Motivations, Values and Emotions: Three Sides of the same Coin

Stan Franklin and **Uma Ramamurthy**
Computer Science Department & Institute for Intelligent Systems
The University of Memphis, Memphis, TN 38152, USA.
franklin@memphis.edu and urmmrthy@memphis.edu

Abstract

This position paper speaks to the interrelationships between the three concepts of motivations, values, and emotion. Motivations prime actions, values serve to choose between motivations, emotions provide a common currency for values, and emotions implement motivations. While conceptually distinct, the three are so pragmatically intertwined as to differ primarily from our taking different points of view. To make these points more transparent, we briefly describe the three in the context a cognitive architecture, the LIDA model, for software agents and robots that models human cognition, including a developmental period. We also compare the LIDA model with other models of cognition, some involving learning and emotions. Finally, we conclude that artificial emotions will prove most valuable as implementers of motivations in situations requiring learning and development.

1. Introduction

Motivations, values and emotions have been studied by philosophers, psychologists and neuroscientists for decades (Busemeyer, Medin and Hastie 1995; Dalgleish and Power 1999; Reiss 2001; Bower 1974; Russell 2003; Aharon et al 2001; Silverta et al 2004; Rolls 1999; Izard 1993; Davidson et al 2004; McGaugh 2003 and countless others). More recently, roboticists and artificial intelligence researchers have taken up these subjects (Sloman 1987; Wright, Sloman and Beaudoin. 1996; McCauley and Franklin 1998; Antunes and Coelho 1999; McCauley, Franklin, and Bogner 2000; Marsella and Gratch 2002; Langley et al 2003; Franklin and McCauley 2004; Avila-García and Cañamero 2005; Shanahan 2005).

Every autonomous agent (Franklin & Graesser 1997), be it a human, some other animal, a software agent or an autonomous robot, must come equipped with built-in primitive motivations. Otherwise, it wouldn't know how to decide what to do next. Evolution sees to these primitive motivations in biological agents; their designers build them into artificial agents, including epigenetic robots. Each such agent "lives its life" via a continual iteration of sense-process-act cycles during which it samples its environment, decides how best to respond to the current situation, and acts in response (Franklin 1995, the action selection paradigm). Motivations play a primary role in these action selection processes. Just as goals may have sub-goals in their service, primitive motivations may have sub-motivations in theirs. Each motivation in any agent must be in the service of one or more primitive motivations.

One primitive motivation for bacteria is to find nutrients. This motivation is implemented causally (mechanically) by chemo-taxis, the ability to follow a positive nutrient gradient (Alon et al 1999). Increasing concentrations of a particular nutrient molecule in sensory receptors causes the bacterium to tumble less and to swim forward more in the direction of the increasing gradient. This is an example of positive tropism, the involuntary response of an organism, orienting it toward an external stimulus. It can be viewed as a causal implementation of motivation.

A second way of implementing motivation is by explicitly including primary motivations in the form of drives in the action selection mechanism itself (Maes 1989). The computational IDA described below was designed in this manner (Negatu and Franklin 2002).

Yet another way of implementing motivations is with values. A value reflects an agent's general preference for a situation or an action, independent of its current beliefs or goals, that is, independent of the current environmental situation and of the agent's current intentions (Antunes and Coelho 1999). Values are often combined into an utility function that is used by the agent to evaluate options (Wahde 2003). Such evaluation of options, often referred to as reasoning or rational agency, can be effected by deliberation, a kind of internal virtual reality (Sloman 1999; Franklin 2000a). Such deliberation must consider the current environmental situation and the agent's current intentions, as well as its values. The object, of course, is to select an appropriate action.

Feelings in humans include hunger, thirst, various sorts of pain, hot or cold, the urge to urinate, tiredness, depression, etc. Damasio views feelings as somatic markers (1999). One feels feelings in the body. Implemented biologically as somatic markers, feelings typically attach to response options and, so, bias the agent's choice of action. We'll see below how this occurs in the LIDA model.

Panksepp posited that the foundation of emotional feelings is contained in the evolved emotional action system of mammalian brains (2005). In the LIDA model (see below), feelings affect the sensory motor automatisms (SMA) of autonomous agents. When one is sad, it affects the actions chosen and how the actions are taken. Similar role of feelings can be observed with other feelings as well -- how one holds a cup when one is angry vs. when one is happy. Feelings affect our facial expressions and our spoken words as well. The feeling manifests in one's body affecting the SMAs. Further, we hypothesize that feelings modulate learning with an inverse U-curve.

Emotions, such as fear, anger, joy, sadness, shame, embarrassment, resentment, regret, guilt, etc., are taken to be feelings with cognitive content (Johnston 1999). One cannot simply feel shame, but shame at having done something. The something done constitutes the cognitive content. Similarly, one must be angry at someone, that someone being the cognitive content. Feelings, including emotions, are nature's means of implementing motivations for actions in humans and other animals. They have evolved so as to adapt us to regularities in our environments.

These general preferences derived evolutionarily from regularities can be viewed as values. Thus feelings become implementations of values in biological agents, providing a common currency for quick and flexible action selection.

Artificial feelings and emotions are beginning to play an increasingly important role as mechanisms for primary motivations in software agents and robots, as well as facilitators of learning in these systems (Marsella & Gratch 2002, Langley et al. 2003; Franklin and McCauley 2004, Avila-García and Cañamero 2005, Ahn and Picard 2006). Here we present a case study of such feelings and emotions playing both roles in an intelligent software agent capable of performing a practical, real world task. In this agent they are actively involved in every instance of action selection, and at least potentially involved in each learning event. The pervasive, central role that feelings and emotions play in the control structure of this software agent mimics the roles they play in human cognition, and gives rise to clarifying hypotheses about human decision making and several forms of human learning.

2. The LIDA Model

LIDA provides a conceptual (and potentially a computational) model of cognition (Franklin 2000, 2001b) implemented as a software agent (Franklin & Graesser 1997) or as an epigenetic robot. The computational IDA "lives" on a computer system with connections to the Internet and various databases, and does personnel work for the U.S. Navy, performing all the specific personnel tasks of a human (Franklin 2001a). In particular, IDA negotiates with sailors in natural language, deliberates, and makes voluntary action selections (Frankllin 2000a) in the process of finding new jobs for sailors at the end of their current tour of duty. IDA completely automates the work of certain Navy personnel agents (detailers) (McCauley and Franklin 2002).

The LIDA (Learning IDA) model implements and fleshes out Global Workspace theory (Baars 1988, 2002), which suggests that conscious events involve widespread distribution of focal information needed to recruit neuronal resources for problem solving. The LIDA implementation of GW theory yields a fine-grained functional account of the steps involved in perception, several kinds of memory, consciousness, context setting, and action selection. Cognitive processing in LIDA consists of continually repeated traversals through the steps of a cognitive cycle (Baars & Franklin 2003, Franklin et al. 2005), as described below.

The LIDA architecture (Figure 1) includes modules for perception (Zhang, et al. 1998), various types of memory (Anwar and Franklin. 2003, Franklin et al. 2005, D'Mello, Ramamurthy, and Franklin. 2006 in press), "consciousness" (Bogner, Ramamurthy and Franklin. 2000), action selection (Negatu and Franklin. 2002), constraint satisfaction (Kelemen, Liang, and Franklin. 2002), deliberation (Franklin 2000a), and volition (Franklin 2000a). The mechanisms of these modules are derived from several different "new AI" sources (Hofstadter and Mitchell. 1994, Jackson 1987, Kanerva 1988, Drescher 1991, Maes 1989). Figure 1 provides the current implementation status of the LIDA model.

The computational IDA senses strings of characters from email messages and databases, and negotiates with sailors via email. The computational IDA is a running software agent that has been tested and demonstrated to the satisfaction of the U.S. Navy. Detailers observing the testing commented, "IDA thinks like I do."

In addition to the computational model, we will also speak of the conceptual LIDA (Learning IDA) model, which includes additional capabilities that have been designed but not implemented, including mechanisms for feelings and emotions, and for various forms of learning.

The LIDA conceptual model contains several different memory systems. Perceptual memory (often called perceptual organization) enables identification, recognition and categorization, including of feelings.

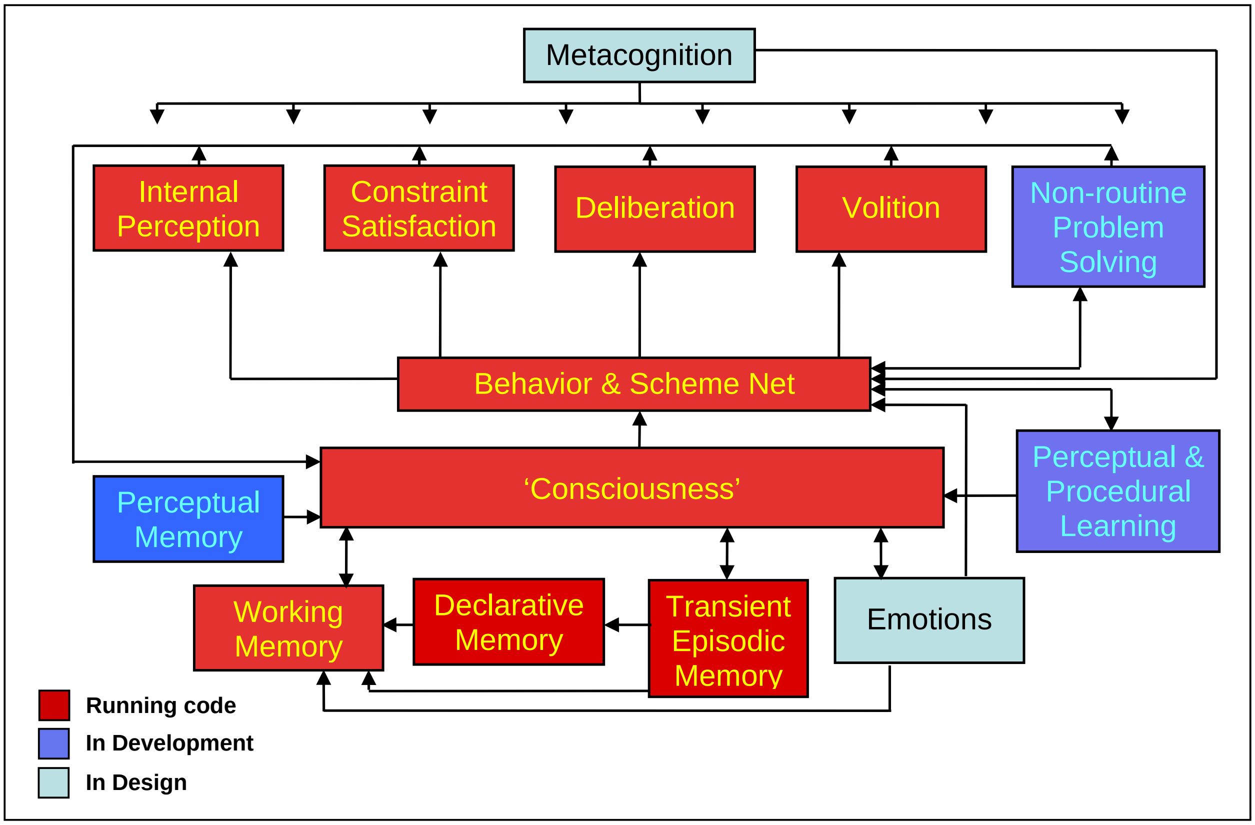

Figure 1: LIDA Architecture

Working memory provides preconscious buffers as a workspace for internal activities. Transient episodic memory is a content-addressable associative memory with a moderately fast decay rate. It is to be distinguished from declarative memory, that is, long-term associative, episodic memory. Procedural memory is long-term memory for skills. (Franklin et al 2005)

Much of the activity within LIDA is accomplished by codelets, small pieces of code that each performs one specialized, simple task. Codelets often play the role of daemons waiting for a particular type of situation to occur and then acting as per their specialization. Codelets in the LIDA model implement the processors postulated by global workspace theory. Neurally they can be thought of as cell assemblies or neuronal groups (Edelman1987, Edelman and Tononi 2000). Various sorts of codelets, including perceptual, attentional, behavioral, expectational, etc., will be described below.

3. The Cognitive Cycle

The LIDA model suggests a number of more specialized roles for feelings in cognition, all combining to produce motivations and to facilitate learning. Here we describe LIDA's cognitive cycle (Figure 2) in nine steps, emphasizing the roles played by feelings and emotions by putting their descriptions in italics. To aid the reader's understanding of the cognitive cycle, we will also carry along a running example of its operation distinguished by a different type font. Imagine that a teenage boy has just stepped through a classroom door into the hall and looked left. Another boy, a bully named Paul, is walking towards him down the hall.

1. **Perception.** Sensory stimuli, external or internal, are received and interpreted by perception constructing the beginnings of meaning. Note that this stage is unconscious.
   - ***a. Early perception:*** Input arrives through senses. Specialized perception codelets descend on the input. Those that find features relevant to their specialty activate appropriate nodes in the slipnet (a semantic net with activation).
   - ***b. Chunk perception:*** Activation passes from node to node in the slipnet. The slipnet stabilizes, bringing about the convergence (binding) of streams from different senses and chunking bits of meaning into larger chunks. These larger chunks, represented by meaning nodes in the slipnet, constitute

the percept. *Pertinent feeling/emotions are identified (recognized) along with objects and their relations by the perceptual memory system. This could entail simple reactive feelings based on a single input or more complex feelings requiring the convergence of several different percepts.*

Mostly visual perception activates the Bully node, representing an individual in the slipnet, along with the Paul node and the Fear node, resulting in their becoming part of the percept.

2. **Percept to Preconscious Buffer.** The percept, including some of the data plus the meaning, is stored in preconscious buffers of LIDA's working memory. These buffers may involve visuospatial, phonological, and other kinds of information. Again, note that this stage is unconscious.
   Feelings/emotions are part of the preconscious percept written during each cognitive cycle into the preconscious working memory buffers.

The Bully node, the Paul node and the Fear node are each part of the percept.

3. **Local Associations.** Using the incoming percept and the residual contents of the preconscious buffers as cues, including emotional content, local associations are automatically retrieved from transient episodic memory and from declarative memory. The contents of the preconscious buffers together with the retrieved local associations from transient episodic memory and declarative memory, roughly correspond to Ericsson and Kintsch's long-term working memory (LTWM) (1995) and Baddeley's episodic buffer (2000). Again, note that this stage is unconscious. *Feelings/emotions are part of the cue that results in local associations from transient episodic and declarative memory. These local associations contain records of the agent's past feelings/emotions in associated situations.*

The Bully and Paul nodes cue local associations (episodic memory) of the last encounter with this Bully including Fear and Pain.

4. **Competition for Consciousness.** Attention codelets, whose job it is to bring relevant, urgent, or insistent events to consciousness, view long-term working memory. Some of them gather information, form coalitions and actively compete for access to consciousness. The competition may also include attention codelets and their coalitions from a recent previous cycle. Again, note that this stage is unconscious. *Present and past feelings/emotions influence the competition for consciousness in each cognitive cycle. Strong affective content strengthens a coalition's chances of coming to consciousness.*

Some attention codelet notes the Bully, Paul, Fear and Pain nodes, among others in LTWM, gathers them, and others, into a coalition, and competes for consciousness.

5. **Conscious Broadcast.** A coalition of codelets, typically an attention codelet and its covey of related information codelets carrying content, gains access to the global workspace and has its contents broadcast. This broadcast is hypothesized to correspond to phenomenal consciousness. *The conscious broadcast contains the entire content of consciousness including the affective portions. The contents of perceptual memory are updated in light of the current contents of consciousness, including feelings/emotions, as well as objects, and relations. Affect modulate the encoding following an inverted U curve.*

The Paul, Bully, Fear and Pain nodes in perceptual memory receive additional base-level activations.

   Transient episodic memory is updated with the current contents of consciousness, *including feelings/emotions*, as events. *Affect modulate the encoding following an inverted U curve.*

The event of stepping into the hall, seeing Paul, remembering pain and experiencing fear is recorded in transient episodic memory.

   (At recurring times not part of a cognitive cycle, the contents of transient episodic memory are consolidated into long-term declarative memory.) Procedural memory (recent actions) is updated (reinforced) with the strength of the reinforcement influenced by the strength of the affect.

The action scheme for stepping into the hall has its base-level activation increases (reinforced).

6. **Recruitment of Resources.** Relevant behavior codelets respond to the conscious broadcast. These are typically codelets whose variables can be bound from information in the conscious broadcast. If the successful attention codelet was an expectation codelet calling attention to an unexpected result from a previous action, the responding codelets may be those that can help to rectify the unexpected situation. Thus consciousness solves the relevancy problem in recruiting resources. *The affective content (feelings/emotions) together with the cognitive content help to attract relevant resources (processors, neural assemblies) with which to deal with the current situation.*

Action schemes whose contexts contain "seeing a Bully" and "experiencing Fear" activate themselves.

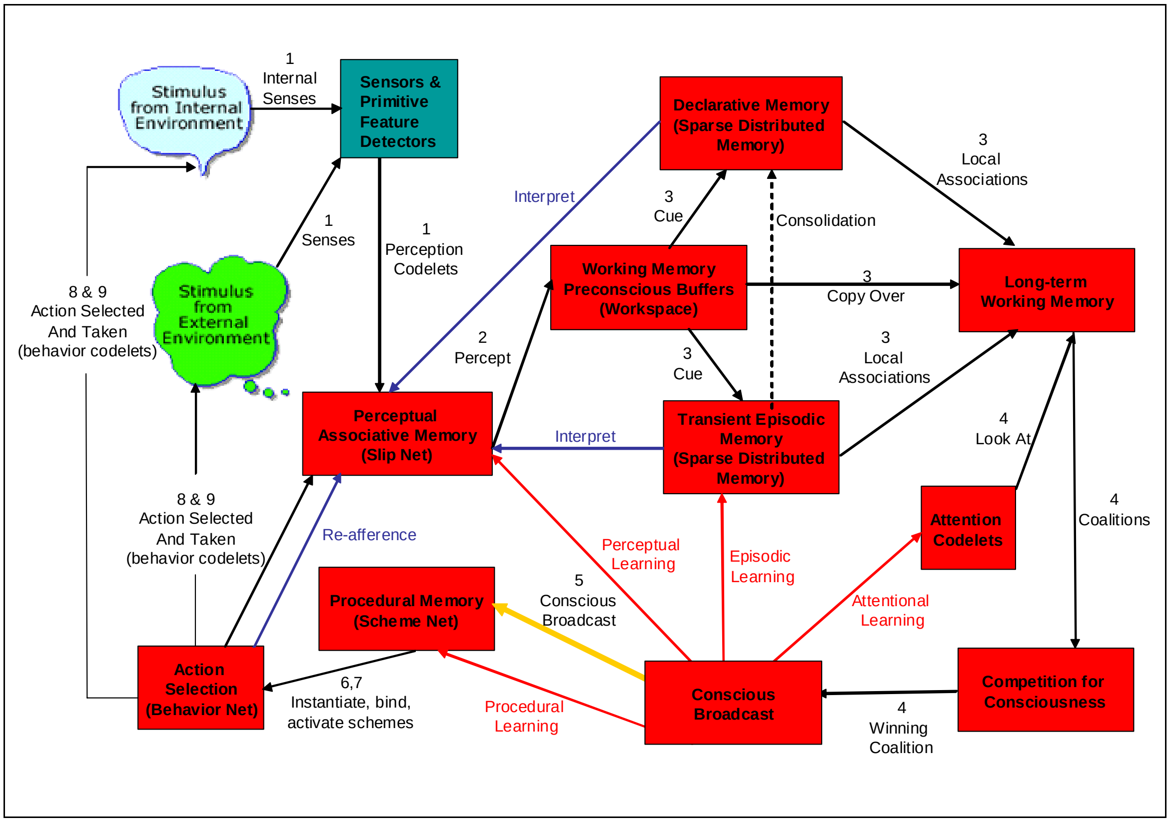

Figure 2: LIDA's Cognitive Cycle

7. **Setting Goal Context Hierarchy.** The recruited processors use the contents of consciousness, including feelings/emotions, to instantiate new goal context hierarchies, bind their variables, and increase their activation. It is here that feelings/emotions directly affect motivation. They determine which terminal goal contexts receive activation and how much. *It is here that feelings and emotions most directly implement motivations by helping to instantiate and activate goal contexts.* Other environmental conditions determine which of the earlier goal contexts receive additional activation.

Schemes with sufficient activation instantiate copies of themselves in the behavior net with their variables bound. One of these is a scheme for stepping back into the classroom.

8. **Action Chosen.** The behavior net chooses a single behavior (goal context), perhaps from a just instantiated behavior stream or possibly from a previously active stream. *This selection is heavily influenced by activation passed to various behaviors influenced by the various feelings/emotions.* The choice is also affected by the current situation, external and internal conditions, by the relationship between the behaviors, and by the residual activation values of various behaviors.

The action scheme to step back into the classroom is selected.

9. **Action Taken.** The execution of a behavior (goal context) results in the behavior codelets performing their specialized tasks, which may have external or internal consequences. This is LIDA taking an action.

The student steps back into the classroom.

The acting codelets also include an expectation codelet (see Step 6) whose task it is to monitor the action, and to try and bring to consciousness any failure in the expected results.

We suspect that cognitive cycles occur five to ten times a second in humans, overlapping so that some of the steps in adjacent cycles occur in parallel (Baars and Franklin. 2003). Seriality is preserved in the conscious broadcasts.

4. Related Work

The LIDA architecture differs significantly from other cognitive architectures such as SOAR (Laird et al. 1987) and ACT-R (Anderson & Lebiere 1998) in that it is not a unified theory of cognition in the sense of Newell (1990). Rather, its various modules are implemented by a variety of different mechanisms including the Copycat architecture, Sparse Distributed Memory, Pandemonium Theory and Behavior Nets (Franklin 2001b). Though the LIDA architecture contains no production rules and no neural networks, it does incorporate both symbolic and connectionist elements. The LIDA architecture allows feelings and emotions to play a central role in perception, memory, "consciousness" and action selection. LIDA's "consciousness" mechanism, based on Global Workspace Theory, resembles a blackboard system (Nii 1986), but there is much more to the LIDA architecture having to do with the interaction of its various modules. Much of this interaction is described in the cognitive cycle detailed above.

The LIDA architecture can be viewed as a specification of the more general CogAff architecture of Sloman (Wright et al. 1996). It has reactive and deliberative mechanisms but, as yet, no meta-management. There is a superficial resemblance between the computational IDA and the ACRES system (Moffat et al. 1993) in that both interact with users in natural language. LIDA and ACRES are also alike in using emotions to implement motivations. Rather than viewing emotions as implementing motivations for the selection of actions on the external environment, Marsella and Gratch study their role in internal coping behavior (2002). From our point of view this is a case of emotions implementing motivation for internal actions as also occurs in the LIDA conceptual model. The ICARUS system also resembles a portion of the LIDA conceptual model in that it uses affect in the process of reinforcement learning (Langley et al. 1991).

5. Conclusions

Being generated from order-of-magnitude one hundred thousand lines of code, IDA is an exceedingly complex software agent. Thus from the usefulness of artificial feelings and emotions in the LIDA architecture one would not jump to the conclusion that they would play useful roles in a more typical order-of-magnitude simpler software agent or robotic control structure. Besides, feelings and emotions are, as yet, only part of the LIDA conceptual model, and have not been implemented. Significant difficulties could conceivably occur during implementation. That artificial feelings and emotions seem to play significantly useful roles in the conceptual version of LIDA's cognitive cycles is not a conclusive argument that they will do so in simpler, implemented artificial autonomous agents.

Still the LIDA model suggests that software agents and robots can be designed to use feelings/emotions to implement motivations, offering a range of flexible, adaptive possibilities not available to the usual, more tightly structured motivational schemes such as causal implementation, or explicit drives and/or desires/intentions.

So, what can we conclude? Note that the computational IDA performs quite well with explicitly implemented drives rather than with feelings and emotions. It is possible that a still more complex artificial autonomous agent with a task requiring more sophisticated decision making would require them, but we doubt it. Explicit drives seem likely to suffice for quite flexible action selection in artificial agents, but not in modeling biological agents. It appears that feelings and emotions come into their own in agent architectures requiring sophisticated learning. This case study of the LIDA architecture seems to suggest that artificial feelings and emotions can be expected to be of most use in software agents or robots in which online learning of objects, categories, relations, events, facts and/or skills is of prime importance. If this requirement were present, it would make sense to also implement primary motivations by artificial feelings and emotions.

Acknowledgements

The authors are deeply indebted to the members of the 'Conscious' Software Research Group at the University of Memphis for their major contributions to this research, and to Mike Wintner for reading and commenting on an early draft.

References

Ahn, H. and R. W. Picard 2006. Affective Cognitive Learning and Decision Making: The Role of Emotions. The 18th European Meeting on Cybernetics and Systems Research (EMCSR 2006), April 18-19, 2006, Vienna, Austria

Alon, U., M. G. Surette, N. Barkai, and S. Leibler. 1999. Robustness in bacterial chemotaxis. *Nature* 397:168-171.

Anwar, A., and S. Franklin. 2003. Sparse Distributed Memory for "Conscious" Software Agents. *Cognitive Systems Research* 4:339-354.

Anderson, J. R., and C. Lebiere. 1998. *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Antunes, L., and H. Coelho. 1999. Decisions Based upon Multiple Values: The BVG Agent Architecture. In *Progress in Artificial Intelligence: Proceedings 9th Portuguese Conference on Artificial Intelligence, Lecture Notes in Computer Science*, vol. 1695, ed. P. Barahona, and J. J. Alferes. Berlin / Heidelberg: Springer.

Avila-García, O. and Cañamero, L. 2005. Hormonal Modulation of Perception in Motivation-Based Action Selection Architectures. In L. Cañamero (Ed.), *Proc. Agents that Want and Like: Motivational and Emotional Roots of Cognition and Action*, Symposium of the AISB'05 Convention, University of Hertfordshire, UK, April 14-15, 2005 (in press).

Baars, B. J. 1988. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.

Baars, B. J. 2002. The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Science* 6:47–52.

Baars, B. J., and S. Franklin. 2003. How conscious experience and working memory interact. *Trends in Cognitive Science* 7:166–172.

Baddeley, A. D. 2000. The episodic buffer: a new component of working memory? *Trends in Cognitive Science* 4:417-423.

Bogner, M., U. Ramamurthy, and S. Franklin. 2000. Consciousness" and Conceptual Learning in a Socially Situated Agent. In Human Cognition and Social Agent Technology, ed. K. Dautenhahn. Amsterdam: John Benjamins.

Bower.G. A. 1974. *The Psychology of Learning and Motivation*. New York: Academic Press.

Busemeyer, J. R., D. L. Medin, and R. Hastie. 1995. *Decision Makings from a Cognitive Perspective. The Psychology of Learning and Motivation*, vol. 32. Burlington, MA: Academic Press.

Dalgleish, T., and M. Power. 1999. *Handbook of cognition and emotion*. Chichester: Wiley.

Damasio, A. R. 1999. *The Feeling of What Happens*. New York: Harcourt Brace.

D'Mello, S. K., U. Ramamurthy, and S. Franklin. 2006. Encoding and Retrieval Efficiency of Episodic Data in a Modified Sparse Distributed Memory System. In *Proceedings of the 27th Annual Meeting of the Cognitive Science Society. Stresa, Italy*.

Drescher, G. L. 1991. *Made-Up Minds: A Constructivist Approach to Artificial Intelligence*. Cambridge, MA: MIT Press.

Edelman, G. M. 1987. *Neural Darwinism*. New York: Basic Books.

Edelman, G. M., and G. Tononi. 2000. *A Universe of Consciousness*. New York: Basic Books.

Ericsson, K. A., and W. Kintsch. 1995. Long-term working memory. *Psychological Review* 102:211-245.
Franklin, S. 1995. *Artificial Minds*. Cambridge MA: MIT Press.

Franklin, S. 2000. Modeling Consciousness and Cognition in Software Agents. In *Proceedings of the Third International Conference on Cognitive Modeling, Groeningen, NL, March 2000*, ed. N. Taatgen. Veenendal, NL: Universal Press.

Franklin, S. 2000a. Deliberation and Voluntary Action in 'Conscious' Software Agents. *Neural Network World* 10:505-521.

Franklin, S. 2001a. Automating Human Information Agents. In *Practical Applications of Intelligent Agents*, ed. Z. Chen, and L. C. Jain. Berlin: Springer-Verlag.

Franklin, S. 2001b. Conscious Software: A Computational View of Mind. In *Soft Computing Agents: New Trends for Designing Autonomous Systems*, ed. V. Loia, and S. Sessa. Berlin: Springer (Physica-Verlag).

Franklin, S., B. J. Baars, U. Ramamurthy, and M. Ventura. 2005. The Role of Consciousness in Memory. *Brains, Minds and Media* 1:1-38, pdf.

Franklin, S., and A. C. Graesser. 1997. Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents. In *Intelligent Agents III*. Berlin: Springer Verlag.

Franklin, S., and L. McCauley 2004. Feelings and Emotions as Motivators and Learning Facilitators. In Architectures for Modeling Emotion: Cross-Disciplinary Foundations, AAAI Spring Symposium Series; American Association for Artificial Intelligence; Stanford University, Palo Alto, California, USA; March 22-24.

Hofstadter, D. R., and M. Mitchell. 1994. The Copycat Project: A model of mental fluidity and analogy-making. In *Advances in connectionist and neural computation theory, Vol. 2: logical connections*, ed. K. J. Holyoak, and J. A. Barnden. Norwood N.J.: Ablex.

Jackson, J. V. 1987. Idea for a Mind. *Siggart* Newsletter, 181:23-26.

Johnston, V. S. 1999. *Why We Feel:The Science of Human Emotions*. Reading MA: Perseus Books.

Kanerva, P. 1988. *Sparse Distributed Memory*. Cambridge MA: The MIT Press.

Kelemen, A., Y. Liang, and S. Franklin. 2002. A Comparative Study of Different Machine Learning Approaches for Decision Making. In *Recent Advances in Simulation, Computational Methods and Soft Computing*, ed. E. Mastorakis. Piraeus, Greece: WSEAS Press.

Laird, E. J., A. Newell, and Rosenbloom P. S. 1987. SOAR: An Architecture for General Intelligence. *Artificial Intelligence* 33:1–64.

Langley, P., K. B. McKusick, J. A. Allen, W. F. Iba, and K. Thompson. 1991. A design for the ICARUS architecture. *ACM SIGART Bulletin* 2:104–109.

Langley, P., D. Shapiro, M. Aycinena, and M. Siliski. 2003. A value-driven architecture for intelligent behavior. In Proceedings of the IJCAI-2003 Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions.

Maes, P. 1989. How to do the right thing. *Connection Science* 1:291-323.

Marsella, S., and J. Gratch; 2002. A Step Towards Irrationality: Using Emotion to Change Belief. *1st International Joint Conference on Autonomous Agents and Multi-Agent Systems*. Bologna, Italy. July 2002.

McCauley, T. L., and S. Franklin. 1998. An Architecture for Emotion. In *AAAI Fall Symposium Emotional and Intelligent: The Tangled Knot of Cognition*. Menlo Park, CA: AAAI Press.

McCauley, L., and S. Franklin. 2002. A Large-Scale Multi-Agent System for Navy Personnel Distribution. *Connection Science* 14:371-385.

McCauley, L., S. Franklin, and M. Bogner. 2000. An Emotion-Based "Conscious" Software Agent Architecture. In *Affective Interactions*, Lecture Notes on Artificial Intelligence ed., vol. 1814, ed. A. Paiva. Berlin: Springer.

McGaugh, J. L. 2003. *Memory and Emotion: The Making of Lasting Memories*. New York: Columbia University Press.

Moffat, D., N. H. Frijda, and R. H. Phaf. 1993. Analysis of a computer model of emotions. In *Prospects for artificial intelligence*, ed. A. Sloman, D. Hogg, G. R. Humphreys A, and A. Ramsay. Amsterdam: IOS Press.

Newell, A. 1990. *Unified Theories of Cognition*. Cambridge MA: Harvard University Press.

Negatu, A., and S. Franklin. 2002. An action selection mechanism for 'conscious' software agents. Cognitive Science Quarterly 2:363-386.

Nii, H. P. 1986. The Blackboard Model of Problem Solving and the Evolution of Blackboard Architectures. *The AI Magazine* Summer 1986:38–53.

Panksepp, J. 2005. Affective Consciousness: Core emotional feelings in animals and humans. Consciousness and Cognition 14 (2005) 30-80.

Reiss, S. 2001. A sensitivity theory of end motivation: Implications for mental retardation. In *Personality and motivational differences in persons with mental retardation*, ed. H. Switzky. Mahwah, NJ: Erlbaum.

Shanahan, M. P. 2005. Consciousness, Emotion, and Imagination: A Brain-Inspired Architecture for Cognitive Robotics. In *Proceedings AISB 2005 Symposium on Next Generation Approaches to Machine Consciousness*.

Sloman, A. 1987. Motives Mechanisms Emotions. *Cognition* 1:217-234.

Sloman, A. 1999. What Sort of Architecture is Required for a Human-like Agent? In *Foundations of Rational Agency*, ed. M. Wooldridge, and A. S. Rao. Dordrecht, Netherlands: Kluwer Academic Publishers.

Wahde, M. 2003. A method for behavioural organization for autonomous robots based on evolutionary optimization of utility functions. *Journal of Systems and Control Engineering* 217:249-258.

Wright, I., A. Sloman, and L. Beaudoin. 1996. Towards a Design-Based Analysis of Emotional Episodes. *Philosophy Psychiatry and Psychology* 3:101–126.

Zhang, Z., S. Franklin, B. Olde, Y. Wan, and A. Graesser. 1998. Natural Language Sensing for Autonomous Agents. In *Proceedings of IEEE International Joint Symposia on Intellgence Systems 98*.