Increasing SIA Architecture Realism by Modeling and Adapting to Affect and Personality

Eva Hudlicka
Psychometrix Associates, Inc.
1805 Azalea Drive, Blacksburg, VA 24060
hudlicka@acm.org

Abstract

A key aspect of social interaction is the ability to exhibit and recognize variations in behavior due to different affective states and personalities. To enhance their believability and realism, socially intelligent agent architectures must be capable of modeling and generating behavior variations due to distinct affective states and personality traits on the one hand, and to recognize and adapt to such variations in the human user / collaborator on the other. In this paper we describe an adaptive user interface capable of recognizing and adapting to the user's affective and belief state (e.g., heightened level of anxiety). The Affect and Belief Adaptive Interface System (ABAIS) is designed to compensate for performance biases caused by users' affective states and active beliefs. The performance bias prediction is based on empirical findings from emotion research, and knowledge of specific task requirements. The ABAIS architecture implements an adaptive methodology consisting of four steps: sensing/inferring user affective state and performance-relevant beliefs; identifying their potential impact on performance; selecting a compensatory strategy; and implementing this strategy in terms of specific GUI adaptations. ABAIS provides a generic adaptive framework for exploring a variety of user affect assessment methods (e.g., knowledge-based, self-reports, diagnostic tasks, physiological sensing), and GUI adaptation strategies (e.g., content- and format-based). An ABAIS prototype was implemented and demonstrated in the context of an Air Force combat task, using a knowledge-based approach to assess and adapt to the pilot's anxiety level.

Introduction

A key aspect of human-human social interaction is the ability to exhibit and recognize variations in behavior due to different affective states and personalities. These subtle, often non-verbal, behavioral variations communicate critical information necessary for effective social interaction and collaboration. To enhance their believability and realism, socially intelligent agent architectures must be capable of modeling and generating behavior variations due to distinct affective states and personality traits on the one hand, and to recognize and adapt to such variations in the human user / collaborator on the other

We have been pursuing these goals along two lines of research: 1) developing a cognitive architecture capable of modeling a variety of individual differences (e.g., affective states, personality traits, etc.) (Hudlicka, 1998; Hudlicka & Billingsley, 1999b; Hudlicka & Billingsley, 1999c), and 2) developing an adaptive user interface capable of recognizing and adapting to the user's affective and belief state (e.g., heightened level of anxiety, belief in imminent threat, etc.) (Hudlicka & Billingsley, 1999a).

In this paper we focus on the area of affective adaptation and describe an Affect and Belief Adaptive Interface System (ABAIS) designed to compensate for performance biases caused by users' affective states and active beliefs. The performance bias prediction is based on empirical findings from emotion research, and knowledge of specific task requirements. The ABAIS architecture implements an adaptive methodology consisting of four steps: sensing/inferring user affective state and performance-relevant beliefs; identifying their potential impact on performance; selecting a compensatory strategy; and implementing this strategy in terms of specific GUI adaptations. ABAIS provides a generic adaptive framework for exploring a variety of user assessment methods (e.g., knowledge-based, self-reports, diagnostic tasks, physiological sensing), and GUI adaptation strategies (e.g., content- and format-based). We outline key research questions, the motivating psychological theory and empirical data, and present preliminary results from an initial prototype implementation in the context of an Air Force combat task. We conclude with a summary and outline of future research and potential applications for the synergistic application of the affect-adaptive and affect and personality modeling methodologies within SIA architectures.

Key Research Questions and Issues

A number of research and design questions arise in modeling of, and adaptation to, affect and personality factors in. These include the following:

- What are the most critical affective and personality factors affecting behavior? What are the effects of these factors on internal information processing and how are these effect manifested externally?
- How can the influences of these factors be best represented within agent cognitive architectures? Which components, and which internal structures and processes, are necessary to model the key affective and personality factors? At what level of resolution must cognitive apparatus be represented to afford the modeling of these factors?

- How can we predict the effects of these factors on behavior for a particular individual within a particular social and task context?
- What are the best strategies for adapting to (compensating for or enhancing) a particular affective state or personality trait?

Selecting Affective States and Personality Traits

The first step for both the modeling and the adaptation research goals was to review existing empirical psychological literature to identify key affective and personality factors influencing behavior. The *affective states* studied most extensively include anxiety, positive and negative mood, and anger. The effects of these states on behavior range from influences on distinct information processes within the cognitive architecture (e.g., attention and working memory capacity, accuracy, and speed; memory recall biases), through autonomic nervous system manifestations (e.g., heart rate, GSR), to visible behavior (e.g., facial expressions, approach vs. avoidance tendencies, aggressive behavior, etc.) (LeDoux, 1898; Williams et al., 1997; Mineka and Sutton, 1992; MacLeod and Hagan, 1992). A wide variety of *personality traits* have been studied, ranging from general, abstract behavioral tendencies (e.g., "Big 5" (Extraversion, Emotional Stability, Agreeableness, Openness, Conscientiousness), and "Giant 3" (Approach behaviors, Inhibition behaviors, Aggressiveness) personality factors), through psychodynamic / clinical personality formulations (e.g., narcissistic, passive-aggressive, avoidant, etc.), to characteristics relevant for particular type of interaction (e.g., style of leadership, preferred style of social interaction, decision making, etc.) (Revelle, 1995). Our initial primary focus in both the modeling and the adaptation research areas was on anxiety, aggressiveness, and obsessiveness.

Recognizing and Adapting to Human Affect and Personality

Methodology

We have developed a methodology designed to compensate for performance biases caused by users' affective states and active beliefs: the Affect and Belief Adaptive Interface System (ABAIS) (Hudlicka & Billingsley, 1999a). The ABAIS methodology consists of four steps: 1) sensing/inferring the individual's affective state and performance-relevant beliefs (e.g., high level of anxiety; aircraft is under attack); 2) identifying their potential impact on performance (e.g., focus on threatening stimuli, biasing perception towards identification of ambiguous stimuli as threats); 3) selecting a compensatory strategy (e.g., redirecting focus to other salient cues, presentation of additional information to reduce ambiguity); and 4) implementing this strategy in terms of specific UI adaptations (e.g., highlighting relevant cues or displays); that is, presenting additional information, or presenting existing information in a format that facilitates recognition and assimilation, thereby enhancing situation awareness.

ABAIS Architecture

The ABAIS architecture consists of four modules, each implementing the corresponding step of the adaptive methodology described above (see figure 1): User State Assessment, Impact Prediction, Strategy Selection, and GUI Adaptation. These four modules are briefly described below.

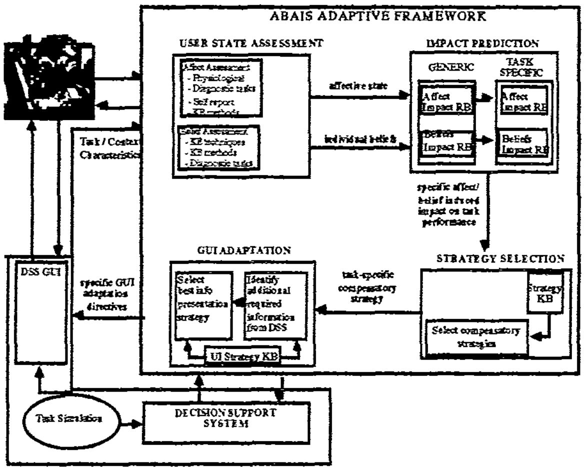

Figure 1: ABAIS Architecture Implementing the Affect-Adaptive Methodology

User State Assessment Module This module receives a variety of data about the user and the task context, and from these data identifies the user's predominant affective state (e.g., high level of anxiety) and situation-relevant beliefs (e.g., interpretation of ambiguous radar return as threat), and their potential influence on task performance (e.g., firing a missile at ambiguous return). Since no single reliable method currently exists for affective assessment, the User Assessment module provides facilities for the flexible combination of multiple methods or data from such methods. These include: physiological assessment (e.g., heart rate); diagnostic tasks; self-reports; and use of knowledge-based methods to derive likely affective state based on factors from current task context (e.g., type, complexity, time of day, length of task), personality (negative emotionality, aggressiveness, obsessiveness, etc.), and individual history (past failures and successes, affective state associated with current task, etc.)

The user assessment module uses knowledge from each of these categories of factors to derive the user's affective state. The assessment process implements a fuzzy rule-based approach consisting of four stages. *First,* a user profile is specified in terms of static and dynamic data, representing task-relevant factors about the user. Examples of static data are personality traits, individual history, and training and proficiency. Examples of dynamic data are physiological data and dynamic task factors. *Second,* the data in this profile are matched against the rules in the user assessment rule-base. *Third,* each relevant factor, represented by an instantiated rule, contributes a numerical weight component to the overall score of the corresponding affect. Individual factors or categories of factors may be weighted differently, to reflect their differential influence on the overall affective state (e.g., static task factors will typically have a lower weight than dynamic factors and real-time physiological signals). *Finally,* after all relevant rules are instantiated, the overall anxiety level is computed and the resulting value is mapped onto a three-valued qualitative variable indicating a low, medium, or high anxiety level.

For the preliminary ABAIS prototype, we focused on a knowledge-based assessment approach, applied to assessment of anxiety levels, to demonstrate the feasibility of the overall adaptive methodology. The knowledge-based assessment approach assumes the existence of multiple types of data (e.g., individual history, personality, task context, physiological signals), and from these data derives the likely anxiety level. Anxiety was selected both because it is the most prevalent affect during crisis situations, and because its influence on cognition has been extensively studied and empirical data exist to support specific impact prediction and adaptation strategies.

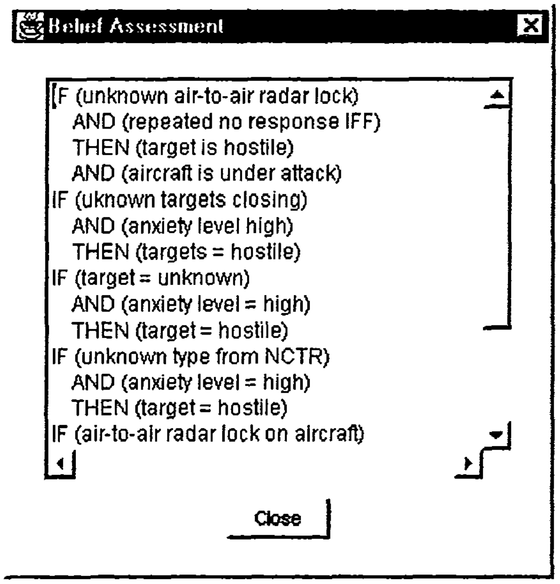

Figure 2: Examples of Belief Assessment Rules from ABAIS Rule Base

The pilot's affective state plays a critical role in his/her belief and situation assessment. By taking into account the current affective state, the ABAIS User Assessment module in effect implicitly models the potential biasing influences of the different affective states and provides a structure which allows the explicit representation of the positive feedback between cognition and affect that is often seen in crisis situations. In other words, increased anxiety contributes to a particular situation assessment (e.g., aircraft is being attacked by hostile aircraft), which then limits the processing of data that could give rise to alternative interpretations and further increases the anxiety level. Examples of rules for belief assessment are shown in figure 2. These rules map the combinations of cues representing external events, individual history, and affective state onto the set of possible situations.

Impact Prediction Module This module receives as input the identified affective states and associated task-relevant beliefs, and determines their most likely influence on task performance. The goal of the impact prediction module is to predict the influence of a particular affective state (e.g., high anxiety) or belief state (e.g., "aircraft under attack", "hostile aircraft approaching", etc.) on task performance. Impact prediction process uses rule-based reasoning (RBR) and takes place in two stages. First, the *generic effects* of the identified affective state are identified, using a knowledge-base that encodes empirical evidence about the influence of specific affective states on cognition and performance. Next, these generic effects are instantiated in the context of the current task to identify *task-specific effects*, in terms of relevant domain entities and procedures (e.g., task prioritization, threat assessment). The knowledge encoded in these rules is derived from a detailed cognitive affective personality task analysis (CAPTA), which predicts the effects of different affective states and personality traits on performance in the current task context. The CAPTA process is described in detail in Hudlicka (2000). The separation of the generic and specific knowledge enhances modularity and simplifies knowledge-based adjustments.

Strategy Selection Module This module receives as input the predicted specific effects of the affective and belief states, and selects a compensatory strategy to counteract resulting performance biases. Strategy selection is accomplished by rule-based reasoning, where the rules map specific performance biases identified by the Impact Prediction Module (e.g., task neglect, threat-estimation bias, failure-estimation bias, etc.) onto the associated compensatory strategies (e.g., present reminders of neglected tasks, present broader evidence to counteract threat-estimation bias, present contrary evidence to counteract failure-driven confirmation bias, etc.). As was the case with impact prediction, the strategy selection module relies on a detailed analysis of the task context via the CAPTA process, which identifies specific strategies available to counteract the possible biases. This analysis then allows the construction of the strategy selection knowledge bases. Table 1 shows examples of task-specific rules for compensatory strategy selection.

Anxiety effects
IF (recent change in radar target status) THEN (emphasize change in status of return)

IF (attention focus = HUD) AND (incoming radar data) THEN (redirect focus to radar)

IF (attention focus = radar) AND (Incoming radio call) THEN (redirect focus to radio)

IF (attention focus = non-radar instruments) AND (incoming radar data) THEN (redirect focus to radar)

IF (likelihood of task neglect for <instrument> = high) & (has-critical-info? <instrument>) THEN (emphasize <instrument> visibility)

IF (target = unknown) AND (target belief = hostile) THEN (emphasize unknown status) AND (collect more data)

Aggressiveness effects
IF (likelihood of premature attack = high) THEN (display all available info about enemy a/c) AND (enhance display of enemy a/c info)

Obsessivness effects
IF (likelihood of delayed attack = high) THEN (display all available info about enemy a/c) AND (display likelihood of attack by enemy a/c) AND (display envelope of vulnerability around own aircraft) AND (display reminders for attack tasks)

Table 1: Examples of Task-Specific Rules for Compensatory Strategy Selection

GUI Adaptation Module This module performs the final step of the adaptive methodology, by implementing the selected compensatory strategy in terms of specific GUI modifications. A rule-based approach is used to encode the knowledge required to map the specific compensatory strategies onto the necessary GUI / DSS adaptations. The specific GUI modifications take into consideration information about the individual pilot preferences for information presentation, encoded in customized user preference profiles; for example, highlighting preferences might include blinking vs. color change vs. size change of the relevant display or icon.

In general, two broad categories of adaptation are possible: content-based, which *provide additional information,* and format-based, which *modify the format of existing information. Content-based adaptation* involves the collection and display of additional data or knowledge to compensate for a particular performance bias. For example, providing additional data about an ambiguous radar signal helps prevent an anxiety-induced bias to identify ambiguous signals as threatening. *Format-based adaptation* involves the presentation of existing data in an alternative format, to enhance visibility and / or to draw attention to neglected displays, and, in general, to facilitate fast detection, recognition, and assimilation of data, thereby improving situation awareness. Figure 3 illustrates examples of ABAIS GUI adaptation in terms of specific cockpit display modifications.

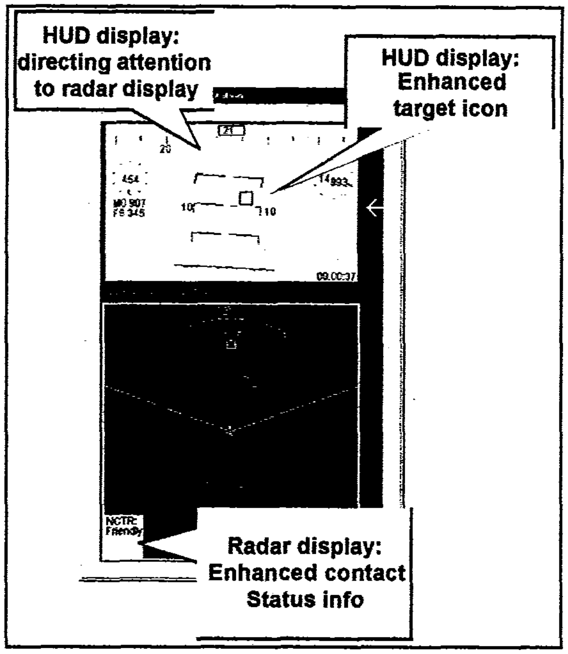

Figure 3: Examples of Generic Classes of GUI Adaptations

Results

The ABAIS prototype was implemented and demonstrated in the context of an Air Force combat mission, used a knowledge-based approach to assess the pilot's anxiety level, and modified selected cockpit instrument displays in response to detected increases in anxiety levels.

Several representative pilot profiles were defined, varying in personality, physiological responsiveness, training, individual history, and adaptation preferences, making it more or less likely that the pilot would reach a specific level of anxiety during the mission.

Once an increased level of anxiety was observed, ABAIS predicted that the heightened level of anxiety would cause narrowing of attention, an increased focus on potentially threatening stimuli, an a perceptual bias to interpret ambiguous radar signals as threats, thus risking fratricide. ABAIS therefore suggested a compensatory strategy aimed at: 1) directing the pilot's attention to a cockpit display showing the recent status change; and 2) enhancing the relevant signals on the radar to improve detection. Figure 4 illustrates these adaptations. Specifically, the blinking, enlarged, blue contact icon on the HUD display indicates a change in status. The blinking blue "RADAR" string displayed on the HUD, the pilot's current focus, directs the pilot to look at the radar display, which shows an enhanced contact icon indicating a change in status, with details provided in the text box in lower left corner of the display.

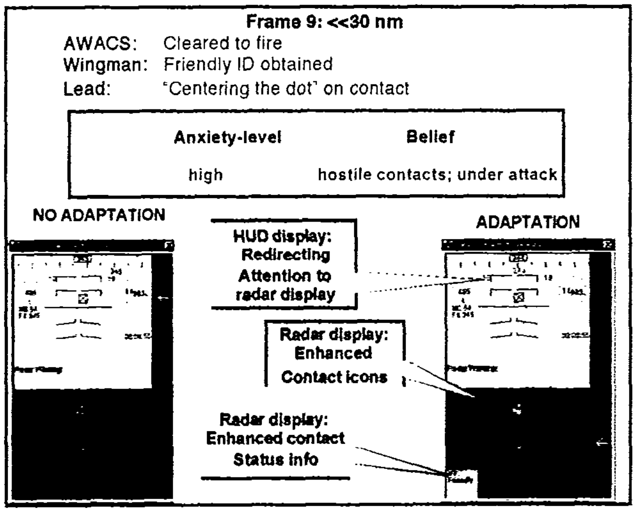

Figure 4: Example of Specific Scenario Adaptation Sequence

Conclusions

We described a research area aimed at producing more realistic behavior in socially intelligent agents, namely the *recognition of and adaptation to a human's affective state*. We developed an adaptive methodology and demonstrated its effectiveness by implementing a prototype Affect and Belief Adaptive Interface System (ABAIS). ABAIS was demonstrated in the context of a simulated pilot task.

ABAIS assessed the user state using a knowledge-based approach and information from a variety of sources (e.g., static task context, dynamic external events occurring during the scenario, pilot's individual history, personality, training and proficiency, and simulated physiological data), predicted the effects of this state within the constrained context of the demonstration task, and suggested and implemented specific GUI adaptation strategies based on the pilot's individual information presentation preferences (e.g., modified an icon or display to capture attention and enhance visibility). The preliminary results indicate the general feasibility of the approach, raise a number of further research questions, and provide information about the specific requirements for a successful, operational affective adaptive interface. Although the initial prototype was developed within a military domain, we believe that the results are applicable to a broad variety of non-military application areas, as outlined below.

Requirements for Adaptation A number of requirements were identified as necessary for affective adaptive interface system implementation. These are listed below:

- Limiting the number, type, and resolution of affective states (e.g., distinguishing between high vs. low anxiety); using multiple, complementary methods and multiple data sources for affective state assessment;
- Providing individualized user data, including details of past performance, individual history, personality traits, and physiological data;
- Constraining the overall situation in terms of situation assessment and behavioral possibilities;
- Providing a wide variety of task-specific data in an electronic format;
- Fine-tuning the rule-bases and inferencing to "personalize" the system to the individual user-task context; and
- Implementing 'benign' adaptations, that is, GUI / DSS modifications that at best enhance and at worst maintain current level of performance (e.g., adaptations should never limit access to existing information).

Key Issues to Address A number of issues must be addressed to further validate this approach and to provide robust affect adaptive systems. These include the following:

- Empirical study demonstrating improved human-machine performance with adaptation
- Demonstration of effectiveness of the ABAIS methodology across multiple task contexts
- Implementation, integration, and evaluation of multiple affect- and belief-assessment methods
- Effectiveness of physiological affect assessment.
- Implementation of non-intrusive physiological assessment methods.
- Effectiveness of knowledge-elicitation and task analysis via the Cognitive Affective Personality Tasks Analysis process as basis for belief assessment across multiple task contexts.
- Feasibility of predicting effects of user state on performance across multiple task contexts.

Future Work

Possible future work in the broad area of user affect and modeling is limitless at this point, as the endeavor is in its infancy. Key questions include issues such as:

- What emotions *should* and *can* be recognized, modeled, and adapted to in human-machine interfaces.
- At what level of abstraction should affect and belief be modeled for different applications.
- When should an agent attempt to enhance the user's affective state, when should it adapt to the user's affective state, and when should it attempt to counteract it.

Canamero offers an excellent summary of some of the affect-related issues that must be addressed by the SIA architecture research community (Canamero, 1998).

Individually, both modeling and recognition of affective states and personality traits provide a powerful enhancement to agent architectures, by enabling socially intelligent, adaptive behavior. The coordinated integration of these two enhancements within a single agent architecture promises even further benefits, by enhancing the realism and effectiveness of human-machine interaction across a variety of application areas. Below we briefly outline applications in three existing and emerging areas: 1) *Education, training,* and infotainment and edutainment industry; 2*) Virtual reality treatment environments* for a variety of affective, cognitive, and personality disabilities and disorders; and 3) *Decision-aids* in real-time, crisis-prone, high-risk decision-making environments.

Infotainment and Edutainment Industry Education and training systems have traditionally focused on cognitive factors. This is in spite of increasing evidence and individual styles, personalities, and affective state *during* learning greatly influences the training outcome. This area thus provides multiple opportunities for agents capable of modeling and adaptation to the trainee's current affective state, by changing the training protocol, altering the mode of information presentation, etc.

Clinical Treatment Settings A key area of applicability for the technologies described in this paper are clinical settings for the treatment of a variety of affect- and personality-induced behavioral problems (e.g., anxiety disorders, social phobias, specialized group therapies, depressive disorders, etc.). Some of these disabilities require lengthy therapeutic interventions where compliance is often difficult due to the intensity of the involvement required by professional personnel and the extensive time required to correct the disorder. Virtual reality training environments incorporating the agent architectures described above have the potential to greatly enhance existing face-to-face treatment modalities.

Decision-Support Systems DSS are increasingly required in a variety of *real-time contexts* subject to 'crisis' situations (e.g., air traffic control, emergency medical operations, disaster management operations, firefighting, etc.). The ability of the system to recognize the user's increased level of anxiety or specific personality traits, and adapt the decision aiding and user interface accordingly, would greatly enhance the effectiveness of these systems. In addition, *virtual reality training environments* designed to identify specific behaviors or situations subject to negative affective state and personality influences, would provide training to develop appropriate compensatory strategies, and *screening environments* to identify potential affect-induced performance biases across a variety of decision-making contexts. *Usability and design evaluation testbed* could be constructed to identify performance variations in specific situations, and the effectiveness of existing systems in accommodating these variations. *Collaborative, distributed decision making environments* where multiple individuals must effectively interact, would benefit from training environments capable of simulating particular personality types, and their effect on the team process and outcome.

Acknowledgements

The research described in this paper was supported in part by US Air Force Contract F41624-98-C5032. We would like to acknowledge the contributions of the following individuals to this project: the technical monitor Dr. Michael McNeese of Air Force Research Laboratory, Mr. John Billingsley of Psychometrix Associates, Inc., Mr. Bob Shaw of NTI, Inc. Dr. Thomas von Wiegand and Prof. Rosalind Picard of MIT, and Dr. Greg Zacharias of Charles River Analytics, Inc.

References

Canamero, D. (1998). Issues in the Design of Emotional Agents. In *Proceedings of "Emotional and Intelligent: The Tangled Knot of Cognition".* AAAI Fall Symposium Series, TR FS-98-03. Menlo Park, CA: AAAI Press., 1998.

Hudlicka, E. (1998). Modeling Emotion in Symbolic Cognitive Architectures. In *Proceedings of "Emotional and Intelligent: The Tangled Knot of Cognition".* AAAI Fall Symposium Series, TR FS-98-03. Menlo Park, CA: AAAI Press., 1998.

Hudlicka, E. and Billingsley, J. (1999a). *ABAIS: Affect and Belief Adaptive Interface System*. Report AFRL-HE-WP-TR-1999-0169. WP AFB, OH: US Air Force Research Laboratory.

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

Hudlicka, E. and Billingsley, J. (1999c*). MAMID: Methodology for Analysis and Modeling of Individual Differences*. Report 9905. Lincoln, MA: Psychometrix Associates, Inc,.

Hudlicka, E. (2000). *Cognitive Affective Personality Task Analysis.* Technical Report 2001. Lincoln, MA: Psychometrix Associates, Inc.

LeDoux, J.E. (1989). Cognitive-Emotional Interactions in the Brain. *Cognition and Emotion*, 3(4), 267-289.

MacLeod, C. and Hagan, R. (1992). Individual differences in the selective processing of threatening information. Behav. Res. Ther., 30, 151-161.

Mineka, S. & Sutton, S.K. (1992). Cognitive Biases and Emotional Disorders. Psychological Science, 3(1), 65-69.

Revelle, W. (1995). Personality Processes. Annual Review of Psychology, 46, 295-328.

Williams, J.M.G., Watts, F.N., MacLeod, C., and Mathews, A. (1997). Cognitive Psychology and Emotional Disorders. NY: John Wiley.