

Affect and Affordance: Architectures without Emotion

Darryl N. Davis and Suzanne C. Lewis

Department of Computer Science

University of Hull

Kingston-upon-Hull

HU6 7RX, UK.

D.N.Davis@dcs.hull.ac.uk

Abstract

General frameworks of mind map across tasks and domains. By what means can a general architecture know when it has adapted to a specific task and/or environment? Our current work on this theme presents an affect and affordance based core for mind. This draws on evidence from neuroscience, philosophy and psychology. However we differentiate between the mechanisms and processes thought to be allied to cognition and intelligent behavior in biological architectures and the foundational requirements necessary for similarly intelligent behavior or cognitive-like processes to exist in synthetic architectures. Work on emotion is a morass of definitions and competing theories. We suggest that we should not further this confused framework with further models of emotion for artificial systems. Rather we should look to foundational requirements for intelligent systems, and ask do we require emotions in machines or an alternative equivalent of use in control and self-regulation?

Introduction

There is a growing consensus among theorists and designers of complete intelligent systems (Minsky 1987; Franklin 2000; Sloman 2001) that synthetic minds, to be complete and believable, require a computational equivalent to emotion to complement their behavioral and cognitive capabilities. This requirement was highlighted in earlier prominent research (Simon 1967; Norman 1980). This paper confronts this requirement for emotion in intelligent systems on a number of grounds. The thesis is that overall the theory of emotion is too disorganized to be of much use in the design of synthetic intelligence and that more pointedly, emotion is not really a requirement for synthetic intelligence. It is suggested that a direction given by the less semantically overloaded term affect is a more appropriate for synthetic intelligence. The argumentation for affect as a control mechanism makes use of the control

state approach to mind, experimental work with computational agents and ongoing designs for intelligent systems. This argumentation has been developed from a number of research degree projects and the ongoing research of the principal author.

The Red Herring that is Emotion

Norman's pivotal paper (Norman 1980) suggested emotion-like processes will be necessary for artificially intelligent systems. This section builds an argument that denies the need for emotion in synthetic systems, while accepting that notable systems have been built based on models of emotion (Ortony, Clore and Collins 1988; Scherer 1994). The arguments presented do not apply to what Picard (1997) calls emotion recognizers.

Theories of emotion can be typified as belonging in one of several types, for example physiological (James 1884; Plutchik 1994), evolutionary (Darwin 1892), expression (Ekman 1994), appraisal (Scherer 2001) or goal based (Oatley 1992). This is partially due to different programmatic objectives within, for example, neurophysiology, psychology and philosophy. Duffy (1962) considers the use of the fuzzy, ambiguous and misleading term "emotion" as fundamentally flawed. Such terms should be abandoned as confusing and new or clearly delineated terms used only where such concepts are clearly and unmistakably identified. There is such a volume of research in this area that a significant academic revolution would be required to pursue such a path with any success. While this may be true of disciplines that study human intelligence, the same does not hold for the study of artificial systems.

Numerous prominent researchers into intelligent systems have suggested that affect-like mechanisms are necessary for intelligence (Simon 1967; Norman 1980; Minsky 1987; Sloman and Croucher 1987). More recently, Sloman (2001) has suggested that while emotion is associated with

intelligent behavior, it may not be a prerequisite. If that is the case and that emotion is a side-effect of mechanisms in sophisticated and complex biological architectures, the mechanisms responsible for intelligence in biological systems are now tightly bound to the control of these side-effects through evolution. The development of control mechanisms to harness and cope with the affective associations of the mechanisms necessary for intelligence, over the diachronic intervals associated with evolution, is such that in effect emotion and affect are now central to intelligence in biological systems. However no such constraint for synthetic (or artificial) architectures is in force or necessary. The theory of synthetic intelligent systems can therefore progress without the need for emotion. As there appears to be a requirement for something analogous to emotion, we suggest the use of affective control states. This can draw on theories of emotion and cognition in biological intelligent systems where necessary but not be limited or constrained to what is plausible or necessary for biological systems. For example, a synthetic system need not model or recognize the emotive state termed fear but merely recognize situations that (potentially) jeopardize its role(s) and tasks in its current environment. Put simply, theories of emotion from the cognate disciplines such as philosophy and psychology can afford functional models of affect for synthetic systems without the need for the theorist or designer of synthetic systems to be concerned with the semantic overloading associated with specific emotions, or indeed the concept of emotion per se. Furthermore most theories of emotion involve body attitude or facial expression changes that are typically inappropriate for machines. There are no machines that necessarily rely on body posture or facial expression for communication other those affective systems that attempt model the emotive state of their user (Picard 1997). An immediate benefit is the researcher interested in intelligent synthetic systems can move away from the definitional morass that surrounds much of the work on emotion. The exception of course is where models of emotion are required for sophisticated man-machine interactions. Even there the interactive system needs only to model the emotive/affective state of its user, and not actually function in terms of emotion.

Emotion and affect in synthetic systems

Two theory types do provide a sound basis for affect in synthetic agents. Emotion in the goal-based theories, for example, can be described as “*a state usually caused by an event of importance to the subject*” (Oatley 1992). Here emotion involves mental states directed towards an external entity, physiological change, facial gestures and

some form of expectation. Appraisal type theories define emotion as “*a sequence of interrelated, synchronised changes in the states of all organismic subsystems (...monitoring subjective feeling..) in response to the evaluation of an external or internal stimulus event that is relevant to central concerns of the organism*” (Scherer 1994). Here emotional processes involve five functionally defined mechanisms involving information processing over perception, regulation of internal states, decision making over competing motives, the control of external behavior and a feedback system across these four. There is considerable overlap between these definitions. In effect, emotions, in socio-biological agents, are affective mental, conative and/or cognitive, states and processes. A coherent hybridization of these theories is possible with a notable exception. A number of researchers use the concept of the basic emotions. Scherer instead allows for modal forms of emotive processing. Of the many modes that an emotion system can take, some are near identical or relatively similar to the states described as basic emotions.

Previous research (Davis 2001a) has indeed used emotional models that include basic emotions based on the work of (Oatley 1992). That work concluded with the current stance that emotion, and therefore basic emotions, are unnecessary for synthetic systems. The alternative developing theory of affect draws on the above two theories and themes such as control states (Simon 1967; Sloman 1993; Davis 2001b) and perceptual affordances (Gibson 1979; Davis and Lewis 2003). Affect is defined in terms of information processes and representational structures. It is qualitatively defined over the values {*negative, neutral, positive*} and numerically over the interval [-1.0,1.0]. Affect becomes the basis of a control language across a synthetic architecture. It allows external events and objects to take valenced affordances, and allows internal mechanisms to prioritize and be prioritized via affective valences.

A salient feature of the above definitions of emotion is their description in terms of goals and roles. Earlier work on agents and control states (Davis 2001b) focused on goal processing. It addressed how goals need to be valenced in a number of different ways, for example intensity, urgency, insistence. Affective values can be used to instantiate and modify aspects of motivators, whether low-level drives or abstract goals, and their associated representations. Furthermore, if an agent is to recognize and manage emergent behaviors, and particularly extreme and disruptive control states, this multi-layer model of affect provides the means for reflective processes to do this. This model of affect addresses the need to integrate different categories of processes in a synergistic manner. The affective valencing of processes and representational

structures can be given or the agent can adapt or learn according to its role and current environment. As an agent monitors its interactions within itself and relates these to tasks in its external environment, the impetus for change within itself (i.e. a need to learn) is manifested as an affect based state. Such a control state can lead to the generation of internal processes requiring the agent to modify its behavior, representations or processes in some way. The modification can be described in terms of a mapping between its internal and external environments measured over the difference between current and desired states. This influences the different categories of cognitive and animated behavior that an agent may use in attempting to satisfy its goals.

Motivators and Synthetic Control States

The control state approach to cognition (Simon 1979, Sloman 1993) builds on an assumption that cognition is an epistemic process that can be modeled using information processing architectures. Such information processing architectures can be in any number of non-exclusive control states. The processing of information can give rise to changes in the current set of extant control states. The nature of information processing is dependent upon the currently extant control states. The same information may be processed differently in different control states (Sloman 1993; Beaudoin 1994, Davis 2001). One current taxonomy provides five broad categories of control state. These broad categories overlap, and indeed the main focus of this research (motivators) are often used as a generic framework.

- Beliefs are internal models, possibly inferred from perceptual acts or from information arising from other control states; these need not have a rational basis.
- Images are control states using mental images, and are alike to Barsalou's (1999) perceptual simulations. These images may relate to any perceptual modality in typical or atypical ways.
- Imaginings are control states that embody alternative ways of constructing internal worlds. These can be related to directed problem-solving and therefore motivators, or make use of for example images.
- Motivators are problem-solving schemas which link internal and external 'reality', involving perception of events and states, representations and paths to modified states of affairs.
- Behaviours are actions on an environment. The most primitive type, reflexes, are ballistic mappings from input (i.e. perception) to output (i.e. behavioural response).

Here the focus is on motivators as a representational form that enables perception, affect and cognition to interact. The representational framework is outlined in this section. Argumentation and experimentation with design and computational models are introduced throughout the rest of the paper. Note that this treatment differs from the earlier references given above.

A motivator is a representational schema (akin to a *blackboard*) that brings together many aspects of perceptual and cognitive processing. The main components include:

- Semantic Content. Typically a proposition P denoting a possible and desired state of affairs.
- The Motivator Attitude to the semantic content P, for example make true, keep true, make false etc.
- An indication of current belief about the status of semantic content P, for example true, partially true, false.
- Actors and Agents referenced by the motivator
- Objects referenced by this motivator
- The plan sets for achieving the motivator. This may be null.
- Behaviours associated with any plan set and/or past similar motivators
- The affect process model associated with the motivator.
- The affordance model this motivator uses to reference external entities (for example a perceptual affordance model used for Actors and Objects referenced by a specific motivator).
- The current (commitment) status of the motivator, e.g. adopted, rejected, undecided, interrupted, completed.
- The dynamic state of the motivator e.g. being considered, nearing completion etc.
- Importance (e.g. neutral, low, medium, high, unknown). This may be intrinsic or based on associated processes, whether affective, cognitive or hybrid.
- Insistence: a heuristic value determining interrupt capabilities. This value results from the Affect model and importance, intensity and urgency components
- Intensity: In effect the value the affective system and processes associates with the difference between the current and desired state of the motivator
- Urgency Descriptor: How urgent is this descriptor – this may be qualitative (e.g. high, low) or quantitative (for example a time-cost function).

- A Decay Function which defines how insistence decreases while motivator is not attended to.

Associated with motivational structures are attitudes to classes of events and entities relevant to that motivator. These are valenced in the same way as affordances and affect. The association of perception, behavior and abstract representations about plans of actions and the relevance of actions and entities in the environment and internal are now defined over the same dimensions of negative, neutral positive with an associated numeric range of [-1, 1]. Affect and affordance become the means by which a sophisticated agent architecture can weigh its processes and control the economics of its processing. It provides a means whereby attention can be directed. These qualities can be defined over processes within a multiple level information processing architecture. Unlike many current architectural theories this need not necessarily be defined in terms of three columns and three or four levels if that proves too constraining (Norman and Shallice 1984; Gat 1997; Sloman 2001; Davis 2001b; Singh & Minsky 2003).

Linking Perception, Affect and Cognition

This section places recent investigations into affect and affordance within an ongoing framework of research into the development of architectures for synthetic intelligence built around the concept of control states, and in particular motivators. In these developing computational systems, activity and behavior at one level is represented and controlled at other layers. The primary conjecture is that the design and implementation of such architectures can proceed using a systematic control language that obviates the need for ad hoc heuristics to direct the processing within an intelligent system. This control language is domain-less to the extent that it can be used across multiple non-contingent domains (and applications) and yet be capable of capturing the control metaphors most suitable for any specific domain. This control language is grounded in the use of affect and needs to be consistent across different domains, tasks and levels of processing. The computational work is being developed with no requirement for emotion but rather a reliance on affect (a valencing of and within internal processes) and affordance (a motivational based perception of external events) that together can be used to guide both internal and external acts.

The central tenet is that behavior is purposeful and directed. Data and information processing is only performed when some motivation, whether conscious, subconscious, explicit, implicit, distributed or centrally represented, requires it. To do otherwise is poor design

and a waste of resource. The results of purposeful computation (sic) should matter to a synthetic system. For the results not to matter means one of three things. Firstly the system is focussed on irrelevant data and its attention has been drawn to irrelevant stimuli. Or secondly, it has performed the wrong type of processing on the information and obtained no relevant computational results. Or thirdly the data was relevant, the processing correct and the result simply poses no motivational gain or loss. An intelligent computational system needs to certainly minimize the first and second alternatives and use the third as feedback into the way it is controlling itself. Put simply there is typically far too much typically relevant data, information and means of processing it for a synthetic system to focus on irrelevant data, information and processing. In any thing other than the most spartan of environments (whether physical or virtual) an intelligent system is in need of powerful representational metaphors that can direct its focus of control to the motivationally relevant. Relying on emergence and tightly coupled environmental behaviours will not suffice.

Some Experimentation

The current work is based on experiments in the theory, design and implementation of affect and emotion based architectures (Davis 1996, 2001a, 2001b, Davis and Lewis 2003). Although earlier work did pursue the concept of computational emotion, this perspective is rejected for reasons already outlined. The continued research into affect draws on the ecological perspectives offered by Gibson (1979), and on the work of Simon's control state theory. Preliminary work (Davis 1996) centered on motivators and goals, how they come into being and how they are managed.

Recent experimental work (Bourgne 2003) revisited the representational structure and processes associated with motivators, but made use of affect and affordances instead of emotion. Theoretical, design and architecture parsimony, has lead to the rejection of emotion as a requirement for intelligence in synthetic systems - affect suffices. The usefulness of affect has been shown in related work (Davis and Lewis, 2003). This work (Lewis 2004) suggests that adding affect to cognitive desire and intention models such as CRIBB (Children's Reasoning about Intentions, Beliefs and Behavior) (Bartsch and Wellman 1989) results in more effective processing and task management. The major difference between CRIBB and ACRIBB agents is the latter use an affect value from their goals to order perceptual propositions. This affects belief order, in particular truth maintenance, and the simple goals the agents act out.

No of Agents	10 Time Intervals	20 Time Intervals	50 Time Intervals
2	32.5	79.5	80
4	29.3	76.5	9.5
8	33.63	76.5	8
10	29.3	65.6	5.5

No of Agents	10 Time Intervals	20 Time Intervals	50 Time Intervals
2	3.85	6.25	5.65
4	3.44	4.6	4.93
8	3.05	3.54	4.66
10	3.12	2.97	4.57

No of Agents	10 Time Intervals	20 Time Intervals	50 Time Intervals
2	92	100	70.5
4	91.4	100	66
8	88.6	99.4	62.15
10	89.5	98.2	58.7

Table 1. Fungus Eater Experiments (From top for Energy, Ore Collected and Life Expectancy) for affective (ACRIBB) and standard CRIBB agents

Tables 1a,b,c show results from using a standard agent test-bed, the Fungus Eater (Toda 1962), as the number of agents exceeds what is environmentally viable. The statistics collected are: Energy left at the end of the time interval; amount of ore collected; and the survival time of the agent. The environment contains 10 fungus (10% bad fungus), 10 ore, 5 small fungus, 1 golden ore and 1 medicine object. Only one type of agent was used at a time. All the results are an average value and the time survived is shown as a percentage of the time interval that the agents survived on average. The results from these experiments show that as the number of agents increases the amount of energy left decreases in both agent types. The ACRIBB agents maintain their energy level between 80 and 60 which corresponds to high and low energy drive thresholds respectively. The amount of ore collected by any one agent decreases as the number of agents increases. The ACRIBB agents out-perform the CRIBB agents in the amount of ore that is collected. ACRIBB agents can survive for 100% of the time intervals, but as the number of agents increase the survival rate decreases. The ACRIBB agents' survival rate reaches 92.64% at its lowest. The CRIBB agents, however, do not survive 100% of any time interval, but the survival rate also decreases as the number of agents increase.

Mechanisms for Affect and Motivation

This section builds on the material presented and develops an architectural theory for synthetic intelligence. This makes use of computational work on the nature of motivation and affect and how these can be co-joined in perception and cognition. This theory is currently incomplete. What follows here, are argumentation, design

and simple design experiments that demonstrate the more cogent aspects.

In the three-layer three-column model of Sloman (2001) and others, there exist reflexes and reactive behaviors that allow a direct response to sensory events. These can provoke processes or being modified at a more abstract level. Other automatic processes *necessitate* the generation of deliberative control states to achieve their goals. The deliberative layer represents those (control state) processes typically studied in thinking, human problem solving etc., plus other processes related to the management of low level actions. The reflective processes serve to monitor cognitive behavior or control it in some other way. The more extreme affective states (symptomatically categorized as a loss of control or perturbation) are effectively clamped by means of self-regulatory processes within the architecture. This model is quite general. The effect of altering the relative size and importance of the layers is an open issue. High level and low level processes coexist and interact in a holistic manner through the use of affect. In effect, goal processing, planning, decision making and other cognitive processes are not purely abstract but exist in relation to other automatic (and affective) processes. They are, in effect, embodied within the context of their interactions with their underlying processes and the agent within its environment.

Variations of the three-column, three layer architecture are used when they suit our research goals. However, there is no equivocal commitment to such architectures as the only frameworks suitable for building artificial minds. For example the given experimentation makes use of a pure deliberative three-column architecture based on cognitive models of reasoning in children (Lewis 2004). It does not matter at this stage in the research of artificial minds that any architecture seems to have many different levels. For current purposes less clearly delineated frameworks are developed using the combination of processes and mechanisms that support the concept of control states; in particular a combination of motivators and affect. There is no a-priori reason to coerce delineation across a related set of epistemic processes in terms of their reflexive, reactive, reactive planning, deliberative, reflective or meta-management mechanisms. To a large extent the coercion of processes into a strict delineation of architectural layers is an unnecessary constraint. Here we consider how small assemblages of processes can work together and then consider how instantiated instances of these process schemas are conjoined in an architectural whole.

The framework being developed here considers an agent situated in an environment. The agent can have a physical (for example, a robot) or virtual presence. Franklin (1997)

considers the principle function of mind to decide what to do next. This clear and unambiguous definition can be accepted as a basic tenet. A similar parsimonious approach to the role(s) attributed to emotion in biological agents from psychology enables a clear functional description for the role(s) affect must play in synthetic agents. The aim is not to mirror all the roles associated with emotion but to determine which are essential to a situated control system that needs to decide what to do next. To do so requires that aspects of the psychological theories of emotion, dismissed in the above sections as irrelevant to the design of synthetic minds, be revisited with this sole purpose; to determine which aspects are essential to a functioning synthetic mind. The concept of basic emotion can be rejected for reasons given above in favor of the modal response of Scherer's theory.

Agents need to determine whether events in their world favor or inhibit their goals. Similarly an agent needs to be able to determine whether an ongoing process within its control infrastructure offers potential solutions to tasks involving the juxtaposition of the agent's goals, the current state of the environment and the activities the agent can pursue in its environment and within itself. At its simplest each event, process and percept can be viewed as positive, neutral or negative. The conjugation of these valences offers a potentially multi-dimensional valence which may also be collapsed over a dimension internal to the agent represented also as negative, neutral or positive. A continuous value range can be mapped over the discrete valence range (i.e. [-1,1]). Frequently occurring or favored patterns of valences are an affect equivalent to the modal responses of Scherer.

If an object, agent, event or process are to be valenced, there needs to be some means by which such valences are instantiated. Motivators provide the schema that enables this. Earlier research (Davis 2001) looked to the philosophy and psychology of both human mind and current artificial intelligence to provide a taxonomy of motivators. Here a new taxonomy can be used that draws on these sources but does so without emotion and other control states not required for this simplest of affect models. The reason that control states such as motivators are needed is that agents are considered to be situated. This not only means that agents are coupled to their environment in terms perceptual and motor activities but situated in relation to the environment in terms of the tasks, roles, drives and goals the agents are expected (or designed) to maintain or achieve. This is analogous to the goal or motivational aspects of psychological theories of emotion. Events in an environment mean something to an agent not only in terms of their logical impact on the rationality of goal maintenance but also in terms of their

positive or negative valence for the agent in its current and any of its intended or possible future goal states.

Consider an agent ($Agent_1$) in a simple environment with a limited repertoire of primitive actions. It is capable of sensing two objects and one further agent (Figure 1). Objects offer behavioral affordances associated with R_1 (pick) and R_2 (push) and Agents R_N (collaborate). The agent must decide between pushing $Object_1$ or $Object_2$, picking $Object_1$ or $Object_2$ by synchronising its actions with $Agent_2$ or doing nothing. If $Agent_1$ sensed only one Object it would have only to decide whether or not to act on R_2 .

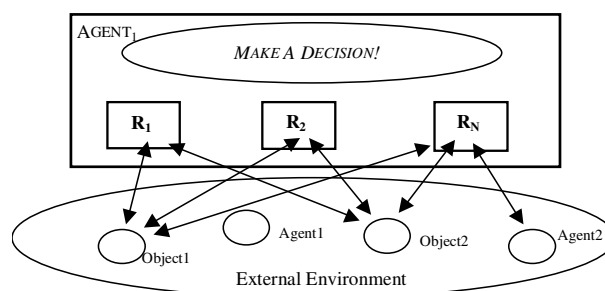


Figure 1. An agent in need of a decision!

The design for the agent can be determined by asking some very straightforward questions. If the question were 'What is the simplest model by which an agent can make this decision?', the response may well be along the lines of a behavior-based architecture (Brookes 1991), or a decision net (Maes, 1991). While this may preclude the use of explicit representations, in either case a valencing mechanism is required, either to relate sensory events to behaviours and also to allow subsumption of behaviours or in the latter case, the criteria by which the agent can decide between actions. If the question were 'What is the simplest model by which an agent can learn to make this decision?', then some form of representation whether connectionist or otherwise (for use in for example reinforcement, genetic or symbolic learning paradigms), would be necessary. Again some means of measuring the difference between actual and desired states would be required to act as a fitness function for the learning mechanism. The question we are asking is slightly different: 'What is the simplest generic mechanism by which an agent can learn to make this decision?'. In any case, across all three questions there is implied a common reliance on a means by which information can be ordered (i.e. valenced in some way) and the relating of sensory input to purposeful behavior (i.e. a motivational state of some description). The motivational schema introduced above provides the representational framework to allow this. The five-aside football scenario addressed by both Nunes (2001; Davis 2001a) and Bourgne (2003), which is

analogous to the above scenario demonstrated the applicability of the approach. Indeed Bourgne's work demonstrated that the inclusion of affect into the model produced more effective agents. Similarly the work of Lewis has shown how CRIBB was made more effective through the addition of motivational and affective states. Here these interlinked pieces of research are advanced through considering how the motivational framework can cross across two different (albeit rather simple) domains. The starting point is Agent₁ in figure 1 but with the addition of facilities that instantiate and manage motivational constructs (figure 2). Agent_M is this agent.

Agent_M has a number of actions it can take upon objects and agents in its environment. It has already learnt that it can associate behaviors R1 and R2 with Object1 and Object2, but can only perform R2 in collaboration with another agent. The simplest action is to push the nearest object. It chooses this and Object2 is now stuck in a corner. Attempting to push Object2 again results in the behavior having no effect. The affordance associated between objects and R2 (push) needs to be weakened in the current instantiation. The agent needs to change the way it views its environment, and makes decisions about which actions to apply. If the overall goal were to place Object2 on Object1 in the center of the environment, some simple plans bundled together within a reactive planner would suffice. The agent would simply use an alternative plan or behavior set. This would not help in answering the third question posed above, nor address the brittleness of such systems when the environment poses novel questions outside of the plan designer's mindset. As shown above in the fungus eater experiments, even in such simple environments the addition of affect results in more effective use of resources, particularly in stressed environments. However to enable an adaptive control model, that can vary the associations between potential actions, objects and other agents as result of feedback requires a more powerful valenced control model. A number of partially instantiated motivational constructs represent the alternatives open to the agent.

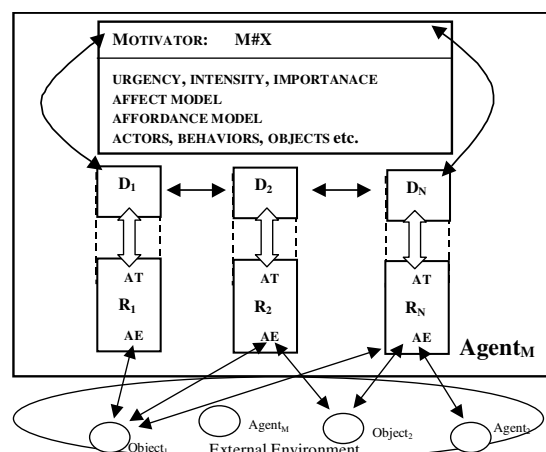


Figure 2. Motivator Enabling the Linking of Reactive (R_1 , R_2 , R_N) and Deliberative (D_1 , D_2 , D_N) Processes through Affect (AT), and the mapping of Affordances (AE) onto Actors and Objects in an external environment

Conclusion

We have presented experimental results and summaries from over 10 man-years research into a comparison of emotion and affect in computational agents. This paper has rejected the use of emotion in synthetic systems in favor of affective models. The argumentation made use of parsimony, computational requirements and computational plausibility of the psychological theories of emotion. It also considered how evolution has intertwined cognition and affect in biological systems to produce complicated interdependent systems. The designers of synthetic intelligent systems may require affective mechanisms in their systems for a number of reasons. These include criteria for sorting information and learning that relate to ongoing and temporarily local motivational states. These affective mechanisms need not be so computationally compromising as emotion as portrayed in any literature that reviews the emotions or the theory of emotion. It is suggested that affect can provide a consistent valencing mechanism across different classes of processes and mechanism in a sophisticated agent. It indeed provides the control language metaphor necessary for such agents. If we consider the role of emotion in Scherer's theory, there is a requirement for five functionally defined systems involving information processing over perception, regulation of internal states, decision making over competing motives, the control of external behavior and a feedback system across these four. The brief presentation here of the developing affect model can cover these requirements in synthetic systems. While the current explanation makes use of explicit representational forms, connectionist versions are not precluded. One future project will indeed address the implications of adopting a

connectionist approach to control state theory. It matters not in this developing theory whether a symbol is made explicit within a symbol reasoning system or whether it is a networked residue within a connectionist scheme.

References

- Barsalou, L.W. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577-609.
- Bartsch, K. & Wellman, H. 1989. Young children's attribution of action to beliefs and desires. *Child Development*, 60, 946-964.
- Beaudoin, L. 1994. *Goal Processing in Autonomous Agents*. Ph.D. Thesis, Cognition and Affect Group, Computer Science, University of Birmingham.
- Bourgne, G., 2003. *Affect-based Multi-Agent Architecture (for a 5-aside football simulation)*, M.Sc. Thesis, Computer Science, University of Hull.
- Brooks, R. A. 1991. Intelligence without representation. *Artificial Intelligence* 47, 139-160.
- Darwin, C. 1872. *The Expression of Emotions in Man and Animals*. London. Murray.
- Davis, D.N., 1996. Reactive and motivational agents. In: *Intelligent Agents III*, J.P. Muller, M.J. Wooldridge & N.R. Jennings (Eds.), Springer-Verlag.
- Davis, D.N., 2001a. Multiple Level Representations of Emotion in Computational Agents, *AISB2001: Agents and Cognition*, University of York.
- Davis, D.N., 2001b. Control States and Complete Agent Architectures., *Computational Intelligence*, 17(4), 621-650.
- Davis, D.N. and Lewis, S.C., 2003. Computational Models of Emotion for Autonomy and Reasoning. *Informatica*, 27(2), 159-165.
- Duffy, E. 1962, *Activation and Behavior*, Wiley & Sons.
- Ekman, P. 1994. *The Nature of Emotion: Fundamental Questions*. Oxford University Press. New York.
- Franklin, S. 1997. Autonomous Agents as Embodied AI. *Cybernetics and Systems* 28:499-520.
- Franklin, S., 2000. A consciousness based architecture for a functioning mind, *Proceedings of the AISB2000 Symposium on How to Design a Functioning Mind*. University of Birmingham.
- Gibson, J.J. 1979. *The Ecological Approach to Visual Perception*, Boston: Houghton Mifflin.
- James, W. 1884. What is an Emotion? *Mind*. 9, 188-205
- Lewis, S.C., *Computational Models of Affect*, Ph.D. Thesis, Department of Computer Science, University of Hull, To Be Submitted, 2004.
- Maes, P. 1991. The agent network architecture (ANA). *SIGART Bulletin*, 2(4), 115-120.
- Minsky, M.L. 1987. *The Society of Mind*, William Heinemann Ltd..
- Norman, D.A., 1980. Twelve issues for cognitive science, *Cognitive Science*, 4, 1-33.
- Nunes, H.A., 2001, *Investigation of Motivation in Agents Using Five-Aside Football*, M.Sc. Thesis, Department of Computer Science, University of Hull.
- Oatley, K. 1992. *Best Laid Schemes*. Cambridge. Cambridge University Press.
- Ortony, A., Clore, G.L. & Collins, A. 1988. *The Cognitive Structure of Emotions*. Cambridge University Press.
- Picard, R., 1997. *Affective Computing*, MIT Press.
- Plutchik, R. 1994. *The Psychology and Biology of Emotion*. Harper Collins. New York.
- Scherer, K. 1994. Toward a Concept of 'Modal Emotions'. In P. Ekman and R. Davidson (Eds): *Nature of Emotion*, Oxford University Press. New York.
- Scherer, K. 2001. Appraisal Considered as a Process of Multilevel Sequential Checking. In K. Scherer, A. Schorr, T. Johnstone (Eds). *Appraisal Processes in Emotion*. Oxford University Press. New York
- Simon, H.A. 1979. Motivational and emotional controls of cognition, Originally 1967, Reprinted in *Models of Thought*, Yale University Press, 29-38.
- Singh, P. and Minsky, M. 2003. An Architecture for Combining Ways to Think, *Proceedings of the International Conference on Knowledge Intensive Multi-Agent Systems*.
- Sloman, A. and M. Croucher. 1987. Why robots will have emotions. *Proceedings of IJCAI87*, 197-202.
- Sloman, A. 1987, Motives mechanisms and emotions. *Emotion and Cognition*, 1(3):217-234.
- Sloman, A. 1993, The Mind as a Control System, In *Philosophy and the Cognitive Sciences*, C. Hookway and D. Peterson, Cambridge University Press, 69-110.
- Sloman, A. 2001. Beyond shallow models of emotion, *Cognitive Processing*, 2(1), 177-198.
- Toda, M. 1962 The Design of a Fungus Eater, *Behavioural Science*, 7, 164-183.