

HUMAN BEHAVIORAL REPRESENTATIONS WITH REALISTIC PERSONALITY AND CULTURAL CHARACTERISTICS

Wayne Zachary, Ph.D.
Jean-Christophe Le Mentec, Ph.D.
CHI Systems, Inc.
1035 Virginia Drive
Ft. Washington, PA 19034
215/542-1400
wzachary@chisystems.com, jclementec@chisystems.com

Lynn Miller, Ph.D., Stephen Read, Ph.D.
Department of Psychology, University of Southern California
Los Angeles, California 90089-1061
Lmiller@usc.edu, read@usc.edu

Gina Thomas-Meyers
US Air Force Research Laboratory AFRL/HECS
2698 G Street
Wright Patterson AFB, OH 45433-7604
(937) 255-9474
Gina.Thomas-Meyers@wpafb.af.mil

Wayne Zachary is President and CEO of CHI Systems, Inc. He holds a Ph.D. in Anthropology & Computer Science from Temple University. His research interests include human-computer interaction, modeling and simulation of human cognitive and organizational behavior, and decision support systems.

Jean-Christophe Le Mentec is Principal Software Engineer at CHI Systems, Inc. He received a Ph.D. in artificial intelligence from the University of Nancy (France). His principal research interests include real-time capabilities, blackboard systems, distributed artificial intelligence, planning, and adaptive reasoning.

Lynn Miller is a Professor at the Annenberg School for Communication and is internationally known for her work on personality and interpersonal communication. She holds a Ph.D. in Personality Psychology from the University of Texas, Austin.

Stephen Read is a Professor of Psychology at the University of Southern California, and is internationally known for his work on computational models of social perception, causal reasoning and personality. He holds a Ph.D. in Social Psychology from the University of Texas at Austin.

Gina Thomas-Meyers is a research psychologist at the US Air Force Laboratory in Dayton, Ohio, where she conducts cognitive engineering research and development in models of human cognition and engineering of work-centered systems to optimize human capabilities.

HUMAN BEHAVIORAL REPRESENTATIONS WITH REALISTIC PERSONALITY AND CULTURAL CHARACTERISTICS

W. Zachary, J.-C. Le Mentec, L. Miller, S. Read, & G. Thomas-Meyers

ABSTRACT

The Personality-enabled Architecture for Cognition or PAC is a new technical capability to create human behavioral representations (HBRs) with pre-defined and specific personality traits and cultural characteristics. This capability meets a current and growing need for human models that exhibit personality and cultural variability. The need arises from multiple sources, but primarily from the increased frequency and complexity of military operations involving coalition forces with great cultural and personality diversity, and the growing trend toward asymmetrical conflicts involving adversaries with poorly understood cultural values, characteristics, and behavior patterns. PAC integrates theory and empirical data from personality psychology, social psychology, cognitive science, and neuroscience. Unlike existing cognitive architectures that attempt to build affective and personality factors as customizations to an underlying formally rational symbolic architecture, PAC uses dimensions of personality, emotion, and culture as foundations for the cognitive process. The structure of PAC allows it to function as a personality/emotional layer that can be used stand-alone or integrated with existing constrained-rationality cognitive architectures. In addition, a set of tools was developed to support the authoring, analysis, and testing of PAC HBRs. Demonstration PAC-HBRs were based on characters from VECTOR, a game-based cultural familiarization trainer.

1. INTRODUCTION/ PROBLEM

In today's multi-national military endeavors, factors of personality, affect, and culture influence decision-making, mission performance, and communication in many ways. The post-Cold War adversary has become a regional or non-governmental force, driven by a cultural/religious (rather than political) ideology, and led by 'strong man' leaders whose personality strongly affects the military and political strategy. At the same time, the rapid movement toward net-centric warfare is introducing complexities into command and control (C2). Operations may occur on short notice and involve novel combinations of forces across services involving allied and coalition participants. As the need for collaboration and effective team-work in C2 is increasing, the operations are staffed by personnel with highly varied cultural, personality and system backgrounds. There is, thus, a very strong need for more realistic representations of these factors in the human behavioral representations (HBRs) that are used in military simulations for C2 training, C2 mission rehearsal/analysis, and even C2 system acquisition.

Prior HBR research has attempted to build cultural and personality factors on top of existing symbolic HBR architectures such as ACT-R, Soar, or COGNET/iGEN[®]. However, a growing body of results from both psychology and neurophysiology provide evidence that emotion and personality are *not* optional, non-essential epiphenomena of cognition, but are rather foundational features of the organization and robustness of human behavior. This argues that a general and reusable approach to incorporating personality, emotion, and culture into HBRs must incorporate these factors *into the core architecture itself*. This paper describes a research effort

to develop a new technical capability for human behavioral representation (HBR) called the Personality-enabled Architecture for Cognition (PAC). The PAC is a general software model of human cognition that is able to produce HBR agents with realistic personality traits, emotions, and to some degree, cultural characteristics. The research integrates goal-oriented personality constructs from psychology and neuroscience with existing symbolic cognitive architectures to yield substantially new capabilities for HBRs of the kind currently widely used in simulations for training, mission rehearsal/analysis, and simulation-based acquisition, as discussed below.

The bulk of the paper discusses the theory behind PAC, its software architecture, knowledge representation, and HBR development tools. Prior to these discussions, we further explore the relevance of this work to the broader world of command and control. The paper concludes with a brief example of a PAC in the context of cultural training, and a discussion of broader implications and future directions for this research.

2. GOALS AND CONSTRAINTS FOR PAC DEVELOPMENT

The PAC development was undertaken in light of the training, mission rehearsal, and system engineering needs summarized above. To maintain a practical focus on meeting these needs, the following requirements were used to guide the PAC research:

1. Be based in psychological data and theory – the core set of concepts, constructs, and operating principles had to be explicitly based on accepted psychological theories of personality, culture, and cognition in the scientific literature, to provide it with an explicit and principled scientific foundation.
2. Provide a generative, parameterized modeling capability – needed to be implemented at a level of abstractness that allows PAC:
 - a) to create HBR models that *generate* behaviors as the (external) situation unfolds (i.e., it needed to be more than a simple scenario-based scripting language);
 - b) to include explicit parameters that allow a given HBR to exhibit different personality traits leading to different behavioral sequences depending on the values set for various parameters (i.e., to have specific personalities 'dialed in' before the HBR model is executed).
3. Incorporate validated parameter settings – this allows the PAC architecture and its various parameters needed to be empirically tuned, verified and validated against human behavioral data so that it can be used with confidence in simulation/game-based training and other C2 applications.
4. Be inter-operable with emerging simulation paradigms, including game-based simulations – A PAC-generated HBR needed to be able to be integrated and executed in the broadest range of simulation environments possible, but including at least standards-based (e.g., HLA-based) constructive simulations and game-engine based simulations, so that PAC-based HBRs could be used whenever and wherever needed.
5. Support model-developer needs – with cost of HBR development a growing concern, PAC also needed to include tools that structure and support the development, testing, and debugging of PAC-based HBR models, so that PAC could be used efficiently and effectively in application development.

3. THE THEORY UNDERLYING PAC

The PAC research has combined theories and data from personality psychology, social psychology, cognitive science, and neuroscience. While cognitive science and neuroscience have long represented and tested theories as computational cognitive models, the same is not true for personality psychology. Translating the core theories from personality psychology into a more computational form and integrating them with the core theories in cognitive science was the major theoretical challenge to the development of PAC.

Personality is defined as enduring tendencies to think, feel, and behave in consistent ways. The current understanding of personality in psychology derives from a lexical approach, that is, from analyses of the words in a given language that describe personality traits, and the semantic structure underlying these. Work on the lexical analysis of trait language (e.g., Saudicer & Goldbert, 1996a) and work on the structure of a variety of different trait scales (e.g., Tellegen & Waller, 1997; Costa & McRae, 1992; and Wiggins & Trapnell, 1996) have given rise to what is commonly called the Five-Factor Model of personality. The Five-Factor model, also called simply the 'Big Five', is a hierarchical model of personality, in which relatively narrow and specific traits are organized in terms of five broad factors: Neuroticism, Extraversion, and Openness to Experience, Agreeableness, and Conscientiousness (McRae & John, 1992). Overall evidence in support of the Big Five taxonomy has grown over the past two decades, and it has become the generally accepted theory for personality psychology. However, even though the Big Five model has demonstrated generalizability across different cultures and languages, it is by itself not a cognitive framework, in the sense that the Big Five provides no insight on the internal information representations and processes that give rise to the categories of behavior that it encompasses.

Two related lines of research, however, enabled us to fill this gap in the PAC research. One of these was the work by Read, Miller and colleagues to identify possible cognitive constructs underlying personality. Read, Jones & Miller (1990), and Read & Miller (1998) noted how personality could be represented as configurations and differential activations of motives¹, plans and beliefs. They noted, for example, that a trait such as “helpful” could be decomposed into the *motive* to help others, *beliefs* about the value of helping others and whether others deserve help, the *plans* to help, and the *resources* to do so. This work tied Big Five theory to the kinds of constructs used in cognitive simulation research. Subsequent research showed that the types of goals that Read and Miller talked about were organized into two levels. The first level (the motive level) contains a set of general motives that correspond to adaptive strategies that individuals use to manage and negotiate their place in everyday life. These strategies/motives include:

- seeking social bonding,
- avoidance of social separation,
- establishment of/conformance to dominance and authority relationships,
- exploration and play,
- caring and parenting,

¹ We use the term motives here to distinguish it from the term 'goals' that is frequently used in mainstream cognitive science. Goals typically are used to organize cognitive processing and behavior with regard to accomplishing a very specific, typically work-based, end-state (e.g., 'open the file', 'send target coordinates'). Motives, as we use them, also organize behavior but in a much more pervasive manner than work goals and are focused on the needs of the individual (e.g., don't get killed/hurt, from friends and other social relationships) rather than on work goals.

- self-preservation and concerns for physical safety, and
- mating.

The second level (the sensitivity level) includes only two overarching sensitivities, which govern how the first level motives are pursued and traded-off in any specific situation. These are *Behavioral Approach Sensitivity* (BAS), which governs sensitivity to reward and rewarding stimuli and a *Behavioral Inhibition Sensitivity* (BIS), which governs sensitivity to punishment and avoidance of threatening stimuli (Gray, 1987; DuPue, 1996). They represent, in essence, the underlying tendencies for approach and avoidance.

The second related line of research tied personality theory to neuroscience constructs (e.g., Panksepp, 1998, 2000). Davidson, R. J., & Irwin, W. (1999), and Davidson, R. J., Jackson, D. C., & Kalin, N. H. (2000) provide extensive evidence that the left and right prefrontal cortex (PFC) are differentially involved in the BIS and BAS systems. The left PFC seems to process positive, approach-related or appetitive motives and emotions, whereas the right PFC processes withdrawal related emotions, such as fear, disgust, and sadness. Interestingly, anger is more typically related to the approach system and seems to be processed in the left PFC, with more approach related emotions (e.g., anger associated with not achieving goals towards which one is striving). Thus, the BAS and BIS seem to be mapped into the left and right PFC, respectively, and may integrate and provide a “read-out” from the lower level *motive* systems (see Cacioppo, J. T., Gardner, W. L., & Berntson, G. G., 1999). These two levels are highly and bi-directionally interconnected such that the lower level systems send activation to the higher level systems, while at the same time the higher level systems can influence the activation of the lower level goal systems.

In addition, there is evidence for a third brain system, the Disinhibition/Constraint system (DCS) that provides for an even more general level (level three) of inhibitory control for the other systems (Tellegen, 1985; Watson & Clark, 1993). Inhibition acts to enforce selectivity among activated concepts by enhancing the differences in their activations (see Nigg, 2000). Higher levels of inhibition result in greater differentiation among concepts, as only the most highly activated concepts will remain active. As a result, variability in strength of inhibition affects the likelihood that various concepts will play a role in cognition, motivation, and behavior. Thus, DCS may govern the extent to which the system is goal-focused (resulting in enacting more goal-directed behavior) versus highly reactive to changing environments (resulting in an individual appearing more prone to distraction).

These above discussions provide the basis for the translation of personality and emotion theory into a computational frame (i.e., PAC). Figure 1 shows the resulting three-level structure which is used in PAC. Individual motives are activated as a result of interactions in and with changing situations. The activations of the motives are determined in part by situation factors and prior experience (i.e., knowledge and memory), but also by innate individual differences – baseline activations which differ from person to person. The changing activation of individual motives is also affected by the overall sensitivities set by the BIS and BAS, which are themselves changing as a result of situational factors from an innate individual baseline activation. Finally, the entire activity level of the system is further focused (or defocused) by the DCS, which again varies in activation from an innate baseline.²

² Although space limitations preclude a detailed discussion here, this overall structure has been systematically mapped back into the lexically-derived Five Factor model, showing a cognitive and neurobiological basis for the five factors.

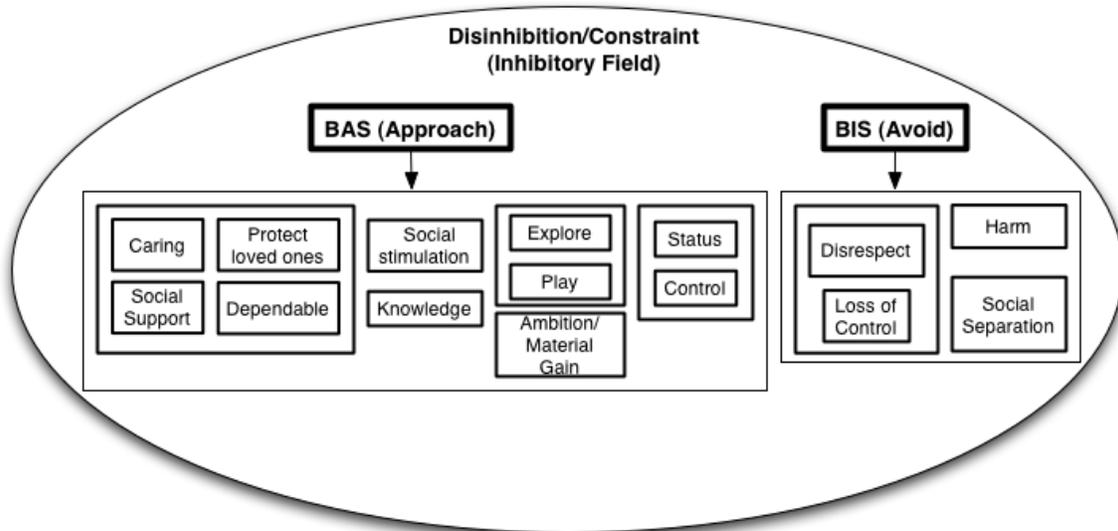


Figure 1. Interconnections of Motivational and Sensitivity Systems

The core constructs and theory from the cognitive side were more readily available. We took the general theoretical framework established in the National Academy of Science Study (Pew and Mavor, 1998), and depicted, at the most general level, in Figure 2.

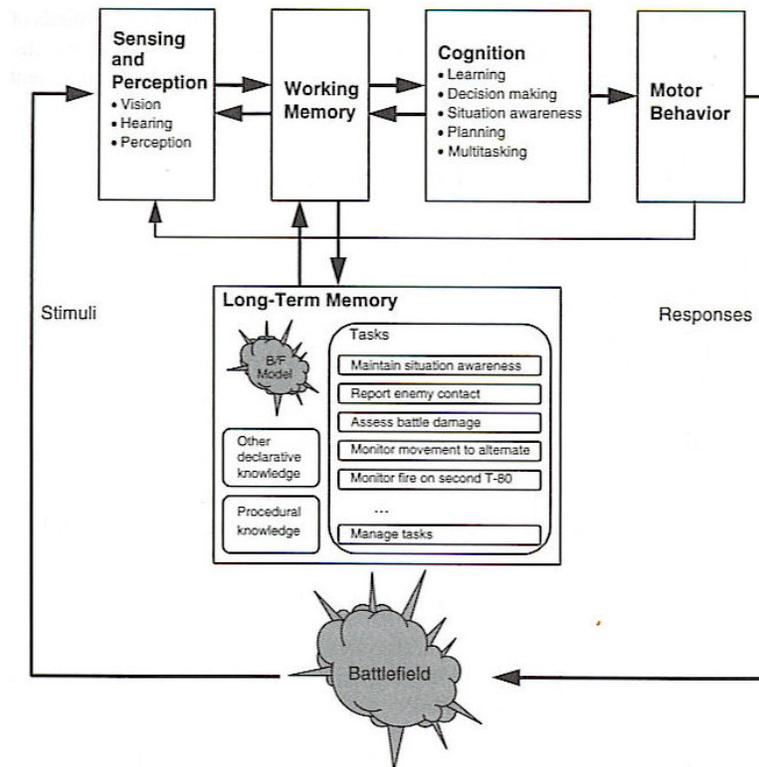


Figure 2. General Constructs and Organization for Current Cognitive Architectures (from Pew & Mavor, 1998)

This framework is used in virtually all commonly-used cognitive architectures, and has several generally-accepted features. In this framework, the person/model:

- is constantly building and maintaining an internal symbolic representation of the external world (e.g., the battlespace) that provides awareness of the external situation, as filtered by the sensory, perceptual, and interpretive processes (and knowledge);
- recognizes opportunities that the current situation (as internally represented) affords for accomplishing work goals and tasks;
- acquires or activates localized work goals from opportunities afforded by the current situation;
- recognizes constraints and conflicts between and among the various activated goals/strategies, and prioritizes or otherwise deconflicts them;
- tailors remembered or learned strategies for accomplishing the prioritized work goals to the current situational context; and
- undertakes physical actions (including verbal actions) in the environment as needed to carry out the tailored strategies.

The resulting behaviors are typically consistent with theories of constrained rationality, and show how a fixed body of knowledge can generate a range of behaviors that accomplish abstractly defined work goals across a range of different situations.

4. ARCHITECTURE AND IMPLEMENTATION OF PAC

The main architectural challenge in creating PAC was to integrate the motive and sensitivity systems framework (Figure 1 above) with the conventional cognitive architecture structure (Figure 2 above). This was accomplished in PAC by integrating the motive-based processes that make up personality as a deeper level set of processes. These personality processes operate in parallel to the constrained-rationality or processes (e.g., the process by which people accomplish work tasks) with one exception – the deeper level process does not presume to have a separate or parallel 'action' component. Instead, it leaves the action aspect under the control of the constrained rationality process. It, thus, ends the processing cycle with a step of integrating the execution of the motive-based behavioral strategies with the tailoring and execution step of the constrained rationality process. This can happen in one of two ways, either by:

- inserting additional actions into the set of actions constructed by the purely work-focused cognitive process, or
- further tailoring the work-related processes to accommodate the strategies activated by the personality-based deeper loop.

The constrained-rationality processes focus on understanding the situation and generating the work behavior (i.e., tasks) required by the situation. The personality-based processes focus on the evolving social situation and on generating and applying strategies to achieve the various personal motives that are activated. Thus, there is an on-going social understanding process (i.e., and analog of situation understanding) which recognizes situational affordances to pursue specific motivations. Whether that affordance results in an activation of the corresponding motivation depends largely on the person's baseline activation for that motive. For example, a person with a low baseline for pursuing dominance is less likely to recognize (or react to) situations which afford an opportunity to increase social status or dominance.

The unfolding social situation and the person's response to it may also result in short term perturbations to activations of the various motivations and deeper-level sensitivities (e.g.,

BIS/BAS/DBC) which, in turn, may temporarily change the behavior of the system. Over time, though, the activations of these personality factors would be expected to return to their baseline activation levels, because of the persistent nature of baseline personality traits.

This strategy for integrating personality and constrained-rationality cognition led to a general architecture for PAC as shown in Figure 3.

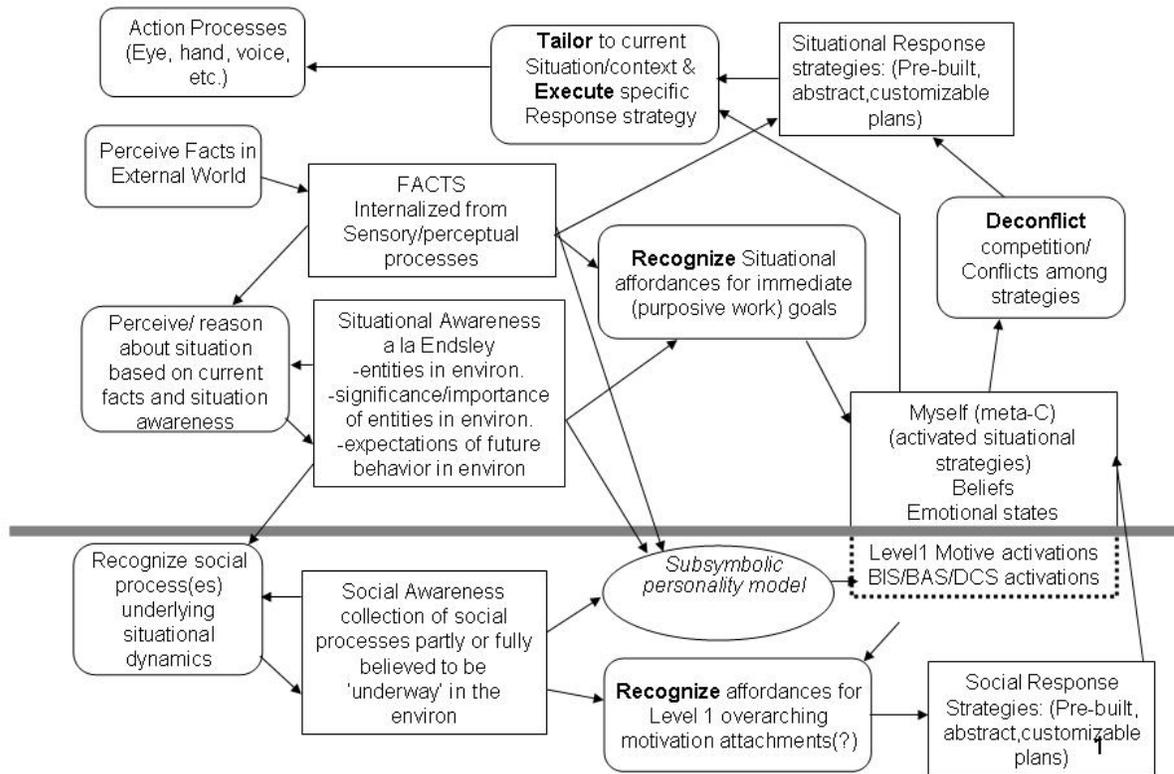


Figure 3. Conceptual Architecture of PAC

The layout of Figure 4 is deliberately organized to place all processes either above or below a line, which divides the overall architecture. The components 'above the line' correspond to the processes involved in conventional constrained rationality cognitive architectures, while the components 'below the line' make up the processes that are unique to the personality-based processes. This allowed the PAC software to be implemented as an entirely separate personality-based layer in a two-layer cognitive system, working in tandem with a constrained rationality layer (which could, in principle, be any existing system). Importantly, the connections between processes above and below the line are limited to two specific data/knowledge stores:

- 1) the contents of the (current) situation awareness, and
- 2) a metacognitive store in which the awareness of the currently active strategies and corresponding goals is maintained.

This simple interconnection between the personality layer and the constrained rationality layer, further structures the relationship between the two layers pictured in Figure 4.

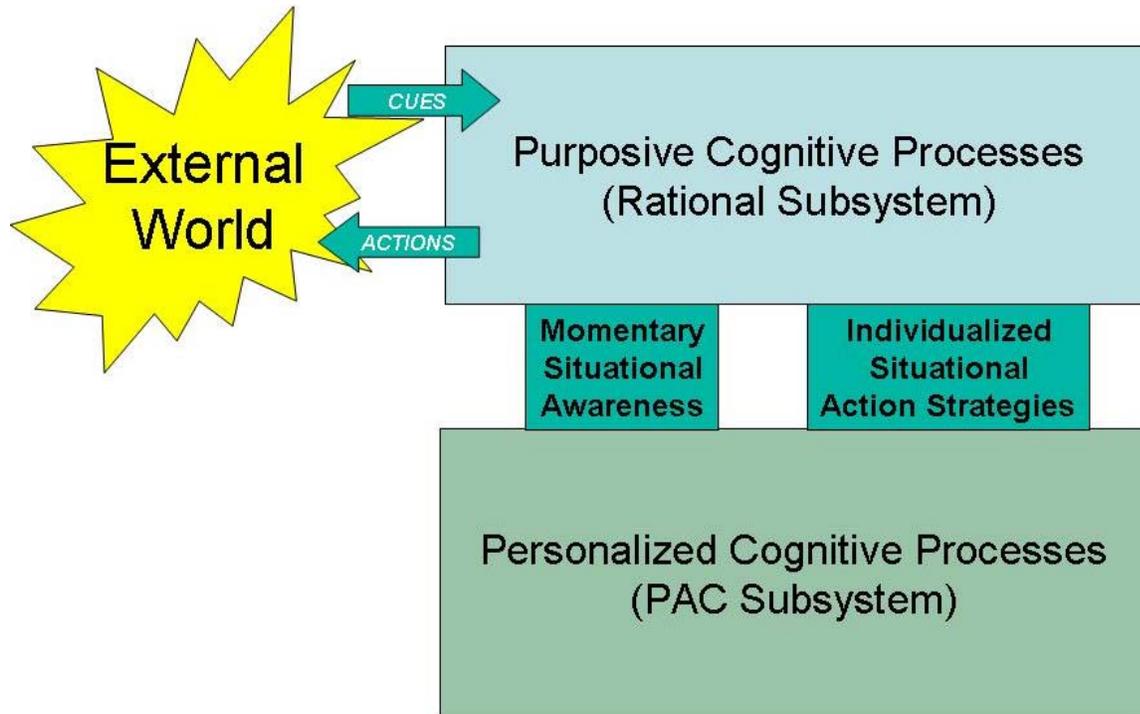


Figure 4. Component Subsystem View of PAC

The key component of the personality subsystem of PAC is the subsymbolic personality model, which integrates situational understanding with baseline activations of the general motives and the BIS/BAS/DCS sensitivities. The architecture of this component is discussed below, but requires a prior discussion of the way in which knowledge in the personality subsystem is organized in PAC.

4.1 Knowledge as Story Structures

Knowledge is represented in an extensible set of generative story structures or templates which are used both to generate behavior in complex environments and to interpret behavior of others. This representation was selected for a number of reasons. Miller and Read (1991) have argued that a simple narrative or story is fundamental to the representation of most traits. Using the example of 'helpfulness', the person must see that someone else is in need, they must intend to do something about that need, they must take action to satisfy that need, and the need must be satisfied. This forms a basic narrative with a *condition* that instigates a *goal*, which leads to the formation and enactment of a *plan*, which leads to the *outcome* of the *plan*, either success or failure. Read and Miller further argue that a story structure is central to the ways in which people represent their understanding of social interaction. This narrative structure of traits thus made a story-based representation a natural fit in our attempts to capture the interpretation of social interaction and events. This makes narrative a natural way to represent our agent's understanding. Additionally, story structures are likely to provide a structured, constrained way for the developer of PAC-based HBRs to represent human social knowledge.

The general representation of a story in PAC is as a collection of units, called Plot Units, which capture a piece of the story line and the various ways in which it might play out. Each Plot Unit is composed of a series of interconnected Action Structures and Behavioral Options.

The ‘Action Structure’ forms a micro-level representation of an intended action in the causal-chronological sequence that makes up the Plot Unit. The Action Structure specifies such elements as the character/agent (WHO), the act-type (DOES-WHAT), the modality of action (HOW), and the setting (WHERE/WHEN). Perhaps most importantly, it also specifies the opportunities that different possible evolutions of the story forward from that point afford for activation or application of the overarching motives that comprise the top level of the PAC personality model (Figure 1 above).

Each 'next step' along an alternative evolution of the story leads to a different Behavioral Option that is connected to the Action Structure. The Behavioral Option represents execution of the action AND the possible changes to the external world that the action may create (once implemented by the motor system). The Behavioral Option also points to the appropriate Action Structure that would occur next along that evolution of the story. Thus, just as an Action Structure can lead to multiple Behavioral Options, a Behavioral Option can lead to multiple future Action Structures. This gives the story structure a tree-like quality, though in formal terms is actually a semi-lattice. This is because two (or more) Behavioral Options can lead to the same future situation in the external world and thus lead back to a single Action Option. Figure 5 shows an abstract version of a full Plot Unit.

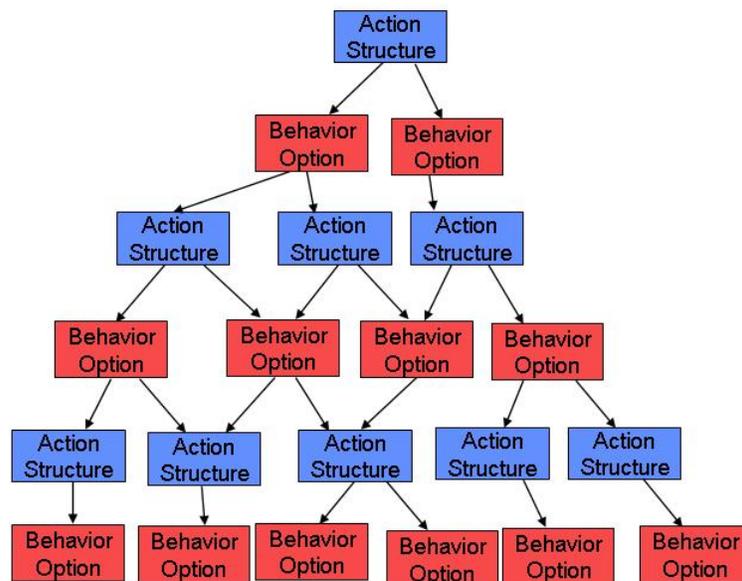


Figure 5. An Abstracted Plot Unit

From the perspective of the component architecture in Figure 4, the Action Structure represents processing that occurs within the *PAC subsystem*, while the Behavioral Option represents processing that occurs within the *rational subsystem*, as well as the consequent events that unfold in the *external world*.

The Plot Unit (and the full stories that the Plot Units comprise) shows all the expected evolutions of the story and are represented from an ego-centric view of the agent/character. Thus, for even a simple interaction to occur, there must be some general mapping of the plot units that are understood by the two actors. This is shown in Figure 6. The story in Figure 6 is joined at the point where the situational state leads to a character to exercise the (single) Action Structure at the top of the figure. Presuming that the left-hand Behavioral Option is selected, the

Behavioral Option is then executed in the external world, where some other agent/character responds with some action. That other individual is presumed to share an understanding of the general story and specific Plot Unit, but to have his/her own specific motives in the interaction. This other agent/character thus chooses the Behavioral Option at this point of the story that maximizes his or her local motives. That action causes a new state of the external world, which leads the original person or agent to undertake the action structure at the far left of the bottom of Figure 6.

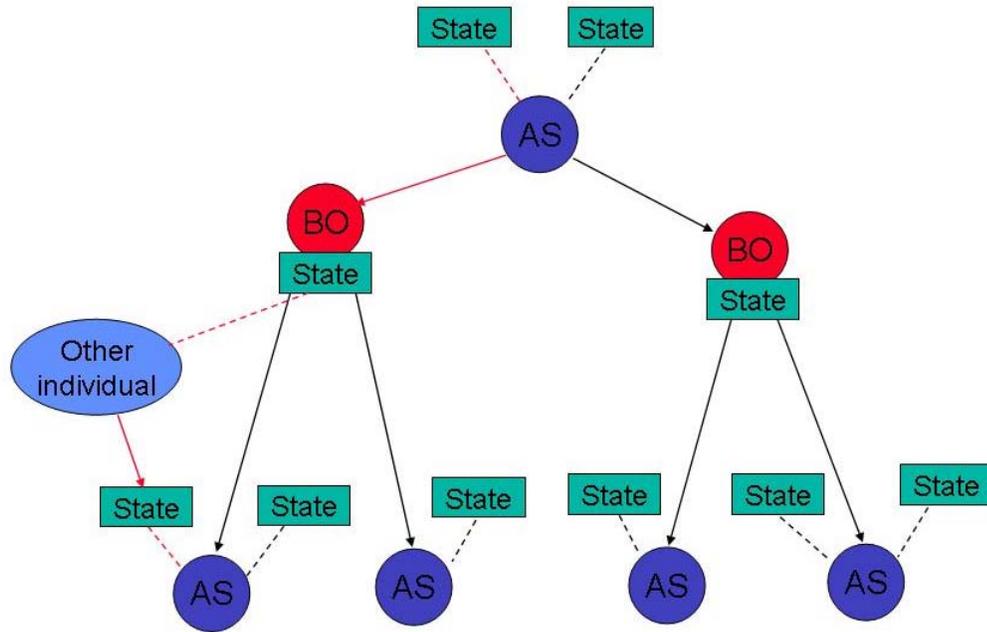


Figure 6. Plot Units as Inherently Interactive Structures

As noted above, PAC views story structures as both a model for generating behavior and a model for understanding behavior as it unfolds. The process described immediately above shows this inter-relationship. Presuming that the story structure is generally shared across the individual characters/actors, each uses the story structure:

- to recognize and interpret what the other character/actor has done by mapping the perceived action as an instance of one of the Behavioral Options available at that point in the story, and
- to generate new behaviors by selecting the Behavioral Options from the current Action Structure that are most consistent with that character/agent's motives.

Story structures are largely culture-specific, and thus are a primary source of cultural variability in PAC.

4.2 Subsymbolic Personality Model

The story structures afford opportunities for PAC HBRs to exhibit specific personality traits or trait combinations. For example, a given part of the story may afford the opportunity for a strongly assertive person to exert leadership, but at the same time may afford opportunity for an insecure person to accept projected authority and be led. The subsymbolic personality

mechanism within the PAC layer of the architecture controls the process by which inherent personality traits of the individual are exhibited in this process. This personality model is described below. This component is termed the PAC Motive Interpreter, and is pictured in Figure 7.

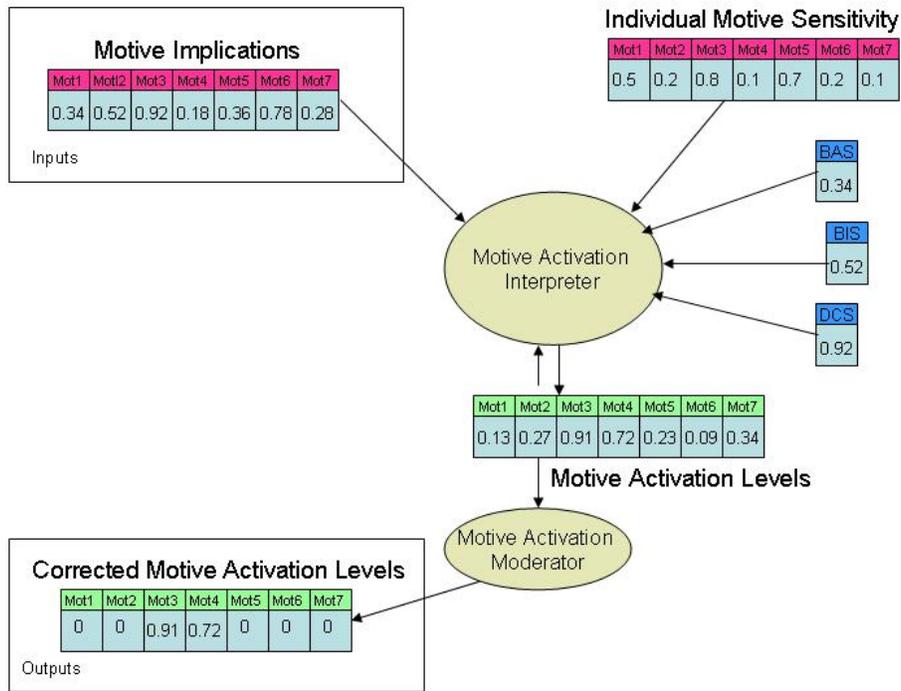


Figure 7. PAC Motive Interpreter

The Motive interpreter calculates the motive activations for each motive as each Action Structure is processed during the evolution of a story structure. At any point in time, the motive interpreter operates on three types of data:

- a set of *motive implications* from the current action structure -- A motive implication is a numerical value between 0 and 1 indicating relevance of an action structure in relation to a specific motive.
- a set of predefined individual motive *baseline activations* -- Each motive type included in PAC has a baseline activation state that represents the innate tendency of the individual being simulated. That is, it represents the degree to which that individual is predisposed to pursue that motive when and if an opportunity arises.
- three current *sensitivity levels* associated with the BIS, BAS, and DCS mechanisms -- The BAS is a general neurobiological system that affects the gain or sensitivity to stimulation of the approach related goals that fall into that system. The BIS is a general neurobiological system that affects the gain or sensitivity to stimulation of any avoid related or punished related goals. The DCS a general inhibitory field that serves to dampen or inhibit the activation of all motives, such that higher sensitivities leads to stronger goal focus and less disorganized behavior.

For each motive, the motive interpreter calculates its (new) level of activation R by using one of two formulas. The first formula is used when a motive implication for this motive is provided:

$$(1) \quad R = 1 - \frac{1}{1 + \gamma [I + S - DvC]_+}$$

Where:

$[x]_+ = x$ if $x > 0$ and $[x]_+ = 0$ if $x \leq 0$;

I is the motive implication as provided by the current action structure,

S is the individual sensitivity for this motive,

gamma (γ) is either the BIS or BAS depending on the type of motive. If the motive is of type avoidance, then $\gamma = \text{BIS}$; for all other motives, $\gamma = \text{BAS}$.

Formula (1) adjusts the activation levels to meet the opportunities afforded in the current Action Structure.

The other formula is used when the current Action Structure has no motive implication is available for a particular motive. It implements a decay mechanism that progressively returns the activation level to the individual sensitivity level:

$$(2) \quad R_n = k(S - R_{n-1}) + R_{n-1}$$

Where:

R_n is the resulting level of activation of the motive for the n cycle and, R_{n-1} the activation level at the previous cycle.

k is a decay factor,

S is the individual baseline activation for this motive.

4.3 PAC Implementation

The PAC software was implemented as an integrated development environment or (IDE) which includes the PAC model-execution software, along with a model editing component, a model execution interface and an environment framework. Thus, PAC was implemented in a way that provided the PAC model-builder with a set of tools for performing various development tasks on PAC models, ranging from model authoring to model execution. Each component is implemented in Java, to facilitate cross-platform execution.

The *IDE platform* is the least architecture-dependent component of PAC. It provides project management functionality, project navigation and windowing. The IDE platform serves as host for model specific rendering modules through a set of defined APIs. The IDE platform has been based on a set of open technologies such as FlexDock (<https://flexdock.dev.java.net/>), XStream (<http://xstream.codehaus.org/>) and the Apache project's XML tools (<http://xml.apache.org>). The *representation authoring component* connects a set of editing components to the IDE platform. Most representation entities have at least one conventional user interface in addition to alternative graphical views. This component also defines contextual help and editing aid for the representation entities. The *execution tracing/debugging component* provides a user interface allowing human-controlled individual input in addition to tracing the execution of every individual. The graphical representation of action structures has been implemented using JUNG (<http://jung.sourceforge.net/>).

4.4 A Simple PAC Example

A series of PAC-based HBRs was developed to demonstrate the technology and to verify and validate the correct operation of the architecture and its software implementation. One of these examples is described below. This example is based on interaction sequence from VECTOR, a game-based training system which allows a human trainee to interact with members of another culture to conduct cultural-familiarization. The non-trainee characters are actually HBRs; in this version, the HBRs are designed around a Kurdish Arabic cultural model. The overall training scenario allows the human trainee to interact with Arab-culture HBRs playing various roles in a military context. The vignette in the example consists of a soldier approaching a shopkeeper in order to gather information about a recent bombing. The interaction can play out in a number of ways, depending on the actions of the actors, as they move from greeting, to information request, and finally request resolution and disengagement. The example is concerned with examining the mapping of personality motive baseline activations to behavioral variations within the vignette's interaction sequence. For example, how would a shopkeeper highly motivated by concerns for personal safety react to an information request compared to a shopkeeper highly motivated by material gain, and how could such personality differences be implemented in a systematic way?

Borrowed and scaled down from several VECTOR scenarios, this model is comprised of two agents only, a shopkeeper and a soldier. As in the original scenario, the setting is the shopkeeper's store. The storyline (Figure 8) is, for the reasons discussed, simple.

The soldier approaches the counter, his main intent to obtain a piece of information. The shopkeeper meets him at the counter, where they proceed to engage in small talk. In this example, small talk is restricted to a single exchange of greetings. Once this pleasantry has been accomplished, the soldier proceeds to request the information of interest, the content of which, for our current purpose, is inconsequential. At this time, based on his particular level of the relevant motives, the shopkeeper will make a choice: he will decide to provide the requested information or deny the soldier's request. This behavioral alternative represents the single point of divergence in the story that was the requirement for the model. In either case, the story resumes with closing courtesies similar to the initial greetings. The soldier thanks the shopkeeper, regardless of whether the information was provided. After the shopkeeper acknowledges thanks and the two exchange goodbyes, the soldier exits the shop, thereby ending the story. It is worth noting that the model could be easily elaborated to provide a branch in the behavior of not merely one, but both agents involved. For instance, based on the shopkeeper's decision to honor or deny the request for information, the soldier could have two resulting behavioral options, 'thank you' or 'thanks anyway' respectively. However, for the intent of this model, a branch point for a single agent sufficed.

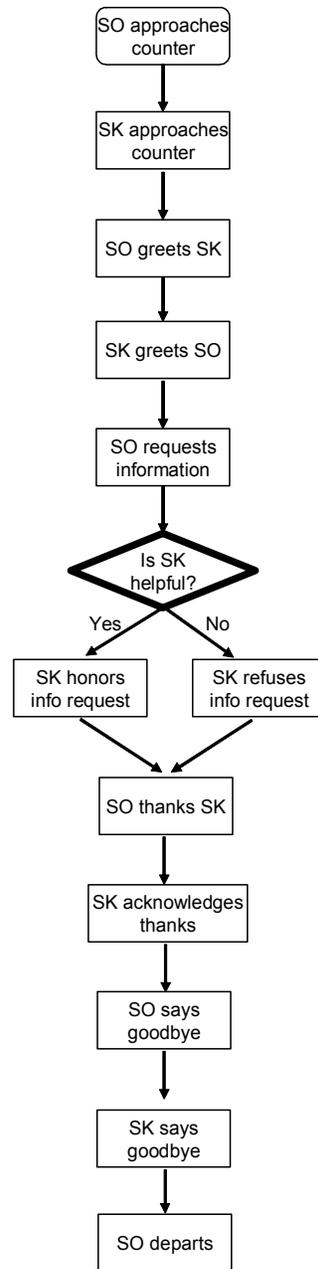


Figure 8. Basic Model Storyline

The basic model used a very simplified subset of motives, shown in Table 1. This subset was considered the bare requirement to meet the model’s objective, to create a single, motivationally-influenced behavioral choice.

Table 1. Basic Model Motive Structure

Motive Type	BIS or BAS
Avoid Harm	BIS
Gain	BAS
Helpful	BAS
Knowledge	BAS
Soc stimulation	BAS

The majority of the motive types involved are those regarded as belonging to the Behavioral Approach System. Only the motive *Avoid Harm* is designated as part of the Behavioral Inhibition System. The most noteworthy motives involved are *Helpful* and *Knowledge*. *Helpful* was used to characterize the shopkeeper’s desire to accommodate the soldier by providing the information he in turn desires. Conversely, *Knowledge* was used to represent the soldier’s desire to gain this information. The remaining motive types were added essentially for completeness. In fact, for the model to serve accurately as a true base-case for evaluating the architecture, it was important that it included other types of motivations aside from the aforementioned central ones. As discussed in previous sections, the current implementation designates the two motives with the highest level of activation as ‘active’ motives, and it is only these that are permitted to affect narrative and behavioral selection. Therefore, it was crucial the model contained other motives that were allowed to compete with the *Helpful* and *Knowledge* motives.

As the table indicates, *Helpful* is the motive type of key significance to this model. It is the shopkeeper’s activation level of this motive that determines whether he chooses to provide or deny the information the soldier has requested. As a result, altering his initial value of the *Helpful* motive causes the simple split in possible behavior we desired. The model as it exists currently has the starting value set to 0.5. Lowering this value results in the shopkeeper denying the soldier’s request; raising the value results in the shopkeeper honoring the soldier’s request. Since this motive is marked as part of the Behavioral Approach System, merely increasing the shopkeeper’s BAS activation level will also increase his propensity to provide the information. Furthermore, as the other competing motives of the model are marked as part of the Behavioral Inhibition System, one can also increase his BIS activation level, causing him to be more likely to deny information. The reverse holds true as well for decreasing both the BIS and BAS levels. The degree to which these values must be changed to cause the respective behavioral change in the shopkeeper depends on the specifics of the model’s configuration. It is these simple effects that have allowed us to constantly evaluate the core mechanisms of the architecture as we iteratively improved certain aspects.

Even within the scope of this simple example, the value and power of PAC can be seen. The idea of simulation/game-based training systems like VECTOR is to give the trainee an immersive and engaging opportunity to practice interactions like the one described above. In the real world, such encounters are repeated many times daily, and each can be a life-or-death

situation. Understanding how to behave in such a way as to obtain mission goals and avoid violence is key, and such understanding is best gained in a realistic but benign (simulation) environment. However, *realistic* is the operative word here. If the vignette is repeated dozens of times on many practice missions in different simulated villages, the shopkeepers cannot all act exactly the same. Each needs to reflect a different personality and set of past experiences (both of which will be unknown to the soldier on patrol, in life and in the simulation). These individual differences will generate different story evolutions that can lead to different outcomes. Only by providing the training simulation with such unpredictable and varying characters can the trainee get the benefit of effective training, in which he or she must be able to understand, predict, and adapt to the characteristics of the persons they are interacting with. This need is as great or even greater when providing training for collaborative tasks in joint/coalition command centers. There, although the danger to life, limb, and mission is not present in the direct interaction, the failure to adapt to the culture and personality of the other teammates can lead to potentially disastrous errors of omission or commission for the forces being commanded and controlled.

5. CONCLUSIONS AND FUTURE DIRECTION

The PAC system is novel and unique in its use of hierarchically organized, biologically based personal motive systems as a means for predictively simulating individual personality and cultural differences in the macro-organization of behavior. The two-layer organization of PAC allows it to leverage the past investment in existing cognitive architectures (e.g., COGNET/iGEN[®], ACT-R, and SOAR) by allowing them to function as the constrained rationality subsystems for the PAC personality subsystem. We believe that PAC will be able to capture most key aspects of culture and personality, by manipulating both the underlying motivate parameters of characters/agents and the story structures they used to interpret and generate behaviors. To the degree that personality and affect are culturally constrained, the PAC framework will be able to generate HBRs that exhibit culture-specific personality variability.

Future directions with PAC are proceeding in two related directions. First, continued elaboration of the architecture and associated toolset is on-going. An implementation of key emotions is underway, based on the theories of emotional appraisal (e.g. Ortony, A., Clore, G. L., & Collins, A., 1988; Lazarus, R. S., 2001). The emotion extensions will allow emotions (as well as motives) to affect the selection of Behavioral Options in a story. Additional extensions of the development tools are also underway, as is the generation of an open interface between PAC and existing cognitive architectures.

The second area of future development is the creation of C2 training simulations and games that will help C2 personnel develop increased collaboration and cooperation skills as are needed for effective net-centric operations involving diverse units from different cultural backgrounds and possessing widely different personalities.

REFERENCES

Cacioppo, J. T., Gardner, W. L., & Berntson, G. G. (1999). The affect system has parallel and integrative processing components: Form follows function. *Jrnl of Personality and Social Psych.*, 76, 839-855.

- Costa, P. T., Jr., & McCrae, R. R. (1992). *Revised NEO Personality Inventory (Neo-PI-R) and NEO Five-Factor Inventory (NEO-FFI): Professional manual*. Odessa, FL: Psychological Assessment Resources.
- Davidson, R. J., & Irwin, W. (1999). The functional neuroanatomy of emotion and affective style. *Trends in Cognitive Sciences*, 3, 11-21.
- Davidson, R. J., Jackson, D. C., & Kalin, N. H. (2000). Emotion, Plasticity, Context, and Regulation: Perspectives from Affective Neuroscience. *Psychological Bulletin*, 126, 890-909.
- Depue, R. A. (1996). A neurobiological framework for the structure of personality and emotion: Implications for personality disorders. In J. Clarkin & M. Lenzenweger (Eds.) *Major theories of personality disorders* (pp. 347-390). New York: Guilford Press.
- Gray, J. A. (1987). *The psychology of fear and stress*. (2nd Ed). New York: Cambridge University Press.
- Lazarus, R. S. (2001). Relational meaning and discrete emotions. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, methods, research* (pp. 37-67). Oxford: Oxford University Press.
- McCrae, R. R., & John, O. P. (1992). An introduction to the Five-Factor Model and its applications. *Journal of Personality*, 60, 175-215.
- Nigg, J. T. (2000). On inhibition/disinhibition in developmental psychopathology: Views from cognitive and personality psychology and a working inhibition taxonomy. *Psych. Bulletin*, 126, 220-246.
- Ortony, A., Clore, G. L., & Collins, A. (1988). *The Cognitive Structure of Emotions*. New York: Cambridge University Press.
- Panksepp, J. (1998). *Affective Neuroscience: The Foundations of Human and Animal Emotions*. New York: Oxford University Press.
- Panksepp, J. (2000). Emotions as natural kinds in the mammalian brain. In M. Lewis & J. M. Haviland-Jones (Ed.), *Handbook of Emotions (2nd edition)*. (pp. 137-156). New York: Guilford Press.
- Pew, R. & Mavor, A. (eds.) (1998). *Modeling Human and Organizational Behavior: Application to Military Simulations*. Washington, DC: National Academy Press.
- Read, S. J., & Miller, L. C. (1998). On the dynamic construction of meaning: An interactive activation and competition model of social perception. In S. J. Read & L. C. Miller (Eds.) *Connectionist models of social reasoning and behavior*. Mahwah, NJ: Erlbaum.
- Read, S. J., Jones, D. K., & Miller, L. C. (1990). Traits as goal-based categories: The role of goals in the coherence of dispositional categories. *Jrnl of Personality and Social Psych.*, 58, 1048-1061.
- Saucier, G., & Goldberg, L. R. (1996a). Evidence for the Big Five in analyses of familiar English personality adjectives. *European Journal of Personality*, 10, 61-77.
- Tellegen, A. (1985). Structures of mood and personality and their relevance to assessing anxiety, with an emphasis on self-report. In A. H. Tuma & J. D. Maser (Eds.), *Anxiety and the anxiety disorders*. (pp. 681-706). Hillsdale, NJ: Erlbaum.

- Tellegen, A., & Waller, N. G. (1997). Exploring personality through test construction: Development of the Multidimensional Personality Questionnaire. In S. R. Briggs, J. M. Cheek, & E. M. Donohue (Eds.), *Handbook of adult personality inventories*. New York: Plenum Press.
- Watson, D., & Clark, L. A. (1993). Behavioral disinhibition versus constraint: A dispositional perspective. In D. M. Wegner & J. W. Pennebaker (Eds.), *Handbook of mental control*. (pp. 506-527). Upper Saddle River, NJ: Prentice Hall.
- Wiggins, J. S., & Trapnell, P. D. (1996). A dyadic-interactional perspective on the five-factor model. In J. S. Wiggins (Ed.), *The five-factor model of personality: Theoretical perspectives* (pp. 88-162). New York: Guilford Press.