



4-1-2006

Human Behavior Models for Agents in Simulators and Games: Part I: Enabling Science with PMFserv

Barry G. Silverman
University of Pennsylvania, basil@seas.upenn.edu

Michael Johns
University of Pennsylvania

Jason Cornwell
University of Pennsylvania

Kevin O'Brien
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/ese_papers

 Part of the [Electrical and Computer Engineering Commons](#)

Recommended Citation

Barry G. Silverman, Michael Johns, Jason Cornwell, and Kevin O'Brien, "Human Behavior Models for Agents in Simulators and Games: Part I: Enabling Science with PMFserv", . April 2006.

Postprint version. Published in *Presence: Teleoperators and Virtual Environments*, Volume 15, Issue 2, April 2006, pages 139-162.

Publisher URL: <http://dx.doi.org/10.1162/pres.2006.15.2.139>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ese_papers/296
For more information, please contact repository@pobox.upenn.edu.

Human Behavior Models for Agents in Simulators and Games: Part I: Enabling Science with PMFserv

Abstract

This article focuses on challenges to improving the realism of socially intelligent agents and attempts to reflect the state of the art in human behavior modeling with particular attention to the impact of personality/cultural values and affect as well as biology/stress upon individual coping and group decision-making. The first section offers an assessment of the state of the practice and of the need to integrate valid human performance moderator functions (PMFs) from traditionally separated sub-fields of the behavioral literature. The second section pursues this goal by postulating a unifying architecture and principles for integrating existing PMF theories and models. It also illustrates a PMF testbed called PMFserv created for implementing and studying how PMFs may contribute to such an architecture. To date it interconnects versions of PMFs on physiology and stress (Janis-Mann, Gillis-Hursh, others); personality, cultural and emotive processes (Damasio, Cognitive Appraisal-OCC, value systems); perception (Gibsonian affordance); social processes (relations, identity, trust, nested intentionality); and cognition (affect- and stress-augmented decision theory, bounded rationality). The third section summarizes several usage case studies (asymmetric warfare, civil unrest, and political leaders) and concludes with lessons learned. Implementing and inter-operating this broad collection of PMFs helps to open the agenda for research on syntheses that can help the field reach a greater level of maturity. Part II presents a case study in using PMFserv for rapid scenario composability and realistic agent behavior.

Keywords

personality and emotion, social and cultural factors, physiology and stress, agent cognition, unified architecture

Disciplines

Electrical and Computer Engineering

Comments

Postprint version. Published in *Presence: Teleoperators and Virtual Environments*, Volume 15, Issue 2, April 2006, pages 139-162.

Publisher URL: <http://dx.doi.org/10.1162/pres.2006.15.2.139>

Human Behavior Models for Agents in Simulators and Games:

Part I – Enabling Science with PMFserv

Barry G. Silverman, Ph.D., Michael Johns, Jason Cornwell, Kevin O'Brien
*Ackoff Center for Advancement of Systems Approaches (ACASA), Department of
Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA
19104-6315, USA. e-mail: barryg@seas.upenn.edu*

ABSTRACT

This article focuses on challenges to improving the realism of socially intelligent agents and attempts to reflect the state of the art in human behavior modeling with particular attention to the impact of personality/cultural values and affect as well as biology/stress upon individual coping and group decision-making. The first section offers an assessment of the state of the practice and of the need to integrate valid human performance moderator functions (PMFs) from traditionally separated sub-fields of the behavioral literature. The second section pursues this goal by postulating a unifying architecture and principles for integrating existing PMF theories and models. It also illustrates a PMF testbed called PMFserv created for implementing and studying how PMFs may contribute to such an architecture. To date it interconnects versions of PMFs on physiology and stress (Janis-Mann, Gillis-Hursh, others); personality, cultural and emotive processes (Damasio, Cognitive Appraisal-OCC, value systems); perception (Gibsonian affordance); social processes (relations, identity, trust, nested intentionality); and cognition (affect- and stress-augmented decision theory, bounded rationality). The third section summarizes several usage case studies (asymmetric warfare, civil unrest, and political leaders) and concludes with lessons learned. Implementing and inter-operating this broad collection of PMFs helps to open the agenda for research on syntheses that can help the field reach a greater level of maturity. Part II presents a case study in using PMFserv for rapid scenario composability and realistic agent behavior.

Keywords: personality and emotion; social and cultural factors; physiology and stress; agent cognition; unified architecture

1. Introduction

The fields of virtual reality and microworld simulation have advanced significantly in the past decade. Today, computer generated personas or agents that populate these worlds and interact with human users are now deployed in many endeavors and avenues of investigation. A few of many example application areas are Hollywood animations for movies, cartoons, and advertising; immersive industrial and safety training simulations; distributed, interactive military war games and mission rehearsals; and personal assistant agents to reduce technologic complexity for the general public, among others.

A common challenge running throughout these applications is to increase the realism of the synthetic agents' behavior and coping abilities. This is not an idle fancy, but a serious

objective that directly affects the bottom line of commercial concerns, mission achievement in non-commercial organizations, and the safety and health of individuals who need to transfer skill sets from virtual to real worlds. Agent-oriented products that are more affective and offer a better cognitive fit tend to sell better, such as the successful games *Tamagotchi* or *Catz and Dogz*. This lesson applies to embedded agents as well as stand-alone products. People are known to anthropomorphize technologic items such as cars, slot machines, computers, ATM machines, etc. A strategy of beating the competition is beginning to emerge by including greater degrees of personality, human modes of interactivity (e.g., voice synthesis for car navigation systems), and emotive features in personas embedded ubiquitously (e.g., lip-synched and facially-accurate expressions) (e.g., see Reeves & Nass, 1996; Hayes-Roth, 1998). Similarly, in training, analytical, and education systems with military applications there is a growing realization that greater social subtlety and behavioral sensitivity in synthetic agents provides human trainees with both (1) more and better opportunities to explore alternative strategies and tactics, and (2) higher levels of skill attainment (e.g., see Sloman & Logan, 1999). These benefits are possible if the tactics, performance, and behavior of synthetic agents change in response to alterations in an array of behavioral variables. As a few examples, one would like agent behavior to realistically change as a function of their assigned native culture (vital for executing missions against forces from different countries), their level of fatigue and stress over time and in different situations, and/or their effectiveness following the loss of an opposing force's leader.

1.1 Challenges for Increasing the Realism of Human Behavior Models

There has been significant progress toward the development of improved human-like performance of synthetic agents. However, several serious problems continue to challenge researchers and developers.

Developers have insufficient behavioral knowledge. To date, models of culture, personality, affect and behavior that have been commercialized still tend to be shallow and unsatisfying. There is no deep model of human-agent interactivity. Synthetic agent forces are naïve and unable to act with the credibility or behavioral variety seen in human operators. Changing this is on the cutting edge of current research: e.g., Silverman et al. (2001, 2002a&b), Laird and vanLent (2001), among others.

Artificial life has focused on low level functionality. Character animators have created virtual life forms (e.g., fish, plants, talking heads, full body characters, and groups) that are physically realistic, geometrically accurate, and kinesthetically natural when moving about within their virtual settings (e.g., see Badler, Allbeck, Zhao, Bunn, 2002). There has even been significant development of architectures to produce animated characters that react appropriately to a small range of emotive and environmental stimuli such as fright and flight, flocking, and lip- and facial-movement-synching to utterances or stimuli: (e.g., Terzopoulos, 1999). However, these tend to be reactive systems that perform no deliberative or high-level decision making or cognitive processing such as has been conceived by the artificial intelligence community.

Artificial intelligence (AI) focuses on high level cognitive functionality. Researchers in the “rational agent” community have created a wide array of methods, often formal and grounded logics, to support agent reasoning (Bradshaw et al., 1999), inter-agent communications (Labrou, Finin, & Peng, 1999), and autonomous planning and learning (Tambe et al., 1995). These methods make it possible for unembodied agents to sense and respond to their virtual environments. However, extensive computing resources are necessary to support such abilities. It remains to be seen whether the necessary computing cycles can also be designed to support similar realistic abilities in relatively more resource-intensive embodied life characters. In addition, care must be taken when applying artificial intelligence to enhance the behavioral and coping fidelity of synthetic characters. It is easy to use these techniques to create capabilities that no real human being would possess. For example, a rational agent can perform its tasks without suffering the effects of fatigue, stress, heat, illness, bio/chemical exposure, or other factors that would likely affect the performance of a real human operator. Surprisingly, until quite recently this also was true of widely respected ‘cognitive models’ (e.g., SOAR; ACT-R, Anderson, 1990) whose AI-based designs tend to ignore much that is known about how cognition varies as a function of individual differences, situational change, and task attributes. Since the late 1990s, the lead author has been researching how PMF approaches can turn this tide. Recent evidence indicates the tide is turning and some researchers on frameworks such as SOAR and ACT-R, among others, are now seeking to add such phenomena a posteriori (Belavkin, 2001, Ritter, 2002, Chong, 2004).

Behavioral researchers tend to ignore integration and implementation. Worse are the "silos." There are well over one million pages of peer-reviewed, published studies on human behavior and performance as a function of demographics, personality differences, cognitive style, situational and emotive variables, task elements, group and organizational dynamics, and culture. This is a potentially rich resource for agent developers. Unfortunately, almost none of the existing literature addresses how to interpret and translate reported findings as principles and methods suitable for implementation or synthetic agent development (Silverman, 1991). Too often, factors described in the human performance literature are only roughly quantified. Informed judgment and/or additional testing is required to parameterize factors as dose-response curves or performance moderator functions (PMFs). It is time consuming and often beyond the abilities of non-psychologist agent builders to determine the validity and generalizability of findings reported in behavioral scientific studies. Worse still, these literatures are highly specialized with deep "silos" separating the many topics one needs to synthesize for agent behavior. Any progress on such a synthesis is hampered by lack of first principles from the literature, and any incentive for behavioral researchers to benefit from computational integrations. Only small pockets of behavioral researchers currently straddle these concerns (e.g., Ness et al. 2005, Payr & Trappl, 2004). PMFserv is meant to help its users eliminate such barriers and to open possibilities for syntheses.

There is a dearth of interchange standards. There are few interchange standards in the AI and game-maker communities. One cannot readily substitute characters or agents from one simulation or game to run in the world of another: e.g., see Bjorkman, Barry, & Tyler (2001). Existing systems cannot easily be modified or extended to incorporate

interesting new factors, findings, or code. Interchange standards are needed to exploit the rich diversity of achievements from various related fields of development. At a basic level, such standards would cover API specifications for plug and play modules. Far more complex standards could be developed to support the exchange of agent knowledge and ontologies, behavior models, and the means to apply them within a diverse variety of systems.

There is a need to validate “useful” models of human behavior. According to folkism, “all models are broken, some are useful.” No model will ever capture all the nuances of human emotion, the full range of stress effects, or how these factors affect judgment and decision making. However, to the extent that a model provides a valid representation of human behavior, it will likely be useful to those who wish to simulate that behavior. Unfortunately, the problem of model validity has no simple solution. It is difficult to run existing simulations against real past events. Many first principle models from the behavioral science literature have been derived within a particular setting, whereas simulation developers may wish to deploy those models in different contexts. Likewise, there are validity issues raised by AI models of agent reasoning that are able to simulate human task performance in the absence of human-like reasoning. Is it even possible to validate the integration of multiple factors (e.g., stress and fatigue) when the research literature is largely limited to the study of independent rather than interactive effects of these factors? What methods of validation should be used, and for what purposes are models to be considered trustworthy? These are just a few of the many issues and questions that must be resolved in order to identify whether and how an agent-based model of human behavior should be implemented.

1.2 Human Behavior Modeling for Virtual Agents

To overcome the obstacles identified above, it would be useful to remove cross-community barriers and create a set of sharable resources for the modeling and simulation community. This goal is considered here, beginning with a discussion in this section of what can be learned from the behavioral science literature and then moving on in Section 2 to the question of how those assets can be integrated with existing, ‘reusable’ models of human performance and cognition. Part II of this paper further explores the reuse question looking into how to add such models to legacy software and simulators.

In terms of what can be learned, there is a voluminous literature, easily numbering in the 10,000s of studies, on the subject of human performance under stress – or “performance moderator functions” (PMFs). The author has abstracted about 500 PMFs from this literature and has begun to explore a common mathematical framework for representing them and how to use them in agent behavior simulations (Silverman, 1999; Silverman, Johns, Shin, & Weaver, 2002; Bharathy, Chung, Silverman, & Cornwell, 2002). One of the earliest studies in this field generated the now classic Yerkes-Dodson “inverted u” curve, which illustrates that as a stimulus or moderator is increased, performance is initially poor, then improves, and then falls off again after passing a threshold (Yerkes-Dodson, 1908). Thus, performance can be better in slightly chaotic, moderately time-pressured settings than in settings absent of stress, though this varies with individual differences and may not apply to all moderators.

In particular, we are interested in *emergent macro-behavior due to micro-decisions of bounded-rational agents* and with developing a framework that permits one to examine the *impacts of biology, stress, personality, culture, emotion, social relations, and decisionmaking* upon human coping behavior. 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. Or to study why one leader attracts followers and another is shunned.

2. Integrating the Many PMFs into a Unified Behavior Architecture

The PMF findings discussed above suggest that there is a large, often untapped resource to assist those who create human performance simulations. Much can be gained the more that simulation developers work to embrace such collections and to implement and test their results. This, in turn, should and could create a forcing function back upon behavioral scientists. Behavioral science has produced some compelling models and PMFs of individual factors, but they still need to do a better job of studying and evaluating integrative frameworks.

This broad study of so many PMFs lead the lead author of this paper in 1999 to postulate a systems synthesis in the form of a set of principles and a unifying architecture for agent coping behavior that have guided the evolution of this research and that are now repeated here:

(1) Inter-relations Between the Parts -- In the system literature by definition, a synthesis must focus on the inter-relations between the parts so as to foster understanding the purposeful behavior of the whole. In this case, the "parts" that need to be synthesized should be codifiable from the major branches of the literature such as shown by the large blocks of Figure 1 (we postpone discussion of many of Figure 1's details). It is widely accepted today that coping behavior (the synthesis we are interested in) is somehow influenced by biology/stress; personality and cultural values and emotion; social relations; and rational cognition; as well as perceptual and memory functions. Yet it is difficult to locate sources that address such a wholism and how these big parts play together.

(2) Subsystems are Systems as Well -- Here again, one is interested in the relationships between the parts. PMFs from diverse sources should be implemented and inter-operated to study modular sub-systems. Thus, the major boxes of Figure 1 themselves should synthesize diverse PMFs. Indeed, Figure 2 shows a number of lower-level PMFs that have been linked together in our architecture and the linkages between them – again we discuss the details in the next section. What is worth noting here is that many of the PMFs from the literature do provide useful pieces of the puzzle, but in order to connect them we had to add new PMFs (rounded edge boxes of Figure 2). In other cases we could only implement loose interpretations (quadrilaterals of Figure 2). In both Principle 1 and 2, the systems researcher is interested in accuracy, not precision since precision is a distraction from better understanding relationships between parts. This means the diverse PMFs of Figure 2 often are implemented as first order approximations (e.g., linear systems), and future implementations might improve on that.

(3) Study Best of Breed PMFs -- The unifying architecture in Figure 1 may seem high level; however, that is necessitated by the state-of-the-art today. The goal is to study best of breed PMFs even if they only implement portions of subsystems (sub-modules). The idea is to study how to get them to inter-operate so that modules and larger wholisms may be assembled, studied, and validated or invalidated. 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. This means that any computer implementation must support plugin/plugout/override capabilities, and that specific PMFs as illustrated in Figure 2 (octagonal shapes) should be testable and validatable against field data such as what they were originally derived from. Our research goal is to avoid locking in on PMFs and, in so doing, thereby creating a monolith. Instead, every PMF explored in this research should be readily replaceable. As an example, the connecting PMFs that we synthesized are workable defaults that we expect our users will research and improve on as time goes on.

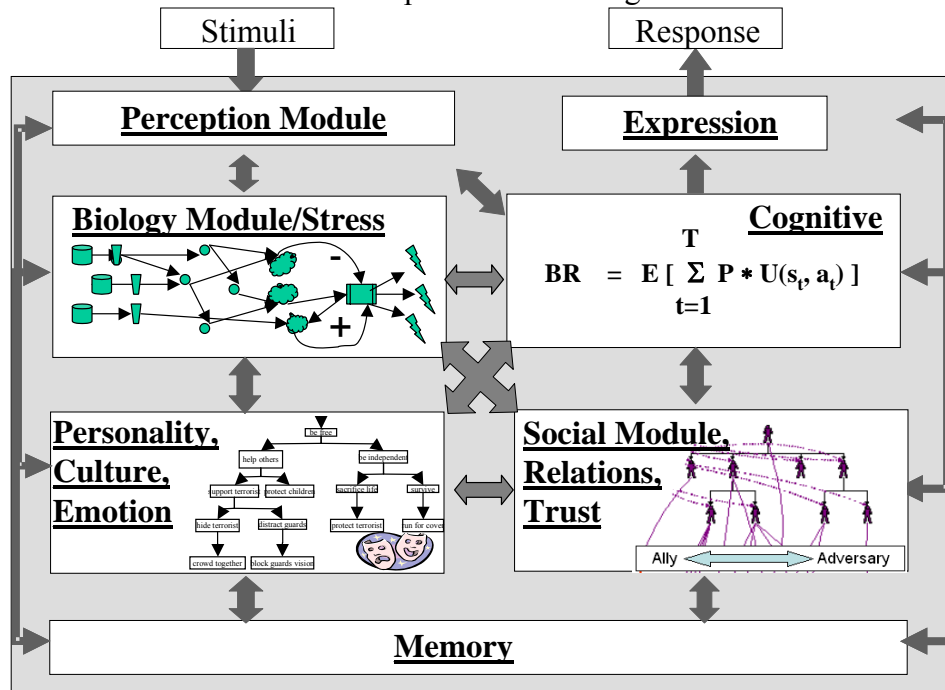


Figure 1. Unified Architecture of Behavior

(4) Agent Archetypes and Individual Differences -- When we synthesize a set of PMFs for each module, one should be able to calibrate and tune the parameters to recreate the coping patterns of archetypical classes of individuals. In this fashion, autonomous agents are capable of relatively realistic coping patterns. For example, in crowd scenes, one might like to have agent archetypes that characterize the mean behavior patterns of orderly protestors, provocateurs, looting hooligans, and so on. For each archetype, what's interesting is not strictly the mean behavior pattern, but what emerges from the collective. To understand that, one expects to instantiate many

instances of each archetype where each agent instance is a perturbation of the parameters of the set of PMFs whose mean values codify the archetypical class of agent they are drawn from. This means that any computerization of PMFs should support stochastic experimentation of behavior possibilities. It also means that individual differences, even within instances of an archetype, will be explicitly accounted for.

(5) Find the Synergy – The unifying behavior architecture described here has an implementation that we refer to as PMFserv as Figure 2 portrays (Silverman et al., 2001; Silverman, Johns, O'Brien, Weaver, and Cornwell, 2002; Silverman, Johns, Weaver, O'Brien, & Silverman, 2002). Systems should be more than the sum of their parts. In developing the synthesis, our goal was that it would support a wide range of agent behavior studies and applications.

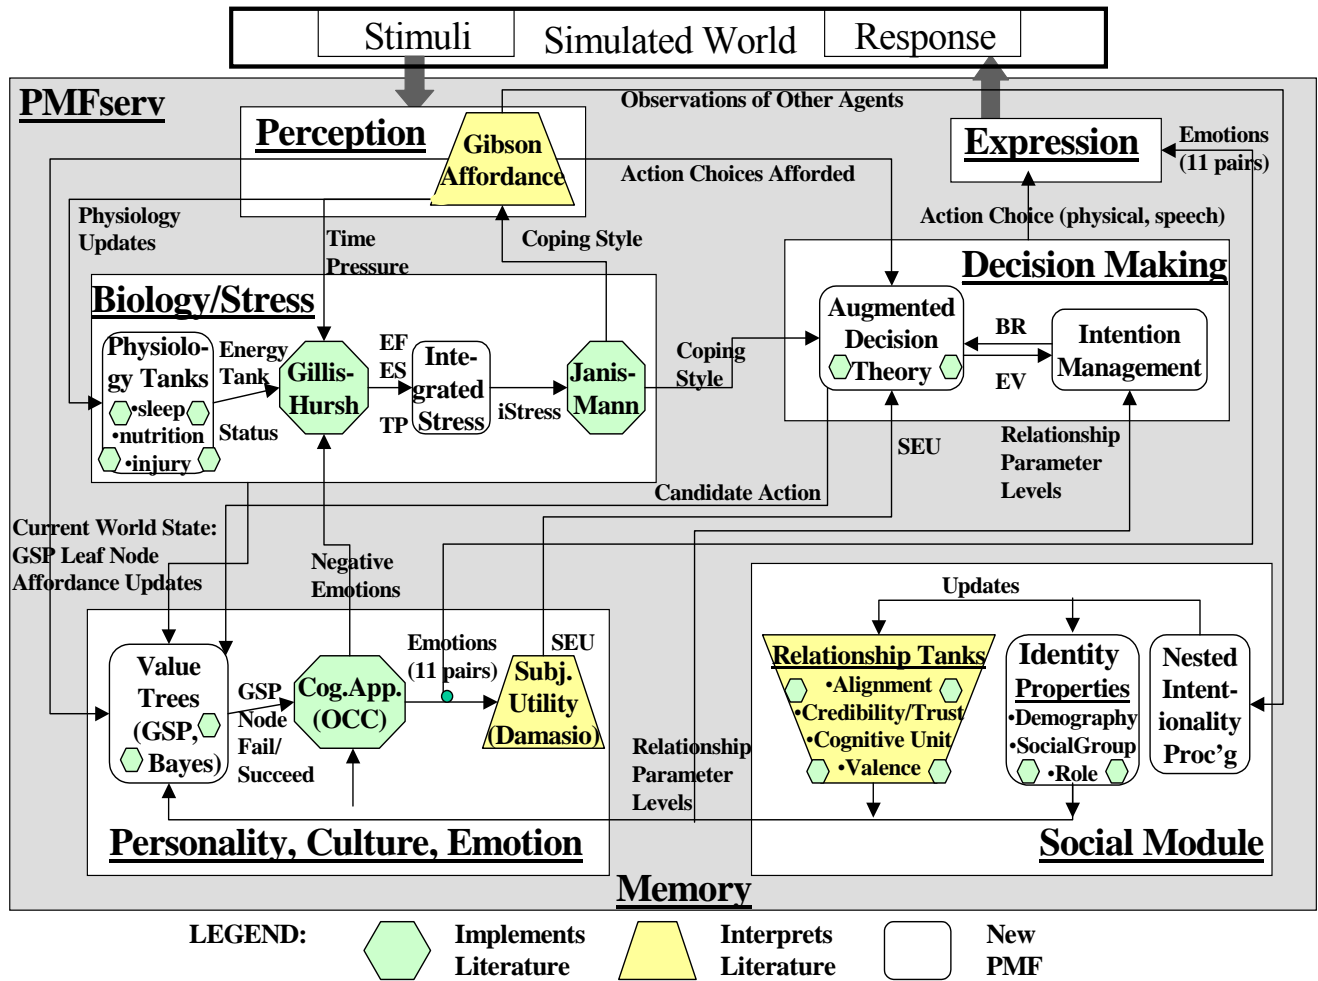


Figure 2. The Current PMFserv Implementation of the Unified Architecture

The next few subsections turn our attention to the modules and details of Figure 2, so it is worth concluding this section with a few words on Figure 1 in total. At the time we started this research in the late 1990s, we were unaware of any implementations straddling the modules of Figure 1, and to this date there still are not any available to the

practice. However, there is a growing interest in this field on behalf of sponsors and of researchers, and one can see many nascent attempts, some of which will be mentioned in what follows along with earlier research relevant to each module of Figures 1 and 2.

2.1 Biology Module: Physiology and Stress PMFs

Following the guidance of Principle 1, one must ask how does biology and stress influence coping behaviors and the other modules of our architecture. If it doesn't contribute, we might be tempted to omit biology. However, even entertainment game characters illustrate the importance of this subsystem and provide "damage reservoirs" and "energy power ups" as crude surrogates. Likewise, many military simulators include crude injury and fatigue PMFs. But how do biology and physiology influence overall behavior, perception, and decision functioning? Our search of the literature in the late 1990s and consults with domain experts revealed a PMF that directly addressed this question, the Janis-Mann coping styles model.


Janis & Mann (1977) provide what is probably the most widely cited methodology of decision strategies for coping under stress, time pressure, and risk. Their methodology has extensive empirical backing and has been replicated and validated by other researchers as well: e.g., see Wohl [1981] among others. Thus it also satisfies our 3rd principle (best of breed PMF). However, despite its stature, we are unaware of any computer agent implementations to use as a plug-in for our architecture, and so we have implemented our own version. The reader can see our interpretation of this methodology as the steps of the inverted U-curve of Figure 3. The methodology 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 3. These coping strategies provide the connection between biology (i.e., stress) and the perception, personality, and decision modules of our overall architecture.

In particular, 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) are regarded by Janis & Mann (1977) 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. (1993) who show that experts work effectively in "near panic", vigilant, and other modes where they immediately recognize a best or near best alternative without vigilant scanning of other alternatives.

The goal of a computerized implementation of this PMF is for overall integrated stress (what we label iSTRESS) to emerge dynamically as events affect the agent's biology. Unfortunately, Janis & Mann (1977) do not provide either (1) precise threshold values that indicate when decision makers trigger a change in coping style (Ω_i), 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 and the Gillis-Hursch (see below) factors into integrated stress as the 2nd from the right PMF box within the Biology Module of Figure 2 indicates. For example, our PMF for Integrated Stress had to have rules that 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. As per Principles 2 and 3, these heuristics are extensions to the original Janis-Mann methodology necessitated by implementation. They are worthy of empirical study on their own, and we have segregated them in our PMFserv implementations so that other investigators may alter or override them as their research warrants. Indeed, a doctoral student is currently exploring these heuristics at present.

a. Theory



b. Implementation




Figure 3. The Classic Performance Moderator Function is an Inverted-U
 (a) Janis-Mann Coping Styles PMF from the Literature
 (b) Janis-Mann PMF as Implemented Visually in PMFserv

As stated above, a major gap in the literature is that very few studies have been conducted to determine how multiple factors combine to produce overall or integrated stress in an individual (e.g., Hammond, 2000). One approach is that of Hendy and Farrell (1997), who adopt an information processing theory and model that focuses on equipment and screen operators and includes factors such as task workload (bits to process) and work rate (bits/sec). They offer an information processing (bit throughput, error rate, decision time) account that attempts to explain the effects of time pressure, task difficulty, mental capacity, fatigue, motivation, anxiety, and the like. However, they offer little data to support their model.

Hursh & McNally (1993) reviewed 1,300 studies to develop a model of decision making in battle that focuses solely on effectiveness under stress. Gillis and Hursh (1999) later

extended this model to account for what they claimed were the prime determinants of (stressed) performance: effective fatigue, event stress, and time pressure. They implemented this approach into a software system that we tried to get ahold of and integrate within PMFserv, however, the Army no longer supports that code. PMFserv thus had to have its own Gillis-Hursch implementation (see Biology Module of Figure 2). In particular, following Gillis and Hirsh, we obtain: (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 can be sent to be emotionally filtered since for example a stoic will construe the same facts differently than a high neurotic (The next section describes the emotional filtering). The quantitative factors that go into these modifiers are then summarized via the following where $f\{\cdot\}$ is our PMF mentioned above and which is currently a linear additivity model since we are interested in accuracy and first order approximations (Principle 2):

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

The interesting thing about this breakout is that it now opens the door to inter-operability of numerous other PMFs that are well researched in their own right. For example as shown in earlier Figure 2, the Event Stress (ES) is provided by our emotion module by counting any events that the agent's value system and personality profile causes negative emotion to be activated (Section 2.2). Time Pressure (TP) is computed in some well-regarded cognitive architectures (e.g., COGNET and iGen in Zachary et al. (2001) and we have worked with them to study best approaches to inter-operate). However, for now TP is afforded to agents by markups on the simulated world.

Finally, Effective Fatigue (EF) is the place where we integrate the remainder of our biology/physiology module as depicted in Figure 4. Specifically, to support rapid implementation, calibration, and study of a wide array of physiological PMFs, we have adopted a visual programming metaphor based on reservoirs with tanks and valves. The biology tanks editor shown in Figure 4 includes a pulldown menu for adding new reservoir tanks, valves, and gates. Through a set of popup menus, one can also set all the parameters of each of these objects, create their dependencies, and spawn GUI controls for them (e.g., as shown on right of Figure 4). Through another pulldown one can start and stop the module. By doing this and by altering the sliders on the right panel, one can test the PMFs alone and in concert, thereby tuning the model in real time. In the implementation actually depicted in Figure 4, there is a stomach tank that computes nutritional input (kcal), a sleep, injury, and temperature tank. Each of these is attached to the wastage valve of the energy tank. That tank is depleted by these interactions as well as by exertion. All these (and other tanks the user may wish to add) are calibrated and tuned to field data for the archetypes being modeled. Drops in the energy tank, as indicated in earlier Figure 2, are monitored by our implementation of the Gillis-Hursch PMF since that leads to increases in Effective Fatigue (EF).

One might pause and ask at this point, is it really important to support the integration and interoperability of so vast a potential array of PMFs? Once again, if our goal is to facilitate expanding the science (and it is), then we need a flexible implementation that allows us to rapidly integrate, test, and update (or replace/override) numerous PMFs, whatever is relevant to the scenario to be studied. For example, if one is interested in the impact of non-lethal weapons on crowds, very different PMFs should be integrated here than if one is interested in trauma and injury due to warfare. Likewise, if one had a better idea for EF, it is easy to set the Gillis-Hirsch PMF to monitor some other external model instead of our energy tank status level. Some users have used this physiology module to study the implementation of various low-level PMFs, how to calibrate the tanks and valves, and how all that inter-operates with the three summary stressors and with integrated stress and the emergence of coping style: e.g., see (a) Bharathy (2002) for a study of sleep, exercise, forced marches, and nutrition and (b) Bharathy et al (2003) for a study of various injury metrics and types of weaponry and trauma effects.




Figure 4. Physiology module uses a hydraulic metaphor for visual editing (and operation) of PMFs as systems of reservoirs, gates, and valves

2.2 Value Trees and Emotion Modeling: Personality and Culture

This section brings us to the questions of: (1) what roles do values (personal, cultural) and emotions play in the overall behavior architecture (earlier Figure 1) and (2) how might we implement and integrate existing PMF models to reasonably recreate actual human behaviors? In the past decade, research has yielded an explosion of literature on the connection between emotional and cognitive processes. Most notably, as suggested in earlier Figure 2, we have been inspired by the connection that Damasio (1994) draws and how emotion guides practical decision making. His "somatic marker hypothesis," though not universally accepted, suggests that body states associated with emotions precipitate somatic markers (gut feelings) that serve to guide decision making.

From our architectural view, this hypothesis implies (1) continual two-way interaction of the emotions with human biology, perceptual, and decision making modules; and (2) the need for an explicit representation of the motivations against which events are appraised. Our architecture and our emotion module are an attempt to implement these concerns, though since Damasio's work is not intended for computational implementation and thus stops short of many of the details required to do so, we view his theory primarily as an inspiration. For example, Figure 2 shows that our affect or emotion module is driven by perception and biology, and in turn, produces the Event Stress (ES) used in the biology which alters focus and perception. Likewise, our emotion module is implemented with a set of value trees representing agent motivation and, with help from the social relations module, events activate nodes on these trees to produce emotions and an overall gut feeling that we treat as subjective expected utility (SEU) to guide the decision module.

Before we explain how all this works, others have implemented emotion models as well. For example, Bates et al (1992) and Elliott (1992) each offer a robust implementation, some elements of which we duplicate here. However, their focus is on agents that provide socially interesting reactions in interactive dramas and fictions. Their agents draw no connections between either biological or decisionmaking functions (and have limited perceptual abilities), and as such these systems are not directly useful to our purposes. Likewise, there are a few agent researchers examining the interplay of emotions and decisionmaking: e.g., see Gmytrasiewicz and Lisetti (2000). Gmytrasiewicz's focus is on prescriptive algorithms and formal rigor, and he has published on the computational complexity his approach imposes. Our work contrasts with such efforts since we focus on descriptive research and scalable human behavior modeling.

Our goal for this module is to explore how well alternative PMFs found in the literature support the requirements of the Damasio-inspired model, although that model itself may be amended in future implementations. In order to support research on alternative emotion theories, this subsystem must include an easily alterable set of activation/decay equations and parameters for a variable number of emotions. Further, since appraisals are based on models of motivation (value trees), this module must include a motivation representation, editor, and processor. Simply by authoring alternative value trees, one should be able to capture the behavior of alternative "types" of people and organizations and predict how differently they might assess the same events, actions, and artifacts in the world around them. This requires that the emotion module be able to derive the gut feel (subjective expected utility and payoff) that the decision module will need to make choices.

We start at the bottom and 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 the importance of one's values or set of value trees (v) and whether those values succeed or fail for the state in question. We express this as follows and make reference to parameters on earlier Figure 2, though not all of these were displayed on that chart due to space limitations:

$$I_{\xi}(s_k) = \sum_{j \in J_{\xi}} \sum_{c \in C_{ijkl}} [W_{ijl}(v) * f_1(r_j) * f_2(O,N)] \quad [2.0]$$

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 ξ . 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 . The relationship parameters (alignment, trust, cognitive unit, etc.) of the next section define these.

$W_{ij}(v_{ijkl})$ = Weighted importance of the values of agent j that succeed and fail in one's i th value tree (see trees and Bayesian weights discussion below for further explanation).

v_{ijkl} = A list of paths through the i th value tree of agent j triggered by affordance ℓ (0=success or 1=failure) by world state k . Here a path refers to the vertical set of nodes and branches that are activated from a given affordance (e.g., from perceiving the current state of the world, or from contemplating a new world state that a given action might precipitate)

$f_1(r_{jk})$ = 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 world state k . (Again these relationship parameters and their possible settings are clarified in Section 2.4).

$f_2(O,N)$ = A function that captures temporal factors of the world state and how to keep apart emotions that might get activated in the current time frame that are about the past vs. the present vs. the future. For example, as time proceeds emotions about the future may need to be recognized as and converted to emotions about the present (e.g., hope being dashed might be more intense than just disappointment in this time tick alone).

This expression captures the major dimensions of concern in any emotional appraisal – values, relationships, and temporal aspects. For the sake of simplicity, we assume linear additivity of multiple activations of the same emotion from the $i=1, I$ different sets of values that the world 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. Specifically, we have examined OCC: Ortony, Clore, and Collins (1988), Roseman et al (1990), and Lazarus (1991). Each of these 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, and we have abstracted them and discuss this in Silverman, B.G., Johns, M., Shin, H., Weaver, R. (2002). At present and as did Bates (1992) and Elliott (1992), we have decided to pursue the OCC model (Ortony et al., 1988) 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 though we send negative emotions to the biology module (ES counting) and an expression module might not treat them as canceling, particularly for agent embodiment.

We have built a number of agents with value trees and find that it is best to represent an agent's values in three distinct trees illustrated in the leftmost box of the Emotion Module of Figure 2. An agent's long term desired states for the world are placed as nodes on a Preference Tree, the Standards Tree captures the actions they believe should or should not be followed in trying to reach those preferred states (e.g., religious commandments, military doctrine, civil codes, etc.), and any short term needs are placed as nodes on a Goal Tree. In regards to the latter tree, we often place physiologic and stressor concerns on this tree which causes the agent to have emotions about its biological functions if the reservoirs in Section 2.1 pass threshold settings. We refer to these three trees as the GSP trees. The branches of these trees are Bayesian weighted to capture relative importance and to help represent behavioral patterns in prior observations of the human archetypes that these agents represent.

The OCC model (middle Emotion Module box in Figure 2) assumes a social agent has three types of value sets that are almost isomorphic to our GSP trees, although they provide little guidance on how they believe these should be implemented: goals for action, standards that people should follow, and preferences for objects. This latter set differs from our own definition, though not seriously enough to require a major change to OCC. Our emotion model is described in great detail in numerous other papers (available at <http://www.seas.upenn.edu/~barryg>) so our discussion here will be brief. The general idea is that an agent possesses the three GSP Trees.

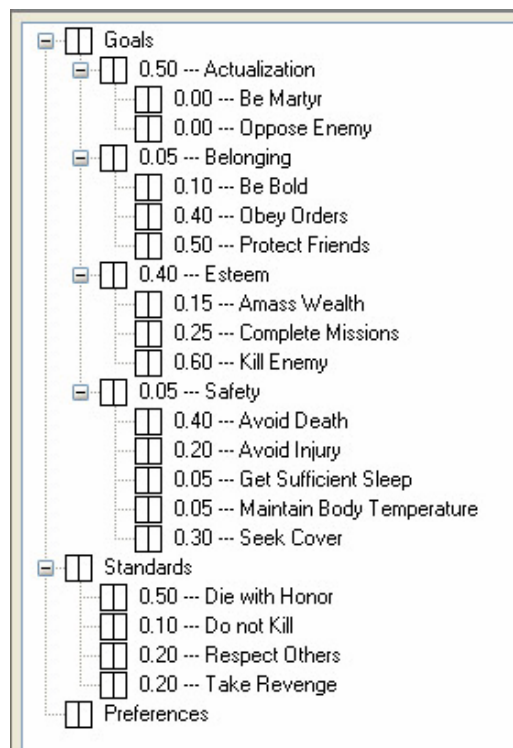


Figure 5. Illustrative shred of GSP trees and values of a sample terrorist

An action can be represented by a series of successes and failures on the sub-nodes of these three trees. Each sub-goal is given a weight that describes how much it contributes to its parent node. When contemplating a next action to take, the agent calculates the emotions it expects to derive from every action available to it, as constrained by perception and coping style. We assume that utilities for next actions, a_k , are derived from the emotional activations. Silverman, Johns, Weaver et al. (2002) describe the set of mathematical equations for the use of the OCC model to help generate up to 11 pairs of emotions with intensities (I_ξ) for a given action. These expressions capture the major dimensions of concern in any emotional model – values, relationships, and temporal aspects. Utility may be thought of as the overall ‘feeling’ and as such, are 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, our Damasio box on the right of the Emotion Module in Figure 2 normalizes the sum as follows so that subjective expected utility varies between -1 and $+1$:

$$SEU = \sum_{\xi} I_{\xi}(a_k) / 11 \quad [3.0]$$

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 appraisal of the utility of each afforded action strategy relative to that agent’s importance-weighted values (GSP trees) minus the cost of carrying out that strategy.

2.3) Perception and Ecological Psychology: The Affordance Approach

For all the PMFs described thus far to work, the agent must use its *perception module*, as constrained by coping mode and emotive needs, to see what’s going on in the world. Perception should be focused based on an agent’s physiology, coping style, prior emotional needs, and any memory elements that might have been created before the current cycle. For example, if the agent’s coping mode is Panic or Unconflicted Adherence, it will not notice anything new in the world. Otherwise, PMFserv applies affordance theory (Gibson, 1979, Toth, 1995) such that each object in the simulated world executes perception rules to determine how it should be perceived by the agent and generates a list of the corresponding actions (a_k) and affordances it can offer that agent (e.g., a rock indicates it can be thrown, which will afford success in hurting an opponent and will consume x units of energy). When these actions are selected as shown in Figure 2, the affordances provide reservoir replenishment or drawdown impacts and GSP tree multipliers for degree of leaf node success or failure. In this fashion, PMFserv agents implement situated ecological psychology (Gibson, 1979).

In earlier revisions of our architecture, each agent stored its own knowledge about the world internally. We ran into problems, however, when we tried to write rapid-scenario

generation tools. We wanted to be able to draw agents and object from a large library of presets that we could drag into a scene and create scenarios with no additional programming. To do so, all of the perceptual data and state information for each object in the scene would need to be generated along with a means for evaluating each option that tied into the individual agents' unique standards, goals, and preferences trees. We could not find an elegant solution that would generate this data on the fly. Entering it all by hand was not a possibility either, given our goal of rapid composability.

To allow for rapid scenario generation and modification we have removed knowledge about the environment from the agents themselves. Instead of imbuing the agents with this knowledge, a scenario designer can build an ontology of objects in the environment that describes how each of those objects can be perceived and used. The popular video game *The Sims* takes a similar approach. Our agents themselves know nothing a priori about the environment or the actions that they can take within that environment. Instead, each object in the environment contains multiple representations (perceptual types) of itself that include its affordances - the actions that can be taken on it, the results of those actions to the object being viewed, and the impact that those actions will provide to the agent in terms of its various biology, stressor, value tree or social tanks. Each object contains perception rules that determine which perceptual type is active for the agent currently perceiving the target object.

The affordance approach was taken to satisfy engineering constraints rather than theoretical concerns, but there is no shortage of theoretical justification for making such a move. In his landmark text *The Perception of the Visual World* (Gibson, 1950), Gibson argued that people perceive the objects in their environment in terms of their affordances, or the opportunities for action that they provide. For example, when we look at a doorknob, the actions that it offers to us – opening a door, for example, or perhaps hanging up a coat – are explicit in our perception. In Gibson's theory, the perception of an object arises from the interaction of the perceptual qualities of that object and the perceiver's past experiences. There is no top-down processing at work. In our implementation of this theory, the perceived object interacts with the perceiver and generates a list of actions along with their anticipated emotional effects (GSP leaf node activations) independently of any cognition on the part of the agent. This is the arrow emanating from the right hand side of the Perception Module of Figure 2.

For example, in the Mogadishu scenario we are currently exploring we have an object that represents an American helicopter (see Figure 6, below). The helicopter has multiple perceptual types, each of which has a different set of actions it affords. One Somali agent might perceive the helicopter as a "Frightening Helicopter" that can be investigated, attacked, or fled from. An agent in the militia might perceive it as an "Enemy Helicopter" that can be attacked, fled from, observed, or taken cover from. A third agent already in combat with the helicopter might view it as a "Combatant Enemy Helicopter" that can be attacked, or taken cover from as well, but with different emotional outcomes. Agents themselves are wrapped in a perception model so that, for example, a suicide bomber in a crowd might be perceived as a normal civilian until he does something to

reveal himself, at which point those around him may shift their perception to see him as a suicide bomber.

To accomplish perceptual shifts of this sort, each object contains a set of perception rules that determine which perceptual type is active for any given agent. These rules take into account the coping style of the perceiving agent, the way the object has been perceived in earlier simulator time ticks, tests of facts about the world (“has agent x pulled out a gun?” etc.), whether or not the object is in the agent’s line of sight, and any other exposed variable within our system that a scenario developer might want to use to determine how an object is perceived.

The affordances themselves are described in terms of their effects on a generic agent’s biological, stress, and social relationship parameters (tanks) as well as upon its value trees (leaf nodes of the goal, standard, and preference trees). Effectively, they convey elements needed to compute the emotional utility that the agent expects to receive from each action. The right-hand column of Figure 6 shows a few examples of the GSP tree affordances offered by a helicopter crash site, as seen by a Somalian civilian:

Helicopter					
Perceptual Type	Action	OCC Results	T	A	
Interesting Object	Approach	G.SatisfyCuriosity	S	1	
		G.StayAlive	F	0.1	
UnguardedCrashSite	Approach	G.SatisfyCuriosity	S	1	
		G.AmassW ealth	S	1	
	Loot	G.ProtectFamily	S	1	
		S.DoNotSteal	F	0.5	
		S.RespectOthersProperty	F	0.5	
		Observe	G.SatisfyCuriosity	S	1
GuardedCrashSite	Approach	G.AmassW ealth	S	0.2	
		G.SatisfyCuriosity	S	1	
	Loot	G.StayAlive	F	0.4	
		P.ForeignSoliders	S	1	
		G.AmassW ealth	S	1	
		G.ProtectFamily	S	1	
LethalCrashSite	Approach	G.StayAlive	F	0.8	
		S.DoNotSteal	F	0.5	
	Loot	S.RespectOthersProperty	F	0.5	
		P.ForeignSoliders	S	1	
		Observe	G.SatisfyCuriosity	S	1
		G.AmassW ealth	S	0.2	
LethalCrashSite	Approach	G.SatisfyCuriosity	S	1	
		G.StayAlive	F	1	
	Loot	P.ForeignSoliders	S	1	
		G.AmassW ealth	S	1	
		G.ProtectFamily	S	1	
		G.StayAlive	F	1	
LethalCrashSite	Approach	S.DoNotSteal	F	0.5	
		S.RespectOthersProperty	F	0.5	
	Observe	P.ForeignSoliders	S	1	
		G.SatisfyCuriosity	S	1	
		G.AmassW ealth	S	0.2	
		G.StayAlive	F	0.4	

Figure 6. Simplified Affordance Structure of A Helicopter Object in the Simulated World

A full listing of the perceptual types and affordances offered by an object in our system is far longer, and also includes descriptions of physiological effects, social relation impacts, and perception rules governing which perceptual type is active for any given agent. More details exist in Cornwell et al., 2003; Silverman, et al., 2003.

Each object in the environment is marked up with as inclusive a set of perceptions and actions as is appropriate for the scenario. Once these objects are marked up, any agent can be dropped into the scenario and automatically have an appropriate understanding of how it can manipulate the objects and agents around it. Likewise, a new object or new agent can be developed independently and instantly “understand” how to interact with the other agents and objects already deployed in the simulation.

In the agent’s decision cycle, each object in the environment evaluates its perception rules against the agent to determine the set of available perceived objects that the agent can perceive and manipulate. We return to discuss this further in the Cognitive Module.

2.4) Social Module: Relations, Trust, Theory of Mind

The social module is just as vital as each of the others in our unifying architecture. For example, in order for human values and emotions to work, there also must be a way to track social relationships and who is aligned with and/or against whom. When something good happens to a family or ingroup member (defined here as close allies), we tend to have positive emotions (though possibly some envy too), and the opposite if it happens to our adversary. If the valence is intense enough, this will noticeably impact our Event Stress (ES) and begin to affect us biologically to the better or worse, respectively. Likewise, decisionmaking and course of action selection is affected by social alignment and trust. Certain classes of actions are rarely contemplated against ingroup members, while others are rarely extended to rivals. Keeping track of roles, alignment, identity, trust, and related social parameters thus helps to sharpen and tighten the focus of our perception and decision processing.

In the social agent simulation community, there are a large number of approaches that we make some use of in this module, however, we are aware of no approaches which satisfy the needs of a social module that operates within a unifying architecture of behavior (Figure 1). For example, artificial life simulations including cellular automata and virtual swarms, among others, provide examples of identity and behavior shifting due to peer pressure and social neighbor influence, however, these omit cognitive or perceptual processing, and rarely include biologies: e.g., see Lustick (2002). Social network analysis methods, in turn, elevate the roles and relations side of this significantly, drawing attention to how members in cliques operate as facilitators, gatekeepers, and so on, however, the agents in these networks have no emotions or cognitive abilities such as lookahead or planning: e.g., Buchanon [2002]. Game theoretic approaches such as iterated prisoner dilemma implement multi-ply lookahead about next moves of other agents based on social alignment and trust, but only at the simplest of levels and with no processing about the motivations driving other agents. Hence they can only shift strategic approach through rebirth in mutational forms: e.g., Cederman [1997]. Even the emotion model implementations that had some value in the prior module offer only limited capability of value here since they tend to focus on a very few parameters that impact emotions about relationships. Still, each of these ideas is a potentially valuable

contribution, and those attempting social modules must either find a way to interoperate some of these softwares, or a way to integrate their contributions so as to study them and drive to a new and useful synthesis. We have chosen the latter approach for now, though it does not preclude subsequent interoperation for scaleup with agents at varying levels of granularity.

Most emotion model implementations suggest social parameters such as the following from OCC and that we implement in PMFserv’s Social Module for each agent (see ‘relationship tanks’ box of Figure 2):

- Agency vs. Object – the degree to which the other agent is thought of as human vs. inanimate object. This shifts how we apply our standards (and how PMFserv applies the standards tree of a given agent). It is far easier to apply hurtful actions if we first objectify our opponent.
- Cognitive Unit – This pertains to a phenomenon in which group identity shifts with context. The canonical example of this involves two sports fans who meet by chance in a rival city. By virtue of being surrounded by others who are different from them, these two identify with one another in a way that would not occur in their home city. In cellular automata artificial society implementations, agents tend to shift identity when surrounded by neighbors of all one identity. In the emotion models of the prior section, the agents tend to reason about and reflect their emotion in such settings. Cognitive unit theory suggests they may become entrenched rather than shift identity.
- Valence – This refers to the strength of a given alignment relationship. How strong is the bond with that friend, or how negative is the link to that rival?

These kinds of variables are necessary but insufficient. Since relations are rarely static, agents need an apparatus to reason about other agents, and how to dynamically reclassify them as events occur and their behavior is revealed. The missing capability is often referred to as “theory of mind”, the mechanism that allows agents to interpret the internal mental states of other agents, rightly or wrongly. A full implementation of a theory of mind mechanism can quickly lead to computational intractability. Our approach to date has been to implement two levels of capability:

- (1) A scalable capability based largely on observation of other agents’ behavior and simple rules for interpreting and relabeling social roles, alignment shifts, credibility changes, trust, etc. The approach is quite simple in concept, though takes some knowledge engineering effort to implement. First, one must delineate discrete levels of the social scales pertinent to the scenario (e.g., alignment, group membership, trust, etc.). An example scale for role and alignment that we often use is as follows:

Closest Ally	Friend	Neutral	Opponent	Worst Enemy
--------------	--------	---------	----------	-------------

Next, the knowledge engineer must delineate the rules for activating a given level or viewpoint as well as the action categories and GSP tree affordances that then prevail when those perceptual types are activated. This is done in exactly the manner of the helicopter markup of the prior section, except now the markup is for the perceptual typing of other humans rather than of inanimate objects. To date we have been able to scale PMFserv with this approach to drive as many as 1,000 agents in crowd

simulations running on a single PC. Part II of this article presents more of the details and some examples.

- (2) A less scalable approach where fewer agents are attempted, but each of these are ‘leader agents’ that seek to dynamically model the motivations and intent of the other leader agents as illustrated in the right hand box of Figure 2. Specifically, we are interested in Agent A constructing a model of Agent B’s GSP trees and then of using that model in the game theoretic sense to decipher B’s actions and speech acts (not natural language) by thinking about underlying motivations. With such a capability, Agent A is able to think more deliberately about when to alter its view of another agent. For example, a component of the trust mechanism must address how to update due to degree of success of a given agent, B, on an action just completed. Was success just a token amount, or was it resounding? And what about a failure beyond agent B’s control or capability? Falcone & Castelfranchi (2004) point out that for some uses, trust might be more properly managed via a cognitive attribution process which can assess the causes of a collaborator’s success or failure. Likewise, they also raise the question of how placing trust in B might in fact alter B’s trustworthiness to the better or worse. These suggestions are compatible with Simari & Parson’s (2004)’s suggestion that approaches are needed which describe how humans make decisions, and their startling finding that these descriptions will reduce computational complexity of intentionality modeling (as opposed to prescriptive formalisms) as one attempts larger scale games. Clearly, this type of dynamic trust process modeling is a vital capability for agents operating in worlds where deception, bluffing, and manipulation are prevalent, as in the case of political leader contexts. We also project the need for leader agents to have to dynamically maintain and reason about other observational data such as, to mention a few areas, on all their relationships, on their personal credibility in the eyes of other agents, and on “tells” that might give away when they are bluffing or being deceitful.

In the PMFserv context, since early 2004 we have been researching and developing ways to give PMFserv agents capabilities to descriptively model the intentions and reputations of other agents, and to manage discourse and speech acts intended to manipulate and sway other agents to their goals. This capability is part of an effort to adapt PMFserv for political leader modeling and so these leader agents can participate in diplomacy games to influence world situations and leaders while simultaneously seeking to “campaign for follower groups” to retain/gain power and authority. In this gameworld, potential followers can move from varying levels of support for the leader (ranging from “in-group member” to “enemy”) depending on the action of the leader. Further details of this activity may be found in: Silverman, Johns, & Bharathy (2004), Johns (2004), Silverman and Bharathy (2005), and Silverman, Rees, Toth, et al. (2005).

2.5) The Decision Making Module

The *decision making module* of Figure 2 serves as the point where diverse emotions, stressors, coping style, memories, and object affordances are all integrated into a decision for action (or inaction) to transition to a new state (or remain in the same state). In essence, at each tick of the simulator’s clock, each agent must be able to process the

following information: current state name (or ID); stress-based coping mode (Ω_i where $i = 1,5$); currently afforded transitions and what action might cause those state transitions (a_{mn} in $A(\Omega)$); and subjective desires for each state based on 11 pairs of emotional scales summed into an overall utility score, SEU. Using all of this information, the agent must select a decision style (Φ , defined below) and process the information to produce a best response (BR) that maximizes expected, discounted rewards or utilities in the current iteration of the world. The decision module is thus governed by the following equation:

$$\text{BEST REPLY (BR}_t) = \Phi_{\text{ISTRESS}, \Omega} \{u_{mn} (s_t, a_{mnt}), p_{mn}\}, \text{ subject to } a_{mnt} \in A(\Omega) \quad [4.0]$$

Where,

$$\begin{aligned} \Phi_{\text{ISTRESS}, \Omega \{.\}} &= \text{as defined below for the alternative values of } \Omega \\ p_{mn} &= \text{perceived probability} = (1 - \Delta) e_m + \Delta_{m\tau} p_{m\tau} \\ u_{mn} &= (1-\delta) \times (\text{SEU from equation 1.0}) \\ \Delta &= \text{memory coefficient (discounting the past)} \\ \tau &= \text{number periods to look back} \\ e_m &= \begin{cases} 0 & \text{action m not situationally relevant} \\ 1.0 & \text{action m is situationally relevant} \end{cases} \\ \delta &= \text{expectation coefficient (discounting the future)} \\ A(\Omega) &= \text{action set available after coping mode-constrained perception} \end{aligned}$$

This is nothing more than a stress-constrained subjective-expected utility formulation. Also, there is a large literature on decision style functions (e.g., among many others see Bradshaw et al., 1999; and Terzopoulos, 1999), and the decision processing style function, Φ_Ω , merely indicates that there is a rich set of possibilities that one can explore within the framework proposed here. Thus, in Vigilant mode one might invoke SOAR, ACT-R, COGNET, or others. Alternatively, simulated experts can adopt the Recognition Primed Decision Making (Klein, Orasanu, Calderwood, and Zsombok, 1993) style, while novices will tend to use a more traditional decision tree.

The algorithm proposed above applies Conflict Theory where appropriate (Janis & Mann, 1977). That is, if the agent's coping mode is Panic or Unconflicted Adherence, no alternatives are weighed and the agent will execute its panic behavior or continue to do what it had already decided to do in the last cycle. Likewise, Unconflicted Change prevents any planning, and the agent must follow the next step of any existing plan. Only when stress increases and the agent's coping mode shifts to Vigilance can the agent re-plan (with any Φ_Ω method as desired). In that case, however, multi-step lookahead planning, cognitive processing with learning, and game theoretic processing can all be fully deployed. At present, we do not implement a cognitive learning model in PMFserv, however, adaptive behavior is guaranteed since various reservoirs in the biology, value tree nodes, or relationship parameters are satiated, those that aren't yet filled will become relatively more important to the agent's behavior, until they are satiated and the others begin to decay.

Also, as shown in Figure 2, we separate the implementation of Φ and BR into the two PMF boxes shown in the Cognitive Module of Figure 2. For reactive agents such as

crowd members, the Intention Management module reduces to one-ply lookahead and generally there are few action choices that need to be processed. The affordance approach suggests the actions available to crowd members and the decision PMF (left side) passes candidate actions one at a time to the Emotion Module to generate its emotions and SEU if the action succeeded and the new world future actually occurred. However, for more deliberative agents, such as leaders, the Best Reply (BR) is not so easily found. The Intention Management PMF must consider multi-step lookahead, other agent motivations (from Nested Intentionality Processing passed through the Emotion), a sequence of possible actions, and how to manage its own reputation and relationship parameters. This becomes a large search space and we are currently researching new approaches for it as mentioned already in the discussion of our leader research.

3) Lessons Learned from the Unified Architecture and PMFserv Studies To Date

This research started out with a unified behavior architecture and several guiding principles about how the systems approach might lead to instantiating such an architecture largely based on synthesizing PMFs found in the literature. Since our goal was to enable the science to be implemented and reused, these next few comments may be thought of as results from the perspective of how to facilitate better behavioral science in models.

1- Exploring Science via the Unified Architecture: Pros – The first versions of PMFserv included 3 behavioral modules (physiology/stress, emotion, and cognition) with the software cobbled together between modules as a sequence. Through usage of PMFserv across diverse scenarios we became keenly aware of the difficulties our implementation posed both as a strict sequence and in terms of replacing PMFs we found to be inadequate or which we no longer wanted in a scenario. This led to evolution of the unifying architecture as shown in Figure 1, and to our view of it as a system of many interacting and highly inter-operating parts. A major challenge of this research, is the validity of the behavioral models we derive from the literature and try to integrate within our architecture. As engineers, we are concerned with validity from several perspectives including the (1) data-groundedness of the models and parameter settings we extract from the literature, (2) accuracy of our coded version relative to the theory it purports to implement, and (3) how the unified result works in terms of correspondence of agent behavior with actual tendencies observed in the real world. In terms of data-groundedness, we conducted an extended review of the behavioral literature and found a great many studies that seem to be legitimately grounded and that possess model parameter significance from a statistical sense. We have tried to provide one such collection of PMFs in this paper. This is not the penultimate implementation rather it is at present a humble structure scientifically. 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 subsystems and within each subsystem. In that regard, our efforts to date are successful. Where feasible, we have tried to encapsulate different models and we now have an integrated fabric stitching together the models of varying groundedness and of different opinion leaders. Via the unifying architecture, we can rather easily plug in

a new opinion leader's model into a given module and run it to study its impact, its properties, and its strengths and weaknesses. Thus earlier versions omitted the value tree reservoirs and perception layers, but now they can be added or not. In modeling political leaders we need less of the physiology or stress modeling and more personality profiling and standards tree elements. If one wanted purely rationalistic reasoning, one could turn off stress and emotion processing entirely and "calculated utilities" and PMFserv's decision unit would still operate. In that sense the unifying architecture is also an "algorithm" of the behaviors that one believes are relevant.

2- Exploring Science via the Unified Architecture: Cons – The largest scientific negative to our efforts thus far is that we reveal the many places where first principles are simply missing in the field at large. Most of the models we implement in PMFserv have no prior implementations and thus are underspecified, requiring us to fill in the missing items as best as currently possible. This is true of Janis-Mann's decision conflict theory, of Damasio's Hypothesis, of the GSP value trees, of the decision theoretic processor that uses descriptive models of coping behavior, and so on. Further, the interrelations between the many parts have more frequently been neglected than attended to in the field. Our implementation thus highlights many of the issues that could form an agenda for principles of synthetic research in behavior modeling. This is not to say that reasonably realistic models of human behavior are unattainable at present. For example, over the past three years, various predecessors to the version of PMFserv framework just described have been used to construct and simulate the people and objects of a number of scenarios, that depicted emergent crowd scenes. Each featured a crowd gathering to protest a social injustice. In one series of scenarios this injustice was a roadblock that kept people from going to work. In several others, protests outside of a prison by various agents (mediated by stress, emotion, and social identity PMFs) lead to rioting and looting of nearby stores and the intimidation of police and protestor's alike. In the various crowd scenarios tested, we evaluated the impact of diverse PMFs (e.g., alternative personal and cultural value levels, impact of chanting and taunting, and diverse security doctrine/orders and behavior) on crowd behavior and on when new crowd equilibria emerged (e.g., peaceful protest, scatter, riot). These efforts enabled us to document a number of lessons learned about the replication of anticipated emergence of different types of crowd behavior (Silverman, Johns, O'Brien, Weaver, and Cornwell, 2002; Silverman, Johns, Weaver, O'Brien, and Silverman, 2002; Cornwell, Silverman, O'Brien, and Johns, 2002; Johns & Silverman, 2001). As an example of correspondence checking, the crowd literature (Horowitz, 2001; McPhail & Wohlstein, 1983) indicates that looting tends to occur when young unemployed males (who rarely join organized activities) take advantage of chaos and distracted security forces. In our simulations, female protesters and employed men tended to flee from riot situations, while unemployed men lurked on the fringes of the protest scene and then proceeded to riot and loot if violence and chaos occurred. Violence and chaos generally occurred when we added provocateur agents to the protest scene, decreased the amount of security agents present, and altered the security agent's standards tree to permit them to strike provocateurs who are overly persistent. This type of result indicates at least surface correspondence and helps to increase confidence in the workings of the PMF

collection.

3- Introducing Visual Programming in the Physiology Module and Beyond – The earliest versions of the physiology module had hard-coded PMFs for eight reservoirs. Updating these, adding new ones, or altering the collection required someone to learn the code, a prohibitive activity. Through the usages mentioned above we had to add new PMFs for the effect of chanting, weaponry, adrenaline, and so on. Eventually it dawned on us that the proper way to do this was to provide this module with a visually programmed user interface, following the model-view-controller design pattern. This has proved to be a valuable time saver for creating new scenarios that have need of new or different PMFs. Due to the success in permitting users to input models without programming barriers, we have since applied this same model-view-controller design pattern to each of the other modules (e.g., integrated stressor tanks, GSP tree editing as in Figure 5, affordance editing, social relation parameter tanks, and so on) though few of the GUIs are shown in this paper. As Part II will explain, we have also found it useful for training developers to be able to use these same ideas in the reuse of digital casts as well as in runtime pausing and editing of characters' PMF settings as scenarios play out.

4- Advancing Affordability Through Affordances – In all our early versions of PMFserv, we were caught in the tension between viewing our agents as finite state machines vs. infinite state machines. In the former approach, since there are only a finite number of actions for a bot to choose from, it was convenient to think of the agents as iterating around Markov chains. A node on the chain would describe the state that an agent was in, regardless of the myriad of physiologic, stress, and/or GSP tree activations that might exist for alternative agents arriving at that state. This way of thinking rapidly began to break down as our scenarios grew in scope and complexity. We were soon creating hundreds of Markov chains in complex hierarchies to capture the many activity sets a given agent might migrate through. Even small changes to scenarios often meant major revisions to dozens of Markov chains, a time consuming task: e.g., see Weaver, Silverman, et al. (2001). Observing that there was little need for finite state machine constraints in our simulated worlds, we abandoned the Markovian approach and implemented the situated cognition form of perception. We believe from our own experiences that this radically improves maintainability -- much easier to add new objects to the world, and the Markov chain gets generated dynamically at runtime as scenario details unfold. It's a technique used in some popular videogames (e.g., The Sims) where it has proven to be highly programmable by the non-programmer public (power users). In an earlier paper we presented an informal proof of this thesis where we compared the efforts of our own scenario programmers under the old and new approaches: [Cornwell, O'Brien, Silverman, & Toth, 2003]. We have yet to conduct a formal proof but we believe such an analysis would demonstrate a non-monotonic relationship between new objects/events in the world and representations in the mind. In the standard, traditional symbolic approaches, the relationship is monotonic increasing, if not exponential, not only with respect to knowledge management but also with respect to

software maintenance. In any event, we have newly switched to this approach and will present some of its capabilities in Part II of this paper.

5- Finding the Synergies -- In bringing together the parts shown in Figure 2 and implemented in PMFserv, we have discovered that we have created a powerful, readily adaptable capability that extends beyond its parts. The engineering advances of lessons 3 and 4 opened up the door for users to achieve this potential. As a few examples:

- Cornwell, et al. (2002) studied the impact of chanting on crowd members and found that music PMFs from the literature could be implemented in a month's time via a combination of existing capabilities of the Perception, Emotion, and Social modules. He summarizes studies of how domestic crowd behavior shifts with and without chanting activity;
- In two separate studies, Bharathy et al. (2002, 2003) worked with our university's sleep center and a trauma surgeon, respectively, and was able to take fatigue models and trauma scoring methods from the literature along with results from the field and to implement these models within PMFserv and tune them adequately for agents in our various crowd emulations. He subsequently added the impact of stimulants such as adrenaline and Khatt.
- Lombard et al. (2003) and Silverman et al (2004) used open sources and content analysis to instantiate the value systems of cultural archetypes that recreate the Bakarra market denizens of Black Hawk Down. That case study is reported in detail in Part II of this article.
- In an 8 week graduate class project, Dobeck, Gadiraju, & Mason (2004) used PMFserv to create a number of household pets and to simulate their personalities, needs, and behaviors in reacting to toys, furniture, food, each other, and human inhabitants in their owner's house.
- In about 2 person months of effort, Bharathy (2005), and Silverman and Bharathy (2005) successfully implemented an off-the-shelf political leader Personality Profiling Tool within the existing value system of the Emotion Module and applied it to construct models of and recreate behaviors of various leaders from the 3rd Crusade (e.g., Richard, Saladin, Emir of Acre, etc.).
- Johns (2004) and Silverman, Johns, & Bharathy (2004) explain how the existing PMFserv collection is being used and extended to support the modeling of world leaders able to play an existing strategic dipolomacy game described in Silverman, Rees, Toth, et al. (2005). This is an example of how agents can use PMFserv to model the value systems of other agents (nested intentionality), manage their own reputation and credibility, and strategize in game theoretic settings.

In all but the last PMFserv usage case study, there was no new programming required. These application builders were able to visually tune the existing software and edit parameters to achieve their study objectives. In the last of these cases, the new capability is being added as PMFs that future users can visually program and benefit from as well.

4. Conclusions and Next Steps

This article has reviewed the results of six years of research on ways to enhance the realism of synthetic agents. We have pursued this agenda with attention to data-groundedness of the models and PMFs we incorporated and with the hope of providing a framework to foster the easy inclusion, replacement, and study of a wide array of physiologic, stress, emotion, cultural, social, and decision models -- enabling better science to be inserted into models of synthetic agent behavior. This article concluded with lessons on how we have tried to improve the engineering of this framework as we responded to new scenario demands that required the rapid updating and swapping of PMFs and their related parameter sets.

One final result of all this framework improvement effort, and of our domestic crowd scene results as mentioned above, is that we were asked to apply PMFserv in an attempt to recreate portions of the crowd and militia behaviors observed in the Ranger operation in Mogadishu as popularized in the book and movie: *Black Hawk Down*. This raised the prospect of modeling different cultural standards and personal value sets. Further, the invitation was to use our PMFserv to drive the characters in a videogame engine called *Unreal Tournament*. This meant we would also be examining how to embed the PMFserv in other vendors' systems and along with other forms of agent modeling. We were quite enthused about this invitation, and present the results in Part II of this article along with the results of a serious validation test of our framework – that of rapidly composing characters from another culture that are faithful to the behaviors observed in that scenario.

Acknowledgement

The PMF related research summarized here and PMFserv were supported by research grants from the Defense Modeling and Simulation Office (DMSO) and the Office of Naval Research (ONR), while the GSP Tree (emotion) subsystem was supported by grants from the Ackoff Center and gifts from the General Motors Corporation. This research has benefited from the help of several teams of students in the lead author's courses and research staff – too many to mention by name. Further, we'd like to thank Eileen Bjorkman, Dexter Fletcher, Capt. Mike Lillienthal, Sue Numerich, Joe Toth, John Tyler, Ruth Willis, and Michael Young for many useful discussions about how to apply PMFserv. Any claims are the responsibility of the authors alone.

References

- Anderson, J.R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Badler, N., Allbeck, J, Zhao, L, Byun, M. (2002). Representing and Parameterizing Human Behavior. *Proc. Computer Animation*, IEEE Press:Los Alamitos, 133-43.
- Bates, J (Winter 1992), Virtual Reality, Art, and Entertainment, *Presence: The Journal of Teleoperators and Virtual Environments*, 1(1):133-138.
- Belavkin, R., Ritter, F., & Elliman, D. (1999). Towards including simple emotions in a cognitive architecture in order to fit children's behavior better. In Proceedings of the 1999 Conference of the Cognitive Science Society. Mahwah, NJ: Lawrence Erlbaum, p. 784.
- Bharathy, G. (2005). *Modeling the Behavior of Synthetic Human Agents in Virtual Worlds. A Systems Methodology for carrying out Knowledge Engineering & Modeling of Cognitively Detailed Agents*, Doctoral dissertation proposal, University of Pennsylvania, 2005 (available from the author at bharathy@seas.upenn.edu).
- Bharathy, G. (2003, May). Physiology Tank Report, prepared for SYS 508, Philadelphia: Univ. of Pennsylvania (available from the author at bharathy@seas.upenn.edu).
- Bharathy, G., Chung, C. Y., Silverman, B. G., & Cornwell, J. (2002). PMF Anthology Addendum: Adrenaline, khat, exertion, injury, and trauma. Prepared for Human Behavior Program, Defense Modeling and Simulation Office, Alexandria, VA. Available at http://www.seas.upenn.edu/~barryg/PMF_Addendum1.doc
- Bjorkman, E, Barry, P, Tyler, J, (Fall 2001) "Results of the Common Human Behavior Representation And Interchange System (CHRIS) Workshop", in *Simulation Interoperability Workshop*, SISO
- Boff, K.R., & Lincoln, J.E. (1988). *Engineering data compendium: Human perception and performance* (HED Tech Report - 3 Volumes). Wright Patterson AFB, OH.
- Bradshaw, J.M., Greaves, M., Holmback, H., Karygiannis, T., Jansen, W., Silverman, B.G., et al. (1999). Agents for the masses? *IEEE Intelligent Systems*, 14 (2), 53-63.
- Buchanan, M. (2002). Nexus: Small worlds and the groundbreaking science of networks. New York: W. W. Norton.
- Lars-Erik Cederman (1997). *Emergent Actors: How States and Nations Develop and Dissolve*, Princeton: Princeton University Press.
- Chong, R. S. (2004). Architectural Explorations for Modeling Procedural Skill Decay. *Proc. Sixth International Conference on Cognitive Modeling*, Pittsburgh, PA,
- Cornwell, J., O'Brien, K., Silverman, BG, Toth, J. (2003). Affordance Theory for Improving the Rapid Generation, Composability, and Reusability of Synthetic Agents and Objects. 12th Conf on Behavior Representation in Modeling and Simulation (BRIMS, formerly CGF), SISO, 12-15 May 2003.
- Cornwell, J.B., Silverman, B.G., O'Brien, K., & Johns, M. (2002). A demonstration of the PMF-extraction approach: Modeling the effects of sound on crowd behavior. *Proceedings of the Eleventh Conference on Computer Generated Forces and Behavioral Representation, Orlando, Florida*, 107-113.
- Damasio, A.R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Avon.

- Dobeck, E, Gadiraju, S, Mason, A, (2004), *Spot the Dog: Using PMFserv to Develop Petworld*, ESE 608 Term Project, Philadelphia: University of Pennsylvania.
- Elliott, C. (1992), *The Affective Reasoner: A process model of emotions in a multi-agent system*, Evanston: Northwestern Univ. (dissertation)
- Gibson, J. J. (1950). *The perception of the visual world*. Boston: Houghton-Mifflin.
- Gibson, J.J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.
- Gillis, P.D., & Hursh, S.R. (1999). Using behavior moderators to influence CGF command entity effectiveness and performance. *Proceedings of the Eighth Conference on Computer Generated Forces and Behavioral Representation, Orlando, Florida, 237-251*.
- Gmytrasiewicz, P, Lisetti, CL (2000), Using Decision Theory to Formalize Emotions in Multi-Agent Systems, *ICMAS 2000*: 391-392
- Hammond, K. (2000). *Judgments under stress*. New York: Oxford Press.
- Hayes-Roth, B. (1998). Animate characters. *Autonomous Agents and Multi-Agent Systems, 1 (2)*, 195-230.
- Hendy, K.C., & Farrell, P.S.E. (1997). *Implementing a model of human information processing in a task network simulation environment* (Report No. DCIEM-97-R-71). Downsview, Ontario: Defence Research & Development Canada.
- Horowitz, D.L. (2001). *The deadly ethnic riot*. Berkeley: University of California Press.
- Hursh, S.R., & McNally, R.E. (1993). Modeling human performance to predict unit effectiveness. In D. Kamely, K. A. Bannister, & R. M. Sasmor (Eds.), *Army science: The new frontiers: Military and civilian applications* (pp. 309-328). Woodlands, TX: Borg Biomedical Services.
- Janis, I.L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment*. New York: The Free Press.
- Johns, M. (2004), *Deception and Trust in Semi-Competitive Environments*, Doctoral dissertation proposal, University of Pennsylvania, (available from the author at mjohns@seas.upenn.edu).
- Johns, M., & Silverman, B.G. (2001). How emotion and personality effect the utility of alternative decisions: A terrorist target selection case study. *Proceedings of the Tenth Conference on Computer Generated Forces and Behavioral Representation, SISO, Norfolk, Virginia, 55-64*.
- Klein, G.A., Orasanu, J., Calderwood, R., & Zsombok, C.E. (1993). Decision making in action: Models and methods. Norwood, NJ: Ablex.
- Labrou, Y., Finin, T., & Peng, Y. (1999). Agent communication languages: The current landscape. *IEEE Intelligent Systems and Their Applications, 14 (2)*, 45-52.
- Laird, J, van Lent, M (2001), Human-Level AI's Killer Application, *AI Magazine*, summer, 15-25.
- Lazarus, R. (1991). *Emotion and adaptation*. Oxford: Oxford University Press.
- Lombardo, J., Pollack, B., George, M., & Brounstein, A. (2003) - Modeling human behavior within the "Black Hawk Down" scenario. University of Pennsylvania, Senior Design Project Report.
- Lustick, Ian S.(2002, June). "PS-I: A User-Friendly Agent-Based Modeling Platform for Testing Theories of Political Identity and Political Stability," *Journal of Artificial Societies and Social Simulations*, Issue 5, no. 3

- McPhail, C., & Wohlstein, R.T. (1983). Individual and collective behaviors within gatherings, demonstrations, and riots. *Annual Review of Sociology*, 9, 579-600.
- Ness, J, Ritzer, D, Tepe, V (eds.), (Jan.2005) The Science and Simulation of Human Performance, NY: Elsevier,.
- Ortony, A., Clore, G.L., & Collins, A. (1988). *The cognitive structure of emotions*. Cambridge: Cambridge University Press.
- Payr,S, Trappl, R (2004), *Agent Culture: Human-Agent Interaction in a Multicultural World*, Mahweh: Erlbaum
- Reeves, B., Nass, C. (1996). *The media equation: How people treat computers, televisions, and new media like real people and places*. New York: Cambridge University Press.
- Ritter, F, Avraamides, MN, Councill, IG (May 2003), Validating Changes to a Cognitive Architecture to More Accurately Model the Effects of Two Example Behavior Moderators, *Proc. 11th Conf. Behavior Representation in Modeling and Simulation*, SISO: Orlando.
- Roseman, I., Spindel, M. S., & Jose, P. E. (1990). Appraisals of emotion-eliciting events: Testing a theory of discrete emotions. *J. Personality and Social Psychology*, 59 (5), 899-913.
- Silverman, B.G., Johns, M., Shin, H., Weaver, R. (2002). Performance moderator function (PMF) anthology for human behavior modeling in military simulations. Prepared for Human Behavior Program, Defense Modeling and Simulation Office, Alexandria, VA. Avail. <http://www.seas.upenn.edu/~barryg/PMFset.zip>
- Silverman, B. G. (1991). Expert critics: Operationalizing the judgment/decision making literature as a theory of “bugs” and repair strategies. *Knowledge Acquisition*, 3, 175-214.
- Silverman, B.G., Johns, M., O'Brien, K., Weaver, R., & Cornwell, J. (2002a). Constructing virtual asymmetric opponents from data and models in the literature: Case of crowd rioting. *Proceedings of the 11th Conference on Computer Generated Forces and Behavioral Representation, Orlando, Florida*, 97-106.
- Silverman, B.G., Johns, M., Weaver, R., O'Brien, K., & Silverman, R. (2002b). Human behavior models for game-theoretic agents. *Cognitive Science Quarterly*, 2 (3/4), 273-301.
- Silverman, B.G., Might, R., Dubois, R., Shin, H., Johns, M., & Weaver, R. (2001). Toward a human behavior models anthology for synthetic agent development. *Proceedings of the Tenth Conference on Computer Generated Forces and Behavioral Representation, Norfolk, Virginia*, 277-285.
- Silverman, B.G., Johns, M., Bharathy, G. (2004, August). “Agent-Based Simulation of Leaders.” ACASA/UPenn, Tech Report, (available from lead author).
- Silverman, BG, Bharathy, G. (May 2005),“Modeling the Personality & Cognition of Leaders,” in 13th Conf on Behavioral Representations In Modeling and Simulation, SISO.
- Silverman, BG, Rees, R., Toth, J, et al. (2005, May), “Athena’s Prism—A Diplomatic Strategy Role Playing Simulation for Generating Ideas and Exploring Alternatives,” to appear in Proc. of First Internat’l Conf on Intelligence Analysis, MacLean, VA: Mitre.

- Sloman, A., & Logan, B. (1999). Building cognitively rich agents using the SIM_AGENT toolkit. *Communications of the ACM*, 42 (3), 71-77.
- Tambe, M., Johnson, W.L., Jones, R.M., Koss, F., Laird, J.E., Rosenbloom, P.S., et al. (1995). Intelligent agents for interactive simulation environments. *AI Magazine*, 16 (1), 15-39.
- Terzopoulos, D. (1999). Artificial life for computer graphics. *Communications of the ACM*, 42 (8), 33-42.
- Toth, J.A. (1995). Book review of *Reasoning agents in a Dynamic World: the Frame Problem* (K.M. Ford and P.J. Hayes, eds.), *Artificial Intelligence Jnl*, 73: 323-369.
- Weaver, R., Silverman, B.G., Shin, H., & Dubois, R. (2001). Modeling and Simulating Terrorist Decision-making: A "performance moderator function" approach to generating virtual opponents. *Proceedings of the Tenth Conference on Computer Generated Forces and Behavioral Representation, Norfolk, Virginia*, 39-44.
- Yerkes, RW, Dodson, JD, (1908) The Relation of Strength of Stimulus to Rapidity of Habit-Formation, *J of Comparative Neurology and Psychology*, 18, 459-482
- Zachary, W., Santarelli, T., Ryder, J., Stokes, J., & Scolaro, D. (2001). Developing a multi-tasking cognitive agent using the COGNET/iGEN integrative architecture. *Proc 10th Conf Computer Generated Forces and Behavioral Representation, Norfolk, VA: SISO*, 79-90.
- Zhang, W. & Hill, R. (2000). Situation Awareness and Assessment: An Integrated Approach. In *Proceedings of the 9th Computer Generated Forces and Behavioral Representation Conference, Orlando, FL*, pp. 365-375.