Modeling the Mechanisms of Emotion Effects on Cognition

Eva Hudlicka

Psychometrix Associates, Inc.
1805 Azalea Drive, Blacksburg, VA , US
hudlicka@ieee.org

Abstract

Emotions exert a profound influence on cognitive processes, both the fundamental processes mediating cognition, such as attention and memory, and higher-level processes including decision-making and learning. A number of emotion effects on cognition have been identified, but their mechanisms are not yet understood. In this paper I describe a methodology for modeling the effects of emotion on cognition, within a symbolic cognitive-affective architecture. The primary objective of the approach is to facilitate the construction of alternative mechanisms of observed emotion effects. The paper describes how the effects of anxiety are modeled and how alternative mechanisms of these effects can be explored.

Introduction

Emotions exert profound influences on cognition in biological agents. This is particularly evident in decision-making. All of the processes mediating decision-making are affected by emotion: attention, perception, situation assessment, goal-management and action selection, as well as the associated memory processes.

Emotion effects, and the associated affective decision biases and heuristics, can be adaptive or maladaptive, depending on their type, magnitude and context. For example, anxiety and fear are associated with preferential processing of high-threat stimuli. This is highly adaptive in situations where survival depends on quick detection of danger and appropriate reaction (e.g., avoid an approaching car that has swerved into your lane). The same bias can be maladaptive if neutral stimuli are judged to be threatening (e.g., a passing car is assumed to be on a collision course and causes you to swerve into a ditch.)

A range of emotion effects and biases has been identified: positive emotions induce a global focus and the use of heuristics, whereas negative emotions induce a more local focus and analytical thinking; anxiety reduces attentional and working memory capacities, biases attention towards the detection of threatening stimuli, and biases interpretive processes towards higher threat assessments, anxiety also induces a self-bias; mood induces mood-congruent biases in recall (Isen, 1993; Mineka et al., 2003).

However, the mechanisms mediating these effects are not yet understood. A better understanding of these mechanisms would enable us to make use of the adaptive aspects of emotion effects in agent architectures, and would enable the design of human-machine interfaces that could help counteract the maladaptive emotion effects.

This paper describes an approach to modeling emotion effects that focuses on identifying the associated alternative mechanisms. The approach is implemented within a symbolic cognitive-affective architecture, MAMID, that models high-level decision-making (Hudlicka, 2007; 2004; 1998). MAMID is a process-oriented model, focusing on emotion, and as such aims to explicitly represent the structures and processes mediating affective processing and emotion-cognition interactions.

Key aspects of MAMID that make it suitable for modeling alternative mechanisms of emotion effects are: (1)high degree of parameterization, enabling manipulation of architecture topology and data flow, and processing within the modules, and (2) testbed environment, within which MAMID is embedded. The testbed facilitates rapid model development and interactive model 'tuning', by providing the modeler access to a range of model parameters, and control of the functions that derive their values. By manipulating these parameters, alternative hypotheses regarding the mechanisms of an observed phenomenon can be rapidly implemented and their behavior evaluated within the context of a specific task.

MAMID is a domain-independent architecture, which has been evaluated in two domains: modeling members of a search-and-rescue team (Hudlicka, 2005), and modeling unit leaders in a peacekeeping scenario (Hudlicka, 2003; 2007). MAMID is currently undergoing validation in conjunction with a parallel set of empirical studies with human subjects, within the search-and-rescue task.

MAMID Cognitive-Affective Architecture

MAMID implements a sequential see-think-do processing sequence (figure 1), consisting of the following modules: *Sensory Pre-processing* (translates incoming data into task-relevant cues); *Attention (*filters incoming cues and selects a subset for processing); *Situation Assessment* (integrates individual cues into an overall situation assessment); *Expectation Generation* (projects current

Copyright © 2008. Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

situation onto possible future states); *Affect Appraiser (*derives a valence and four of the basic emotions from external and internal elicitors); *Goal Management* (identifies high-priority goals); and *Behavior Selection (*selects the best actions for goal achievement).

These modules map the incoming stimuli (cues) onto the outgoing behavior (actions), via a set of intermediate internal structures (situations, expectations, and goals), collectively termed *mental constructs*. This mapping is enabled by long-term memories (LTM) associated with each module, represented by belief nets. *Mental constructs* are characterized by their attributes (e.g., familiarity, novelty, salience, threat level, valence, etc.), which influence their processing; that is, their rank and the consequent likelihood of being processed by the associated module within a given execution cycle; (e.g., cue will be attended, situation derived, goal or action selected).

The availability of the mental constructs from previous execution cycles allows for dynamic feedback among constructs, and thus departs from a strictly sequential processing sequence. Note also that all constructs derived in a given execution cycle are available to subsequent modules for processing, within that cycle.

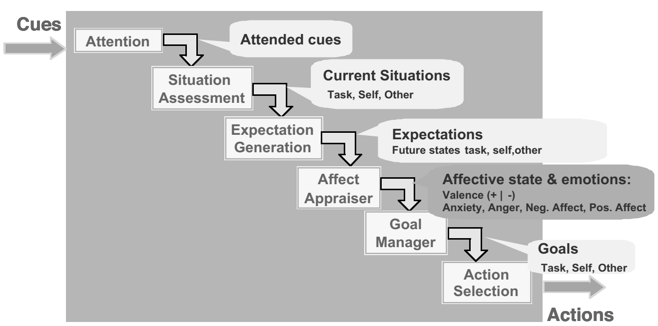

Figure 1: MAMID Cognitive-Affective Architecture

Modeling Affective Processes Within MAMID

MAMID's focus is the modeling of emotion. In an earlier paper, I have suggested that it is helpful to divide affective processes into two categories of core processes: those mediating the dynamic *generation of emotion* and those mediating the *effects of the resulting emotions*: both the visible manifestations and the internal effects on attention, perception and cognition (Hudlicka, 2008). This categorization facilitates the identification of the computational tasks necessary to implement the wide variety of observed affective phenomena, and provides a basis for a more systematic approach for the design and evaluation of computational models of emotion.

MAMID models both of these processes. Emotion generation is modeled via a dedicated *Affect Appraiser* module (see figure 2), which integrates external data (cues), internal interpretations (situations, expectation) and desires and priorities (goals), with stable and transient individual characteristics (traits and emotional states), and generates an affective appraisal at two levels of resolution: a *valence* (corresponding to an undifferentiated positive or negative evaluation, generally referred to as *affective state* in the psychological literature), and one of the four basic emotions (fear/anxiety, anger, sadness, joy).

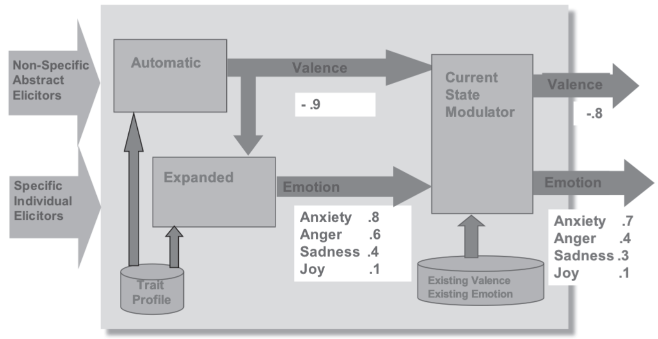

Figure 2: Structure of the Affect Appraiser Module

Generation of basic emotions represents more differentiated processing, where the intensity of each emotion is influenced by task- and individual-specific criteria. This involves a consideration of a variety of idiosyncratic criteria that determine, for example, whether a high-threat situation or en impending goal failure will lead to anger or anxiety, in a particular agent; e.g., a particular situation may affect anxiety positively or negatively, depending on the individual's history and experience. Such differentiated processing requires correspondingly complex inferencing and knowledge, which is implemented in MAMID in terms of belief nets.

Emotion intensities are determined from four contributing factors: *Trait bias factor* – reflecting a tendency towards a particular emotion, as a function of the agent's trait profile (e.g., high neuroticism/low extraversion individuals are predisposed toward negative emotions). *Valence factor* – reflecting a contribution of the current valence, where negative valence contributes to higher intensities of negative emotions and vice versa. *Static context factor* – reflecting the agent's skill level and contributing to the anxiety level if skill level is low. *Individual factor*- weighted sum of the emotion intensities derived from the emotion-specific belief nets, reflecting the idiosyncratic contributions of specific elicitors.

The Affect Appraiser module incorporates elements from several appraisal theories: *domain-independent appraisal dimensions, multiple-levels* of resolution, and *multiple stages* (Leventhal & Scherer, 1987; Smith & Kirby, 2000).

MAMID models emotion effects using a generic methodology for modeling multiple, interacting individual differences, both *stable traits and dynamic states* (Hudlicka, 2007; 2003; 1998). A particular configuration of emotion intensities and trait values is mapped onto a corresponding set of architecture parameter values, which then control processing within the architecture modules, as well as the data flow among the modules (figure 3).

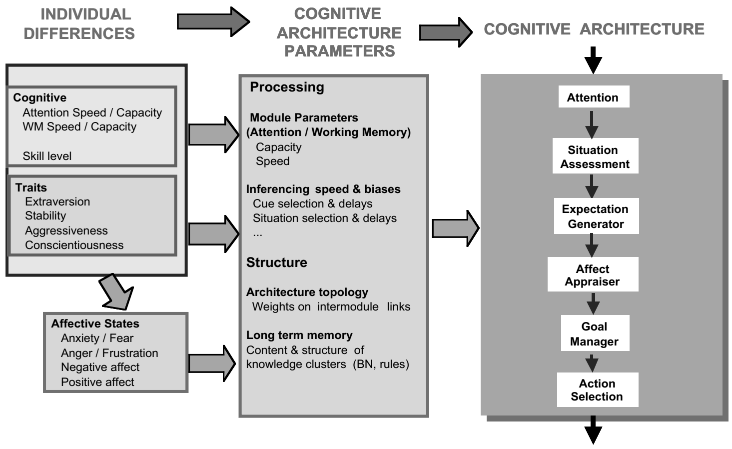

Figure 3: Schematic Illustration of MAMID Methodology for State and Trait Modeling

Functions implementing these mappings are constructed on the basis of the available empirical data. For example, reduced attentional and working memory (WM) capacities, associated with anxiety and fear, are modeled by dynamically reducing the available processing capacities of each MAMID module, which then reduces the number of constructs processed (fewer cues attended, situations derived, expectations generated, etc.). Currently, the parameter-calculating functions consist of weighted linear combinations of the factors that influence each parameter. For example, working memory capacity reflects a normalized weighted sum of emotion intensities, trait values, baseline capacity, and skill level.

Modeling Alternative Mechanisms of Emotion Effects

We illustrate the MAMID modeling methodology with an example demonstrating how multiple mechanisms of anxiety effects can be modeled within MAMID, for anxiety intensities ranging from low to extreme states, such as a panic attack. Anxiety was selected because of its direct relevance for decision-making and behavior, because robust empirical data regarding its effects are available, and because anxiety is emerging as the most significant effect in on-going empirical validation studies (Matthews, 2008).

Panic attack is an interesting state to explore because its extreme nature provides a useful context in which to model the effects of anxiety on cognition, and cognition-emotion interaction in general. Panic attack is a state where the confluence of multiple anxiety effects produces a type of a 'perfect storm', frequently inducing behavioral paralysis. Three anxiety-linked effects are involved: *threat processing bias, self processing bias*, and *capacity reductions in both attention and working memory.* MAMID models all three of these effects, and provides parameters that control their relative contributions to the overall effect on processing.

Recall that a given parameter value is the results of a linear combination of the weighted factors influencing the parameter. The same overall effect (e.g., reduced module capacity) can thus be obtained from multiple combinations of factor values and weights. These alternative configurations then provide the means of defining alternative mechanisms mediating specific effects. MAMID provides facilities that support the rapid construction of these alternative mechanisms, via interactive manipulation of the factors and weights, which allow the modeler to control the magnitude and contribution of each influencing factor.

Threat bias is modeled by first calculating the threat level of each cue, situation and expectation, from factors that include an a priori 'fixed' threat level (e.g., low level of resources is inherently more threatening than adequate resources), state and trait anxiety factors, and individual history. The threat level is then used as a weighted factor in the function calculating the overall construct rank, which determines the likelihood of its processing. In states of high-anxiety, high-threat constructs have a higher ranking, and are thus processed preferentially (cues attended, situations derived) (refer to figure 4).

Self bias is modeled by including a weighted factor reflecting the self vs. non-self origin of each construct in its rank calculating function. High levels of state or trait anxiety then induce a higher ranking for self-related constructs, contributing to their preferred processing.

The *capacity reduction* effects on attention and working memory are modeled by dynamically calculating the capacity values of all modules during each execution cycle, from weighted factors representing the emotion intensities, the four traits represented in MAMID, baseline capacity limits, and skill level.

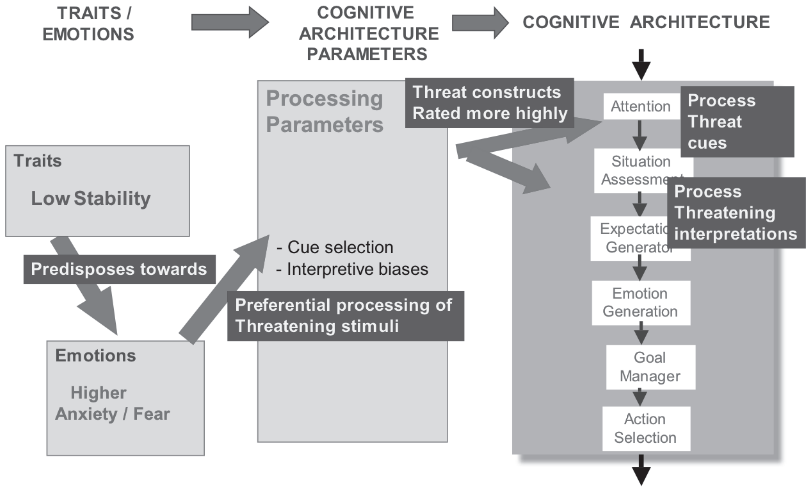

Figure 4: Modeling Threat Bias Within MAMID

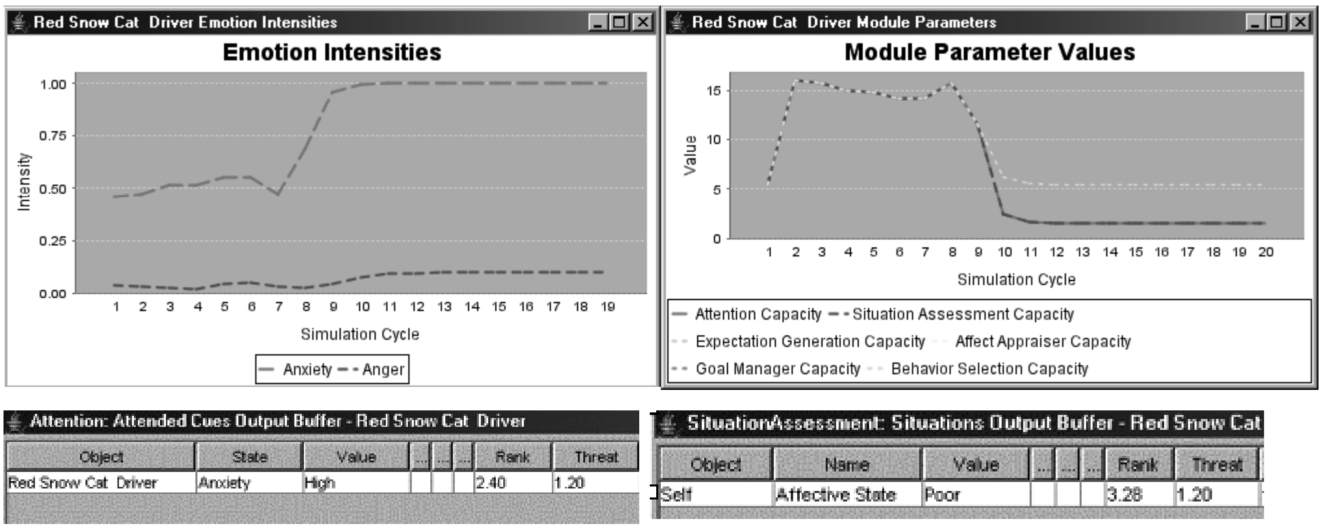

Figure 5: Affective Dynamics and Module Capacity Parameters (top) Associated with a "Panic Attack" State (Cycle 8), and Subsequent Reduction in the Number of Constructs Processed (bottom)

MAMID's ability to model alternative mechanisms of anxiety effects was demonstrated in the context of specific task, within the search-and-rescue domain. Briefly, the agent's task is to find a "lost party" in an inhospitable terrain, where "emergency situations" arise unexpectedly. The agent may need to obtain supplies from available "supply stations", to maintain adequate resources (fuel, first aid kits). In the experiment described below, the agent approaches a difficult "emergency situation", and lacks the required resources. The agent's state of anxiety, dynamically calculated by the Affect Appraiser module, is high; in part because of a trait-induced tendency towards higher anxiety, and in part because of the difficult task ahead and lack of adequate resources.

Within this context, MAMID models a panic attack state as follows. Stimuli, both external and internal, arrive at the Attention Module, whose capacity is reduced. Because of the threat- and self-bias, self-related high-threat cues are processed preferentially, in this case resulting in the agent's focus on a self-related anxiety cue (see figure 5, lower left). This cue, reflecting the agent's anxious state, consumes the limited module capacity, leading to the neglect of external and non-threatening cues (e.g., proximity of a supply station). This results in a continued self- and threat-focus in the downstream modules (Situation Assessment and Expectation Generation). No useful goals or behaviors can be derived from these constructs, and the agent enters a positive feedback-induced vicious cycle (an endless self-reflection), where the reduced-capacity and biased processing excludes cues that could lower the anxiety level and trigger adaptive behavior. Figure 5 shows a diagram of the emotion intensities and module capacities, and representative contents of the cue and situation buffers, providing input to Attention and Situation Assessment modules, respectively.

The model parameters are then modified to increase attentional and processing capacities, thereby enabling the processing of additional cues. This allows the agent to begin processing a larger set of incoming cues, which eventually result in a decreased state of anxiety, and trigger task-related goals and associated task-relevant behavior

A number of factors can be modified to induce the effects described above, simultaneously or sequentially, reflecting multiple, alternative mechanisms mediating the anxiety biasing effects. In the case of the capacity parameters, alternative mechanisms can be defined from the agent's overall sensitivity to anxiety (reflected in the weights associated with trait and state anxiety intensity factors), the baseline, 'innate' capacity limits (reflected in the factors representing the minimum and maximum attention and working memory capacities), and the anxiety intensity itself. This factor can be further manipulated via the set of parameters influencing the affect appraisal processes, including the nature of the affective dynamics (e.g., maximum intensity, and the intensity ramp-up and decay functions).

Related Work

A number of researchers have independently proposed a broader theory of mechanisms mediating emotion-cognition interaction, where parameters encoding various affective factors (states and traits), influence a broad range of cognitive processes and structures (e.g., (Hudlicka, 1998; Matthews & Harley, 1993; Ortony et al., 2005). These parameter-based theories are consistent with recent neuroscience theories, which suggest that emotion effects on cognition are implemented in terms of systemic, global effects on multiple brain structures, via distinct patterns of neuromodulation, corresponding to distinct emotions (Fellous, 2004). Several recent models of emotion effects use an approach that shares similarities with the MAMID methodology described here, that is, the use of processing parameters to encode emotion effects (Belavkin & Ritter, 2004; Ritter & Avraamides, 2000; Sehaba et al., 2007). The approaches vary in terms of specific parameters represented, the functions used to calculate their values, and whether or not the parameters correspond to psychological functions and processes, or serve a generic 'noise' parameters that globally degrade architecture performance to simulate states such as stress.

Conclusions

MAMID's abilities to model alternative mechanisms of emotion effects are continuing to be evaluated, with a focus on anxiety and anger. The on-going empirical study with human subjects is demonstrating anxiety biases in both information seeking and behavior selection, within the search-and-rescue task context. The results are being used to tune the MAMID parameters, as outlined above. Since the parameters correspond to specific psychological variables or functions, we hope that this parallel empirical-computational approach will provide a useful means for validating MAMID models of emotion effects on cognition.

References

Belavkin, R. V., & Ritter, F. E. (2004). OPTIMIST: A new conflict resolution algorithm for ACT-R. In *Proceedings of 6th Intl. Conference on Cognitive Modeling*, Pittsburgh.

Fellous, J-M. (2004). From Human Emotions to Robot Emotions. In *AAAI SS: Architectures for Modeling Emotion*, TR SS-04-02. Menlo Park, CA: AAAI Press.

Hudlicka, E. (1998). *Modeling Emotion in Symbolic Cognitive Architectures.* In *Proceedings of the AAAI FS: Emotional and Intelligent I*, Orlando, FL.

Hudlicka, E. (2003). *Modeling Effects of Behavior Moderators on Performance.* In *Proceedings of BRIMS-12,* Phoenix, AZ.

Hudlicka, E. (2004). *Two Sides of Appraisal.* In *Proceedings of* the AAAI SS: Architectures for Modeling Emotion, Stanford, CA.

Hudlicka, E. (2005). *MAMID-ECS: Application of Human Behavior Models to Risk-Analysis and Risk-Reduction Strategy Development in Human-System Design*. Report 0507. Blacksburg, VA: Psychometrix.

Hudlicka, E.(2007). Reasons for Emotions. In W. Gray (Ed.), *Advances in Cognitive Models and Cognitive Architectures*. NY: Oxford.

Isen, A. M. (1993). Positive Affect and Decision Making In J. M. Havliand & M. Lewis (Ed.), *Handbook of Emotions.* NY: Guilford.

Leventhal, H. & Scherer, K.R. (1987). The relationship of emotion to cognition. *Cognition and Emotion*, 1, 3-28.

Matthews, G.A. (2008). Personal communication.

Matthews,G.A., & Harley,T.A. (1993). Effects of Extraversion Self-Report Arousal on Semantic Priming. *Jnl. of Personality & Social Psychology, 65*(4), 735-756.

Mineka, S., Rafael, E., & Yovel, I. (2003). Cognitive Biases in Emotional Disorders. In R.J. Davidson, K.R. Scherer, H.H. Goldsmith (Eds.). *Handbook of Affective Science*. NY: Oxford.

Ortony, A., Norman, D., & Revelle, W. (2005). Affect and Proto-Affect in Effective Functioning In *Who Needs Emotions?* NY: Oxford.

Ritter, F. E., & Avraamides, M. N. (2000). *Steps Towards Including Behavior Moderators in Human Performance Models*: College Station, Penn State University.

Smith, C. & Kirby, L. (2000). Consequences require antecedents: Toward a process model of emotion elicitation. In J. P. Forgas (Ed.), *Feeling and Thinking: The role of affect in social cognition*. NY: Cambridge.