Modelling Emotion in Computational Agents

Darryl N. Davis[1]

Abstract. This paper presents our recent investigations into computational emotion using agent architectures. An artificial life architecture is presented, using cellular automata as the base for a multi-layer heterogeneous processing infrastructure. In this sort of computational system, activity and behavior at one level is represented and to some extent directed at other layers. The primary conjecture is that through designing and implementing such architectures, emergent properties as studied in artificial life (and elsewhere) can be harnessed in a harmonious manner. Rather than consider emotion solely as an emergent property of such systems, an accepted psychological model of the basic emotions is used as a model for the processing performed by the agent. The work presented here allows the development of a fully integrated heterogeneous agency agent with possible applications in creative domains such as computer music. A further more encompassing possibility is that such research will prove useful in designing computer workstations where the affective interactions and emotive state of the user can be identified. This will enable the workstation to behave in manner more appropriate to the user's needs. The use of autonomous agents provides a framework with which to consider the alternative research directions that such projects entail.

Keywords. Affective Computation; Autonomous Agents; Cognitive Modelling.

1. INTRODUCTION

This paper brings together research in the computational modeling of emotion and agent architectures. The starting points for this ongoing research are the relation between emotion and cognition and the philosophical and computational work of Sloman. From the perspective of psychology, there seems to be a growing agreement on what are the basic emotions [13]. These emotions are defined in terms that are amenable to computational modeling. This recent work offers a useful model for the design of computational systems capable of identifying and autonomously monitoring and managing their internal state. Sloman's conjecture [14] is that complex computational systems will display emotion-like qualities, whether by design or as an emergent and sometimes perturbant behavior. This type of behavior may prove to be particularly troublesome if we consider affective computation [11] as part of the next generation of silicon systems. Therefore when designing complex or sophisticated systems, we should be aware of this possibility. Affective computation will require these next generation computational systems to be able to model emotion at a computational level in order to recognize and identify the affective state, and hence emotional stance, of their users. Furthermore, designs for intelligent systems should be such that the resulting computational system is capable of monitoring itself and catching perturbant behaviors before they become disruptive. Similarly an intelligent system should be capable of recognizing and harnessing beneficial emergent behaviors. Emergence is a major theme within the field of artificial life. Indeed Gilbert and Conte [5], in an overview of research in computational (artificial life) systems that model social life, note that only by designing computational systems with the appropriate representational frameworks can those systems systematically take advantage of emergent behavior. These are the issues that this research and paper address.

2. PSYCHOLOGY AND THE EMOTIONS

Emotion can be described as "*a state usually caused by an event of importance to the subject. It typically includes (a) a conscious mental state with a recognizable quality of feeling and directed towards some object, (b) a bodily perturbation of some kind, (c) recognizable expressions of the face, tone of voice, and gesture (d) a readiness for certain kinds of action*" [8]. Hence emotions, in socio-biological agents, are affective mental (conative and cognitive), states and processes, and any computational model of emotion must attempt to meet similar specifications.

A number of leading psychological researchers in the study of human emotions, for example Oatley and Jenkins [8], Power and Dalgleish [13], appear to be reaching an agreement on what are the basic emotions:

- Fear defined as the physical or social threat to self, or a valued role or goal.
- Anger defined as the blocking or frustrations of a role or goal through perceived other agent.
- Disgust defined as the elimination or distancing from person, object, or idea repulsive to self and to valued roles and goals.
- Sadness defined as the loss or failure (actual or possible) of valued role or goal.
- Happiness defined as the successful move towards or completion of a valued role or goal.

It is suggested that these five suffice as the basic emotions as they are physiologically, expressively and semantically distinct plus they have a biological basis. Oatley suggests that there may also be a case to consider surprise, hatred, contempt and desire (or interest) as further basic emotions. For the present work, the set of five will suffice, and in fact can be simplified to a four dimensional model. An alternative perspective on emotions [9] considers valence (i.e. positive-neutral-negative) and appraisal (cognitive reflection of these valencies) as the primary basis for emotion. They suggest that there are types of emotion, with specific emotions being instances and blends of these instances and types. Hence, sadness and happiness are antipathetic, being reflections of each other, or extremes on one dimension. Here we

---

[1] Department of Computer Science, University of Hull, Kingston-upon-Hull, U.K. email: D.N.Davis@dcs.hull.ac.uk

shall use the term sobriety, with sadness and happiness either side of a neutral state. The resulting four dimensional model is computationally tractable and maps onto our ideas for cognitive architectures. A further salient feature of these definitions of emotion is that they are described in terms of goals and roles. Earlier work on agents was focussed on goal processing [1] and considered how goals can be emotionally valenced. Hence, this four dimensional model provides a sound and consistent framework with which to investigate computational models of emotion and cognition.

Continuing on the theme of modeling behaviors, in a complete agent emotions are defined over different levels in terms of conative (pre-cognitive) and cognitive behaviors. Hence, if we are to model an emotion in a computational agent at the deliberative level (i.e. where they are semantically distinct), we also need to model it in terms of reactive or reflexive behaviors. Furthermore, if the agent is to recognize and manage emergent behaviors (that are related to emotions) then we need to ensure that there is design synergy across the different layers of the agent. The meta-management (reflective) processes that we had not fully addressed in earlier research [2] can be focussed using this model of emotion.

3. AGENT DESIGN

Others, for example [6], have suggested that complex social dynamics can be capably modelled using cellular automata, and that these computational processes provide a convenient modelling environment for the different microfoundations in human decision making. Our work on Go playing agents [4] used sets of homogeneous cellular automata. Here an insect community metaphor is employed, using sets of heterogeneous cellular automata. This provides a foundation for the agent model. The behaviour of any cellular automaton is directed by a (reactive) behaviour chosen by a deliberative reasoning process responsible for any specific cellular automata community. Other deliberative processes and a meta-management process are responsible for unifying the global behaviour of the agent.

Five categories of cellular automata are used. Insect guards (G) that act as data filters, insect workers (W) which act as information gatherers, insect drones (D) which maintain an ongoing (emotional) state within their community, and insect queens (Q) which act as the hub for the cellular automata communities and reflect the state of the hive. All are 4 connected, or potentially so, for reasons of computational simplicity and efficiency. Cellular automata communities are grouped as hives or mobiles. A mobile cellular automata community is capable of moving between hives and communicating hive states at the reflexive level (in effect synthetic-chemical communication). The computational nature of the cellular automata processing is modelled at the reactive level in terms of eleven sets of behaviors. Six sets of behavior mirror the possible connection combinations of cellular automata used in the design and the nature of their communication (types of cell connections), four sets are used to model the internal behavior of the automata, while the last is used to model cell community interactions.

At the deliberative level exist five types of deliberative agencies. The deliberative agencies are responsible for managing and manipulating declarative representations that reflect the goals, states and intentions of the agent. Some deliberative agencies are focussed on the management of very specific processes occurring in other agencies within the overall architecture. Other deliberative agencies, while perhaps capable of specific categories of processing, can turn its focus to a number of other processes within the overall architecture. This type of agency by redirecting its attention reflects changes in the concerns and focus of the agent.

D1 is a deliberative agency responsible for hive management. Exemplars of this type of agency are required for each hive. D2 is a deliberative agency responsible for the management of a mobile cellular automata community. Exemplars of this type of agency are required for every mobile. D3 acts a representation transformer, mapping agency states across different levels and representational schemes. It can, for example, focus its attention on the D2-reactive-mobile coupling, and represent the current focus of that agency in a more abstract manner of use in monitoring the emotional state of the agent overall. Agency D4 is responsible for the management of memory within the agent. It acts as a memory storage and recall filter, communicating with other deliberative agents about current and past states. D5 is responsible for the management of attention thresholds. Through communication with D4, and input to and from the reflective agency, it acts as a deliberative process observer and communicates, for example, control knowledge and control state suggestions to D1 and D2.

The reflective (meta-management) agency used in the present design is a recurrent multi-layer neural classifier. It takes as inputs the emotive state of the individual D1-hive and D2-mobile agencies, plus an overall state as given by the D3 agency. It outputs to the D1-hive agencies in terms of valencies over a four-dimensional emotion vector. These are mapped by the D1 agencies into changes in the reactive behaviours directing the CA communities.

The reactive agencies are sets of behaviors that are appropriate for different kinds of interactions at the reflexive level. At any one time, one or more behaviors from any of the eleven sets of behaviors will be active; the others remain dormant. It is the deliberative agency responsible for a cellular automata community that determines which are the currently extant sets of behaviors for the cellular automata. It is possible to view the reactive behaviors as the nature of communication between the individual and communities of cellular automata. Agency D1 is responsible for selecting the intra-hive behaviors, while agency D2 is responsible for the selecting the intra-mobile behaviors. As we allow interchanges between cellular automata communities via their most external cellular automata, further behaviors are required for mobile to hive exchanges. The required behavior sets are:

b1 - $\{ {}^{1}W_i\text{-}W_j, \ldots, {}^{k}W_i\text{-}W_j \}$ - worker with worker
b2 - $\{ {}^{1}G_i\text{-}W_j, \ldots, {}^{l}G_i\text{-}W_j \}$ - guard with worker
b3 - $\{ {}^{1}D_i\text{-}W_j, \ldots, {}^{m}D_i\text{-}W_j \}$ - drone with worker
b4 - $\{ {}^{1}G_i\text{-}G_j, \ldots, {}^{n}G_i\text{-}G_j \}$ - guard with guard
b5 - $\{ {}^{1}G_i\text{-}D_j, \ldots, {}^{r}G_i\text{-}D_j \}$ - guard with drone
b6 - $\{ {}^{1}D_i\text{-}Q_j, \ldots, {}^{s}D_i\text{-}Q_j \}$ - drone with queen
b7 - $\{ {}^{1}W_i, \ldots, {}^{a}W_i \}$ - worker internal
b8 - $\{ {}^{1}G_i, \ldots, {}^{b}G_i \}$ - guard internal
b9 - $\{ {}^{1}D_i, \ldots, {}^{c}D_i \}$ - drone internal

b10 - $\{ {}^{1}Q_j, \ldots, {}^{d}Q_j \}$ - queen internal
b11 - $\{ {}^{1}H_i\text{-}M_j, \ldots, {}^{x}H_i\text{-}M_j \}$ - hive with mobile

An explanation of these behaviors is given below in describing how emotion and communication is modelled at and between the various layers.

So why do we needs four kinds of cellular automata and what are their roles with regard to emotions? If we pursue the line taken by others (for example Ortony and Oatley) in their study of emotion in humans, we need to model emotion at both a cognitive (deliberative) level and at a conative (reactive or reflexive) level. For simplicity, we are initially modelling the valence of each of the 4 emotions as negative-neutral-positive. A discrete four-element vector can then be used to represent a specific emotional state. At the reflexive level, the queen cellular automaton reflects the emotional state of the hive. Hence, we need a four element vector to represent the queen's and hence the hive's emotional state. The hive's queen is, in effect, modelled as a 3^4 state cellular automaton. Each drone is responsible for maintaining the state of one emotion. For example, the drone to the north of the queen reflects the queen's anger, the one the east the queen's disgust, the one to the south the queen's fear and the one to the west the queen's sobriety. Each drone can be in one of three states, reflecting the three discrete states for the modelled emotion. Workers are two state automata; they hold a piece of data or they do not. Guards act as filters. They can accept or refuse input from up to four directions (being at most 4-connected); i.e. they are, in communication terms, 2^4 state automata. This disparity in cell type and enumeration of possible states means that the communication between automata of different types is necessarily not reflexive. At any given instant, the nature of the communication between automata is described by the currently active behaviors from the sets *b1* to *b6*.

In communities of homogeneous cellular automata, relatively simple rules govern the behavior of any cell [16]. For example, in a community of 4-connected Boolean (ON-OFF) automata, the next state of a cell depends on a count of the number of neighbor cells that are currently on. For example, if the count is two the cell stays in its current state, if the count is three it is switched ON, otherwise it is switched OFF. Such simple rules will not work here, as neighbor cells are not necessarily homogeneous. Even if such simple rules were applicable, we require these reflexive communities to be adaptive. Hence, we need a range of behaviors (or rules) for determining the next state of a cell when connected to cells of a specific type. This set of behaviors is extended to cope with the nature of communication between, for example, drone (three-valued) and worker (two-valued) cells. For example, behavior ${}^{1}b3$ states that a drone reads the state of a worker as ON or OFF, while the worker reads a drone as ON when ON, otherwise OFF. Behavior ${}^{2}b3$, on the other hand, states that a drone reads the state of a worker as ON or OFF, while the worker reads the drone as ON unless actually OFF. Further combinations are possible, for example the drone reads the worker as ON when OFF and vice-versa. This results in ten possible communicative behaviors for the drone and worker combinations. Only two possible communicative behaviors are allowed between workers; they read each other as is (i.e. ON when ON and OFF when OFF) or the inverse (ON when OFF and OFF when ON). Similar analyses define the remaining cell communication behavior sets b2 (four behaviors), b4 (two behaviors), b5 (ten behaviors) and b6 (13 behaviors). The last case (b6) requires further explanation. Both the drone and queen automata are three-valued (i.e. ON, Neutral or OFF). Hence they can communicate about their internal states (a 1-to-1mapping), or their internal states can be interpreted as in the brief description given for b3 above (a further twelve communicative behaviors).

Sets *b7* to *b10* define the internal behavior of the automata. All of the cellular automata can be four-connected, exceptions being on the outside of the hive or mobile, which can be two or three-connected (or if a hive and mobile are connected with each other, three and four connected). In the following COUNT refers to the number of connected cellular automata that are in state ON. Both b7 and b8 define the internal behavior of (internally) two-state cellular automata (the worker and guard cells), and similar internal behaviors are applicable in both cases, given by the following two rules:

R1: $\text{COUNT} = \text{T1} \quad \Rightarrow W_i \leftarrow \text{ON}$
$\text{COUNT} = \text{T2} \quad \Rightarrow W_i \leftarrow W_i$
$\Rightarrow W_i \leftarrow \text{OFF}$
R2: $\text{COUNT} > \text{T} \quad \Rightarrow W_i \leftarrow \text{ON}$
$\Rightarrow W_i \leftarrow \text{OFF}$
where $\text{T}, \text{T1}, \text{T2} \in \{0,1,2,3,4\}$

Rule R1 states the cell switches ON if the number of ON neighbor cells is exactly T1, stays in its current state if that count is T2 else it switches OFF. An alternative behavior, given by Rule R2, states a cell switches ON if the number of ON neighbor cells is greater than T, else it switches OFF. Other rules for defining the next state of a cell, governed by b7 and b8, are possible. Again two alternative rules define the two behaviors forming set b9 (the drones):

R3: $(Q_i = \text{ON}) \wedge (\text{COUNT} > \text{T1}) \Rightarrow D_i \leftarrow \text{ON}$
$(Q_i = \text{N}) \wedge (\text{COUNT} > \text{T2}) \Rightarrow D_i \leftarrow \text{ON}$
$(Q_i = \text{ON}) \wedge (\text{COUNT} \leq \text{T1}) \Rightarrow D_i \leftarrow \text{N}$
$(Q_i = \text{OFF}) \wedge (\text{COUNT} > \text{T3}) \Rightarrow D_i \leftarrow \text{N}$
$\Rightarrow D_i \leftarrow \text{OFF}$
R4: $\Rightarrow D_i \leftarrow Q_i$
where $\text{T1}, \text{T2}, \text{T3} \in \{0,1,2,3,4\}$

Rule R3 states that a drone cell (i.e. an emotion at the reflexive level) is positively valenced if the queen cell is similarly valenced or neutral (N) towards that emotion and sufficiently other neighbor cells are ON, otherwise the drone cell becomes neutral or OFF. Rule R2 simply states that a drone cell takes the same state as the queen cell for any specific emotion. Again other rules defining the next state of a drone cell are possible, and future experimental work will determine which are the most appropriate rules for which emotion given that the overall agent needs to maintain some overall role or achieve some objective or goal. Two similar rules define the internal state transformation (behavior b10) of the queen cell for any one of its four emotion elements:

R5: $(D_j = \text{ON}) \wedge (Q_{I,j} = \text{ON}) \Rightarrow Q_{I,j} \leftarrow \text{ON}$
$(D_j = \text{ON}) \wedge (Q_{I,j} = \text{N}) \Rightarrow Q_{I,j} \leftarrow \text{ON}$
$(D_j = \text{ON}) \wedge (Q_{I,j} = \text{OFF}) \Rightarrow Q_{I,j} \leftarrow \text{N}$
$(D_j = \text{OFF}) \wedge (Q_{I,j} = \text{ON}) \Rightarrow Q_{I,j} \leftarrow \text{N}$
$(D_j = \text{OFF}) \wedge (Q_{I,j} = \text{N}) \Rightarrow Q_{I,j} \leftarrow \text{OFF}$
$\Rightarrow Q_{I,j} \leftarrow Q_{I,j}$

R6: $\Rightarrow Q_{I,j} \leftarrow D_j$

Rule R5 states that a queen cell (i.e. the emotional state of a hive community) is positively valenced on some emotion if the drone responsible for that emotion is ON and the queen is either ON or neutral on that emotion. The queen cell becomes neutrally valenced towards that emotion if the appropriate drone is ON and the queen is OFF or if the drone is OFF and the queen is ON. The queen becomes negatively valenced on that emotion if drone is OFF and the queen neutral, otherwise the queen cell remains in the same internal state. Rule R6 (a mirror of rule R4) states that a queen cell takes the same state as the drone cell for any specific emotion.

Set *b11* is qualitatively different from any of the above rule sets in that it governs more global behavior at the reflexive level. It defines the nature of the interaction between a hive and a mobile community of cellular automata. Two constraints define the behaviors here. If no cell in hive H_i is vertically or horizontally connected with a cell in mobile M_j then there is no interaction between the two communities. If at least one cell in hive H_i is vertically or horizontally connected with any cell in mobile M_j, then one of 128 possible interactions is possible. The 128 interactions are given by 8 hive dorsal to mobile ventral, 8 hive ventral to mobile dorsal, 8 mobile posterior to hive anterior, 8 mobile anterior to hive posterior combinations with four possible rotations of the mobile. We have now described communication and internal state transformation within the agent at the reflexive layer. At the deliberative layer a different form of communication is used, based on our earlier research into decision-making deliberative and mobile agents that made use of a sub-set of the KIF-KQML protocol for the passing of messages between agents [3]. The content type of the messages depends on the nature of communicating agencies. D1 and D2 agencies communicate about hive-mobile couplings. Both D1 and D2 agencies communicate about their (overall) emotive state with the memory agency (D4) and the behaviour sets in place on previous similar emotive states.

There is however another category of communication; inter-layer communication between the different layers. The reflexive and reactive layers do not explicitly communicate with any other layer. The behaviors of the reflexive cells are set by the deliberative agencies at the reactive level. However, these deliberative agencies are capable of reading the internal state of any cell that they administer, particularly the queen cell in any hive. The deliberative agencies also provide input to the reflective agency, and receive messages from that agency. In effect, upward moving messages provide emotional state information, while downward moving messages act as behavioral and emotional directions. Some of our ongoing experiments are investigating the effect of provoking an emotional perturbation at one level and seeing how that effects the agent. For example, if we arrange for the agent to be emotionally neutral (across all the emotions and different layers), what happens if the agent experiences anger at just one level? Does it make a difference if it is at the reflexive, deliberative or reflective levels? What combination of states across the different layers causes the agent, as a whole, to experience an emotional perturbation? There are a considerable number of similar questions, all requiring further and extensive experimental work.

At a more global level, we can consider how these different agencies interact and their effect on the activities of the agent as a whole. We take our lead on this from the ongoing research into cognition as a dynamic process [12]. Here, mind is viewed as a multi-faceted activity, with many types of ongoing, stalled and dormant processes, some conscious, some sub-conscious, whether at the cognitive or pre-cognitive levels. Given that an agent is involved in some task involving mental activity, the agent should, though not necessarily always, be able to focus on that task. The knowledge, knowledge manipulating and other categories of processes activated or used in performing the task will vary over time, space and the agent's predisposition to perform that task. The metaphor used is not that of a complex computer algorithm making use of held and new knowledge and capabilities to solve the task as modelled in much computational cognitive science and artificial intelligence. A more fluid, less mechanistic metaphor is needed. This metaphor needs to take account of elements of the mind such as emotion, motivation (these are perhaps synonymous), attention and distraction. Any architecture for such a synthetic mind would provide only the starting point, as that architecture itself should be dynamic and reflect the current focuses and states of the agent. Achieving a believable computational model of this (if not impossible using today's computer systems) requires novel approaches. We suggest that the current design makes a small move in such a direction. By modeling emotion in a locally stochastic manner at the reflexive layer and allowing the behaviors at that level to be dramatically modified through interactions across the different levels of the agent architecture, a dynamic non-linear model of emotion (and processing) can be achieved. Furthermore not all the agencies need be active or exist all the time. If required the agent should be able to create new (deliberative-CAreflexive) agencies on demand to cope with new scenarios, and then disassemble them when no longer needed. To do this the agent needs to be aware of its ongoing cognitive processes; i.e. at some (albeit superficial level), the agent needs to be conscious of its own internal state. Through the use of the reflective-level emotional model, the agent should be able to classify its current emotional state as some point or space within the 4 dimensional model outlined earlier. This can then be used by agency D5 to vary the threshold that attributes specific (positive-neutral-negative) symbols to the real-valued output of the reflective agency. The agent can relate this reflective (meta-management) awareness of its emotional state to a deliberative classification as discussed in earlier work by Sloman [15] or the author [2]. The current (set of) deliberative control state(s) and ongoing goals are related using the four basic emotions which are defined in terms of goals at the deliberative level. The reflective classification is also related to the currently active reactive behaviors at the hive level, and in a wider sense to the currently active agencies. We suggest that these relations across the different agencies can be represented, at the deliberative level, using a relatively simple propositional logic, as suggested by McCarthy's work on robot consciousness [7]. Such an agent could be said to be aware of its emotional state. We do accept that the current design is limited in a number of ways and whether it can display the plasticity associated with human experience of emotion is a moot point. Whether it will suffice to provide "believable" emotional behavior is unclear at this time.

4. IMPLEMENTATIONS

Some recent work has focussed on the use of homogeneous cellular automata communities as components in an artificial life agent for playing the game of GO. Three related and incremental experiments, that integrate reflexive and deliberative behaviors, have been described elsewhere [4]. In short we moved from using communities of homogeneous cellular automata, with each community responsible for different but related functions, to a model of a multi-cellular organism based on similar-state cells within these communities. We are currently mapping this work on multiple cellular automata communities onto the SIM_AGENT toolkit. We will be able to compare the earlier approach to playing Go with the use of the emotion agent described in this document. Our present implementations of the current model are not grounded in any application. The experiments that we are running are investigating how the different control mechanisms interact. We are finding that any particular initial combination of cell states and reactive behaviours rapidly maps onto a steady state at the hive level. As the deliberative agencies change the currently active behaviours (from sets b1 through b11), minor changes in the steady state occur. The effect of coupling these agencies with the reflective classifier is to be investigated shortly. How these internal states relate to cognitive behaviours such as goal management is for future research.

The current work is an ongoing research program into how to design and build sophisticated agents that make use of past research efforts. In effect it portrays a component-based approach to building agents, analogous to the modular design work by NASA [10], but at a smaller scale. This approach allows us to consider for example a single cellular automaton as an agent in itself. We then select and combine a number of these agents to build a community of cellular automata, which are added to the set of prototype reflexive agents. We then build reactive agents that comprise of behaviors suitable for different categories of cellular automata communities. An agent of this class can then be combined with a cellular automata agent to produce yet another prototypical agent class. An instance of this prototypical agent is then combined with a deliberative agency to produce the hive-D1 coupling as described above. Similar design and implementation processes allow us to produce the D2-mobile coupling. Having produced prototypical agent classes for all the agencies required for our design, we then combine and refine their synthesis in a holistic manner. If at any point, any of the prototypical agent classes fail to meet their design requirements, particularly at the integration stages, that class of agent must be redesigned and the design process reiterated.

5. FUTURE WORK

The impetus in producing this current agent design was not given by any specific application and there is no pressure to keep it as anything but domain-less. We are using it as a computational model of a synthetic mind that can monitor and manage peturbance as emergent emotional states. For example, a user may interact via a graphical front-end and color code the agent and its various states to reflect specific emotional states within different agencies. The agent then responds by moving through various states in attempt to harmonize its internal state. At each point the various agencies and all the cells display their emotive state in the user-specified colors.

However, in pursuing the full implementation of this design it may be advantageous to ground it an application. For example, we can advance our earlier cellular automata work on artificial-life agents that play Go. While that work made no use of planning, we can use the deliberative agent's to devise strategic goals in playing the game (for example *laddering*). We can then relate emotional states to activities throughout the agent. For example, the agent's opponent in blocking a certain goal would cause the agent, at some local or global level, to experience anger. In other scenarios, the agent may experience fear at potentially losing the game. We can then experiment to see if the use of emotions produces a more capable game player. A further avenue that looks promising is the use this agent architecture (or analogous designs) in a more creative role. On a wider front, the research described here may provide a way to incorporate emotional indicators into workstations that work with affective logical input devices.

6. CONCLUSION

In this paper, we have presented an overview of some long-term research into the design of agent architectures for the computational modeling of cognition and affect. Earlier research provides the foundations on which to design an artificial life agent that computationally models emotion in a synergistic and holistic manner. We are currently mapping this design onto an agent toolkit to enable experimental work and so further our research on computational frameworks for the development of complete or cognitive agents. Progress in the theory, design and engineering of cognitive or complete agents is painstaking and difficult. Modern (cognitive) psychology is providing substantial theories of many of the facets that we need model. By designing agents such as the one described here, we can begin to test and use these models using synthetic agents.

7. REFERENCES

[1] Davis, D.N., Reactive and motivational agents. In: *Intelligent Agents III*, J.P. Muller, M.J. Wooldridge & N.R. Jennings (Eds.), Springer-Verlag, (1996).

[2] Davis, D.N. Synthetic Agents: Synthetic Minds, *IEEE SMC-98, San Diego, (1998)*.

[3] Davis, D.N., Agent-based decision support framework for water supply infrastructure rehabilitation. *Computers, Environment and Urban Systems*, **24**, 1-18, (2000).

[4] Davis, D.N., Chalabi, T. and Berbank-Green, B., Towards an architecture for artificial life agents: II, In: M. Mohammadian (Editor), *New Frontiers in Computational Intelligence and Its Applications*, ISO Press, (1999).

[5] Gilbert, N. and Conte, R., *Artificial Societies: The computer simulation of social life*, UCL Press, (1995).

[6] Hegselmann R. and Flache, A., Understanding complex social dynamics: A plea for cellular automata based modelling. *Journal of Artificial Societies and Social Simulation*, Vol. 1, No3, (1998).

[7] McCarthy J., Making Robots Conscious of their Mental States. *Machine Intelligence 15*, Oxford, (1995).

[8] Oatley, K. and Jenkins, J.M, *Understanding Emotions*, Blackwell (1996).

[9] Ortony, A., Clore, G.L. & Collins, A., *The Cognitive Structure of Emotions*. Cambridge University Press, (1988).

[10] Patterson Jr., F.G. and Evans, R.P, Perfecting the Image, *IEEE SMC-98, San Diego, (1998)*.

[11] Picard, R. *Affective Computing*, MIT Press, (1997).

[12] Port R.F., Van Gelder T. (Editors), *Mind As Motion*, MIT Press (1998).

[13] Power, M. and Dalgleish, T., *Cognition and Emotion*, LEA Press, (1997).

[14] Sloman, A. and Croucher, M. Why robots will have emotions. *Proceedings of IJCAI7*, 197-202, (1987).

[15] Sloman, A. Explorations in design space. In *Proceedings 11th European Conference on AI*, Amsterdam, (1994).

[16] von Newmann J., General and Logical Theory of Automata, In: A.H. Taub (ed), J. *von Newmann, Collected Works*, Pegamon, (1963).