An Architecture for Emotional Decision-Making Agents

Eric Chown
Bowdoin College
8650 College Station
Brunswick, ME 04011
207-725-3084
echown@bowdoin.edu

Randolph M. Jones
Colby College & Soar Technology
5847 Mayflower Hill Drive
Waterville, ME 04901-8858
207-872-3831
rjones@soartech.com

Amy E. Henninger
Soar Technology
3600 Green Ct., Suite 600
Ann Arbor, MI 48105
407-207-2237
amy@soartech.com

ABSTRACT
Our research focuses on complex agents that are capable of interacting with their environments in ways that are increasingly similar to individual humans. In this article we describe a cognitive architecture for an interactive decision-making agent with emotions. The primary goal of this work is to make the decision-making process of complex agents more realistic with regard to the behavior moderators, including emotional factors that affect humans. Instead of uniform agents that rely entirely on a deterministic body of expertise to make their decisions, the decision making process of our agents will vary according to select emotional factors affecting the agent as well as the agent's parameterized emotional profile. The premise of this model is that emotions serve as a kind of automatic assessment system that can guide or otherwise influence the more deliberative decision making process. The primary components of this emotional system are pleasure/pain and clarity/confusion subsystems that differentiate between positive and negative states. These, in turn, feed into an arousal system that interfaces with the decision-making system. We are testing our model using synthetic special-forces agents in a reconnaissance simulation.

Categories and Subject Descriptors
I.2.0 [**Artificial Intelligence**]: General – *cognitive simulation.*

General Terms
Human Factors.

Keywords
Emotion, arousal, pleasure, pain, clarity, confusion.

1. EMOTIONS AND DECISION-MAKING
This paper describes a framework for modeling emotions in an interactive, decision-making agent intended to provide a nearly human level of competence in a focused task domain. We regard emotions essentially as subconscious signals and evaluations that inform, modify, and receive feedback from a variety of sources including higher cognitive processes and the sensorimotor system. Because our project focuses on decision making, our model emphasizes those aspects of emotion that influence higher cognition and not those that affect, for example, the immune system. We are integrating a connectionist model of emotions from Chown [1] with Rosenbloom, Laird, and Newell's [4] Soar architecture. The application area incorporates emotions and individual differences into the behavior models of synthetic special-forces agents in an army reconnaissance simulation.

In our framework, symbolic assessments of a small set of "emotional attributes" reside in a working memory. Portions of working memory are accessible by the deliberate cognitive process, and portions are accessible by the emotion mechanisms, so memory serves as the interface between the two. These working memory elements combine with background knowledge to generate strategies, reasoning, and external behavior. At the same time, the cognitive model creates working interpretations of the environment and status of internal goals (situational awareness). Some of these interpretations and assessments feed into the connectionist model, which in turn continuously computes new values for each emotional attribute. Because the theory underpinning our model assumes that these responses were ultimately provided by evolution, we assume that these constraints are, on *average*, beneficial to decision making. Clearly, not all emotional responses are always beneficial. Thus, we recognize the need to demonstrate such tradeoffs in our experiments. We have implemented the architecture within an existing Soar model, and have just begun testing the implementation.

Emotions are often seen as being disruptive to rational thought. However, emotions can also be viewed as an efficiency measure to change the decision making process in beneficial ways. For example, in a dangerous world, agents cannot afford to spend time considering every possibility, but must respond quickly. Kaplan [3] suggests that humans needed to develop numerous ways to process information efficiently, especially since human survival relies upon information processing rather than sharp claws or teeth. Although it is clear that emotions sometimes impede deliberative decision making, out implementation adopts Kaplan's view that emotions also provide a way of coding and compacting experience to enhance fast response selection. In evolutionary terms, it is better to respond immediately to the sight of a large animal, perhaps by fleeing, than to take the time to rationally consider the best course of action.

AAMAS'02, JULY 15-19, 2002, BOLOGNA, ITALY.

In our model, emotions are useful for quickly providing an organism with three critical assessments of an its current state. 1) How important is the current situation? 2) Is the current situation dangerous or beneficial? 3) How effectively can the situation be dealt with? Each of these questions corresponds to a different mechanism in our model.

In humans, arousal is a general measure of how important a situation is. What is commonly called arousal is actually a collection of related responses including, among others, increased heart rate and respiration, and changes in levels of dopamine, norepenephrine and other chemicals in the brain. These changes modulate responses to the world, whether the responses are cognitive or physical. Since we are focussing on decision making, our focus is on how arousal impacts cognition. Among the effects of arousal are that learning is increased with increased arousal and cognition becomes more focussed and reactive as background cortical noise is suppressed.

Through evolution, situations that are replenishing, or have benefits in terms of reproduction are coded as pleasurable. Situations that are directly damaging or dangerous are coded as painful. Experiencing or anticipating either pleasure or pain provides an organism with a general course of action. For example, when experiencing pain an organism will tend to become more active with the idea that something must be done to stop the source of the pain.

Finally, clarity and confusion provide an organism with an assessment of how effectively it can cope with the current situation. Confusion is a sign that the organism's cognitive structure is not adequate to cope with the current situation, while clarity comes when an organism's internal model is in accord with what is happening in the world. Since confusion is a sign of danger it is painful. Since clarity indicates the safety of good decisions, it is pleasurable.

In our model, these mechanisms (pleasure/pain, arousal, and clarity/confusion) form the core of the emotional system, shown in Figure 1. Instead of positing separate systems for fear, anger, etc., it is our view that humans attach emotional labels to various configurations of these mechanisms (in addition to other, more cognitive, factors). For example, fear comes from the anticipation of pain. Anxiety is similar to fear, except the level of arousal is lower. Joy comes from either directly experienced pleasure or the anticipation of the same. The advantage of such a generalized system is that it does not require specialized processing for every conceivable situation or emotion.

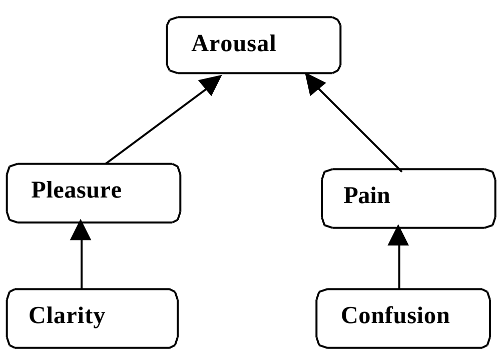

Figure 1: **The emotional subsystem.**

The primary way that the emotional system interacts with the cognitive system is through arousal. Memory and attention are the cognitive components most affected by changes in arousal. Highly aroused people are likely to fall back on well-learned knowledge and habits, even when they might have more relevant knowledge available.

In our implementation, the Soar rules that comprise the knowledge base have additional conditions added such that different kinds of rules only fire at differing levels of arousal. For example, highly cognitive rules will not fire at high levels of arousal, while more purely emotional rules may *only* fire at such levels. This system allows for very general approach to emotions. Emotional decisions, such to flee, might fire whenever arousal is high. This does not necessarily mean that the agent will flee, however, as we use a preference scheme to prefer rules that have more conditions (i.e. are less emotional). In a sense arousal provides an index to the rule-base. At low levels of arousal the rules likely to be indexed will be more purely cognitive and less purely reactive to perception (extremely low arousal corresponds to sleep). At the highest levels of arousal, by contrast, the rules indexed will be almost purely reactive with very little deliberative character.

The decision making agent is based on interactive real-time expert systems that are used for training simulations by the US military [2]. The Soar behavioral model used to evaluate the emotion model was Special Operations Forces (SOF) Soar. This task involves a 6-man team inserted deep within enemy territory for reconnaissance purposes. Our work parameterizes this model to make it susceptible to the emotional attribute levels generated by the connectionist component described above. One of the goals of this work is to demonstrate how emotions can affect decision making, and how the effects can vary according to the personality type of the agent.

2. ACKNOWLEDGEMENTS

This work was sponsored by the U.S. Army Research Institute's Emotional Synthetics Forces STTR – Ph II, contract number: DASW01-99-C-0037and by the NSF, under Grant No. 0092605. We gratefully acknowledge that support. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation (NSF).

3. REFERENCES

[1] Chown, E. (1994) Consolidation and Learning: A Connectionist Model of Human Credit Assignment. Doctoral Dissertation, The University of Michigan.

[2] Jones, R. M., Laird, J. E., Nielsen, P. E., Coulter, K. J., Kenny, P., & Koss, F. V. (1999). Automated intelligent pilots for combat flight simulation. *AI Magazine, 20*(1), 27-41.

[3] Kaplan, S. (1991). Beyond rationality: Clarity-based decision making. In T. Garling and G. Evans (Eds.), *Environment, Cognition, and Action: An Integrative Multidisciplinary Approach*, pp. 171-190. New York: Oxford University Press.

[4] Rosenbloom, P., Laird, J., and Newell, A (1993). *The Soar Papers: Research on Integrated Intelligence.* MIT Press, Cambridge, MA. 1993.