the somatic marker theory. One learns a single impression or feeling about each state and about actions that might bring about or avoid those states. The utility term, in turn, is derived dynamically during each iteration from an emotional construal of the utility of each action strategy relative to that agent's importance-weighted concern ontology minus the cost of carrying out that strategy. We further introduce a modifier on the emotional construal function – the first is a discount factor, δ , that more heavily weights game achievement the closer the agent is to the end of that stage of the game. Thus an agent might be conservative and construe survival as more important early in the game, yet be willing to make more daring maneuvers near the end point: e.g., see Anderson, 2001.

It is useful to now turn to the discussion of the decision processing style function, Φ_{μ} , iSTRESS, Ω . There is a large literature on decision style functions (e.g., among many others see [19], [20], [15], [16] and [17]), and the discussion here is merely to indicate that there is a rich set of possibilities that one can explore within the framework proposed here.

For example, under perfect conditions, humans are presumed to be rational and behave according to Bayes Theorem and expected utility, yet as conditions degrade, they initially follow the dictums of subjective expected utility theory [19] and, eventually, of Recognition Primed Decisionmaking [17] or panic. Cognitive Continuum Theory [20] and Conflict Theory [16] provide compelling explanations of when each decision model is likely to prevail, and we adopt and adapt the latter for now.

2.4 Motor/Expressivity Subsystem

We complete the discussion of earlier Figure 1 by turning now to the motor/expressive subsystem. This module contains libraries of stored procedures that allow the agent to interact with the microworld and that allow it to display its motor and expressive outputs. Based on stimuli from all the other subsystems, the motor subsystem recalls, activates, and adjusts the relevant stored procedures so it can perform the actions intended to reach the (best reply) next state. In attempting to carry out the actions the motor system seeks to carry out best reply actions and perform up to the limits that the physiologic system imposes and by expressing the emotions that currently dominate. To support this effort, those procedures include functions that

allow them to portray alternative behaviors (e.g., fatigue leads to slower rate of movement across the screen). Also, the motor system serves as a stimuli to the other systems. For example crouching for a long period might cause fatigue, pain, emotive distress, and so on.

3. Case Study: Emergent Crowd Behaviors

We have attempted an initial, prototype implementation of our cognitive agent architecture to demonstrate how one might apply it to model the impact of alternative personas and motivations upon crowd behavior. This is not the final word on how to model crowd motivations and behavior, rather this is an attempt to illustrate the range and flexibility that the architecture supports.

The population of the city is initially meandering among random places along set paths. Upon becoming aware of the protest via a message broadcast to all agents, each individual decides whether or not to attend the protest, and if so in what capacity – either to observe or to participate. Figure 4 shows a small group at the outset of the protest, marching around in picket line formation in front of a security guard. While the choice of action here is often the same among various agents, the motivations for doing so can vary significantly. In fact, one agent is attending on a mission for the guerilla group, with the express purpose of causing a public disruption. This simple scenario requires one to model terrorists, defenders, civilians, crowd dynamics, population opinion evolution, and so on.




Figure 4 – Screen Shot of the Protest Scene Showing Observers on the Road, Picketers Holding Placards, and a Sole Security Agent Facing the Crowd.

To support viewing the internals of all these agents, on the left side of Figure 4 is a set of agent identifiers/pulldowns and window tabs. One of the pulldowns allows the user to select a terrorist, defender, or civilian (including up to five types of civilians such as unemployed male, employed male, female, etc) to inspect. For the selected agent, there are several tabbed windows also on the left side of Figure 1 – general, accessors, physiology, stress, emotion, and strategy – that allow one to inspect what that agent experiences, feels, and thinks about the microworld. From these various tabs one can thus piece together the agent’s beliefs, desires, and intentions of the moment.

Let us examine a portion of the scenario in more detail so one can better see how the diverse agents determine their motives, and carry out their actions. The instigator agent sent by the guerilla group to the protest has a portion of his Markov chain that deals with encountering security, taunting them, and precipitating violent reactions from them. Using sources such as [18], [21] and [22] we have derived a representative concern ontology as shown in Figure 5 that includes strong weightings on his goals for belonging (to his terrorist cell), esteem from taking action (they tend to be young males who are action-prone), and self-actualization due to reaching for ideals of freedom. Each of these is grounded in lower level goals that are positively aroused by taking action against the security forces.

Social psychologists have studied factors that contribute to aggressive crowd behavior: e.g., see McPhail [23] and Horowitz [24] among others. There is not uniform agreement on the particulars, but in general the common factors that tend to contribute include: presence of weapons, authoritarian government, lining up behind a barricade, drawing lines between “us” and “you”, dramatizing issues (e.g., in a speech) and making victims, large spatially concentrated crowds, and presence of television camera and crew. Also, rioters do not tend to be criminals, but they do tend to be the unemployed, single, young males without children. The bulk of attendees drawn to participate in the example protest are those very folks. They are therefore susceptible to crowd effects, and to a tipping event that sets them on a rampage including rioting and looting.

The tipping event occurs when the instigator is struck by the checkpoint guard (a neophyte in proper crowd dispersal tactics), an event that is observed by those near the front and communicated loudly. Coupled with the increasingly real possibility of becoming the target of violence themselves, the majority of agents strongly concerned with personal and family safety (employed males, females, etc) are prone to simply disperse. There

are insufficient security forces present for the relative size and density of this crowd, a fact of which the young males are aware. The erupting chaos provides a perfect diversion for, the young unemployed males to target nearby stores and loot them for material items.




Figure 5 – Overview of the Goal Portion of the Concern Ontology of an Agent Provocateur

4. Conclusions and Next Steps

To summarize, diverse communities are interested today in building realistic human-like behaviors into virtual personas. The animation and graphics approaches have lead to kinesthetically appealing and reactive agents. A few such investigators are now seeking to make them more behaviorally and cognitively realistic by reaching out to the artificial life, evolutionary computing and rational agent approaches. These approaches offer many benefits, but they need to be grounded in the behavioral literature if they are to be faithful to how humans actually behave and think. The behavioral literature, however, while vast, is ill-prepared for and cannot be directly encoded into models useful in agent architectures. This sets the stage for the goals and objectives of the current research.

A major challenge of this research, is the validity of the concern system ontologies and behavioral models we derive from the literature and try to integrate within our framework. As engineers, we are concerned with validity from several perspectives including the (1) data-

groundedness of the models and ontologies we extract from the literature, and (2) correspondence of behavioral emergence and collectives with actual dynamics observed in the real world. In terms of data-groundedness, we conducted an extended review of the behavioral literature [15] and found a great many physiological studies that seem to be legitimately grounded and that possess model parameter significance from a statistical sense. However, these tend to be restricted to the performance moderator functions that feed into the individual reservoirs or components of the physiological subsystem. As soon as one tries to integrate across moderators and synthesize the iSTRESS (or even effective fatigue), one rapidly departs from grounded theories and enters into the realm of informed opinion. The problem only grows worse for the emotion subsystem, and for the cognitive layer if one hopes to incorporate behavioral decision theory, crowd models, and the like. And the informed opinions one encounters in the behavioral literature are not consistent. One must choose one's HBMs and opinion leaders.

We have tried to provide one such collection of HBMs in this paper. This is not the penultimate integrative HBM, rather it is at present a humble structure. We have striven initially for satisfying a workability test. That is, we set out to attempt to learn what we could gain by having viable models integrated across all 4 subsystems and across factors within each subsystem. In that regard, our efforts to date are successful. We now have an integrated fabric stitching together the models of varying groundedness and of different opinion leaders. We can rather easily plug in a new opinion leader's model and play it within our framework to study its impact, its properties, and its strengths and weaknesses.

Finally, we offer no defense at present for our failure to have conducted correspondence tests. It's true that the agents may be observed to progress through various Ω levels (unconflicted adherence during daily routine, vigilant upon arriving at the protest, and panic during the looting) and the OCC model makes use of the reservoirs, crowd proximity, and an array of goals, preferences, and standards to generate emotions that appear consistent with what crowds probably feel. However, we simply haven't matured this research to the point yet where we are able to recreate specific historical crowd events from the real world, and to see how well our models are able to simulate actual emergent behavior. That is, however, a vital next step for benchmarking and tuning our models.

Despite validity concerns, there have been some lessons learned to date:

- The literature is helpful for improving the realism of behavior models – We have completed an in-depth survey of the literature and have found a number of models that

can be used as the basis of cognitive models for agent behavior. In fact the problem is less that there aren't any models, so much as the fact that there are too many and none of them are integrated. The bulk of the effort we undertook to date is to document those models, and to figure out how to integrate them into a common mathematical framework.

- There are benefits (and costs) of modeling stress-emotion-decision processing as an integrated topic – In attempting to create an integrated model, the benefits of this approach are that it is more realistic to try and deal with the interplay. Certainly these dimensions are connected in people, and the ability to address all of them in simulations opens up a large number of possibilities for improving agent behavior and for confronting trainees with more realistic scenes.

- Concern ontologies are vital but require ontological engineering– The approach we presented in this paper relies on a common mathematical framework (expected utility) to integrate many disparate models and theories so that agents can assess preferences and standards and determine next actions they find desirable subject to stress induced limitations and bias tendencies. However, to do this properly for any given simulation will also require extensive ontological engineering to flesh out the lower levels of the concern ontologies. Our current efforts are aimed at adding a set of tools for authoring, maintaining, and visualizing these ontologies.

- Emotion models are useful for utility and decision making not just for expressivity – A related contribution of this paper lies in the use of ontology-derived emotion to help derive utilities dynamically. In standard decision theoretic models there is no basis for agents to compute their own utility functions. Instead these are derived by subject matter experts and inserted directly into the agent's decision module. In the approach postulated here, the subject matter experts would interact at a stage earlier, at the stage of helping to define the concern ontologies so that the agents can derive their own utility functions, values, and tradeoffs. This approach frees experts from having to infer utilities, and it places the debate more squarely on open literature accounts of value sets and concern ontologies.

ACKNOWLEDGEMENT

The authors would like to acknowledge the financial support of The Pentagon, Defense Modeling and Simulation Office (DMSO), in general, and the help of Ruth Willis, Lt. Col. Eileen Bjorkman and Phil Barry, in particular. More recently, we thank John Tyler of Mitre

and Dexter Fletcher and Joe Toth of IDA for useful feedback and discussion.

REFERENCES

- [1] R.W. Pew and A.S. Mavor: Modeling Human and Organizational Behavior: Application to Military Simulation, National Academy Press 1998.
- [2] A. Sloman and B. Logan: "Building Cognitively Rich Agents Using the SIM_AGENT Toolkit" Communications of the ACM, Vol. 42(3), pp. 71-77, 1999.
- [3] J. Laird and M. Van Lent: "Human-Level AI's Killer Application, Interactive Computer Games" Artificial Intelligence Magazine, Vol. 22(2), pp. 15-25, Summer 2001.
- [4] J. Bates: "The Role of Emotion in Believable Agents" Communications of the ACM, Special Issue on Agents, July 1994.
- [5] C. Elliot: "The Affective Reasoner: A process model of emotions in a multi-agent system" Doctoral Dissertation, Northwestern University, Evanston, Illinois 1992.
- [6] B. Hayes-Roth: "Animate Characters" Autonomous Agents and Multi-Agent Systems, Vol. 1(2), pp. 195-230, 1998.
- [7] B.G. Silverman, J. Holmes, S. Kimmel, et al.: "Modeling Emotion and Behavior in Animated Personas to Facilitate Human Behavior Change: The Case of the HEART-SENSE Game" Health Care Management Science, Vol. 4, pp. 213-228, 2001.
- [8] A.R. Damasio: Descartes' Error - Emotion, Reason, and the Human Brain, Avon, New York 1994.
- [9] A. Ortony, G.L. Clore and A. Collins: The Cognitive Structure of Emotions, Cambridge University Press, Cambridge 1988.
- [10] I. Roseman, M.S. Spindel and P.E. Jose: "Appraisals of Emotion-Eliciting Events: Testing a Theory of Discrete Emotions" J. Personality and Social Psychology, Vol. 59(5), pp. 899-913, 1990.
- [11] R. Lazarus: Emotion and Adaptation, Oxford University Press, Oxford 1991.
- [12] J. Gratch: "Modeling the Interplay Between Emotion and Decision-Making" Proceedings of the 9th Conference On Computer Generated Forces and Behavioral Representation 2000.
- [13] Seif El-Nasr, T.R. Ioerger and J. Yen: "A Web of Emotions" J. Velasquez (Ed.), Emotion Based Agent Architectures Workshop of Agents '99 Conference, proceedings available at <http://www.ai.mit.edu/people/jvelas/ebaa99/ebaa.htm>.
- [14] L.M. Crumley and M.B. Sherman: "Review of Command and Control Models and Theory" DTIC No. AD-A230 105, September 1990.
- [15] B.G. Silverman: (2001). "More Realistic Human Behavior Models for Agents in Virtual Worlds: Emotion, Stress, and Value Ontologies" Technical Report, Philadelphia, PA, Univ. of Penn/ACASA Technical Report 2001.
- [16] I.L. Janis and L. Mann: Decision Making: A Psychological Analysis of Conflict, Choice and Commitment, The Free Press, New York 1977.
- [17] G.A. Klein, J. Orasanu, R. Calderwood and C.E. Zsombok: Decision Making in Action: Models and Methods, Ablex, Norwood, NJ 1993.
- [18] M. Johns, B.G. Silverman: "How Emotion and Personality Effect the Utility of Alternative Decisions: A Terrorist Target Selection Case Study" Proceedings of the 10th Conference on Computer Generated Forces and Behavioral Representation, SISO, May 2001.
- [19] W. Edwards (Ed.): Utility Theories: Measurements and Applications, Kluwer Academic Publishers, Boston 1992.
- [20] K.R. Hammond, R.M. Hamm, J. Grassia and T. Pearson: "Direct Comparison of Intuitive, Quasi-rational, and Analytical Cognition" (Report No. 248), Center for Research on Judgment and Policy, University of Colorado, Boulder 1983.
- [21] R. Weaver, B.G. Silverman et al.: "Modeling and Simulating Terrorist Decision-making" Proceedings of the 10th Conference on Computer Generated Forces and Behavioral Representation, SISO & IEEE, May 2001.
- [22] C.J.M. Drake: Terrorists' Target Selection, St. Martin's Press, New York 1998.
- [23] C. McPhail: The Myth of the Madding Crowd, De Gruyter, New York 1991.
- [24] D. Horowitz: The Deadly Ethnic Riot, UC Press, Berkeley 2001.