PAC - Personality and Cognition: an interactive system for modelling agent scenarios

Lin Padgham and Guy Taylor
Department of Computer Science
Royal Melbourne Institute of Technology
Melbourne, VIC 3001, Australia
<linpa@cs.rmit.edu.au> <guy@yallara.cs.rmit.edu.au>

Abstract

PAC is an interactive system for experimenting with scenarios of agents, where the agents are modelled as having both cognition and personality, as well as a physical realisation. The aim of the system is to provide an environment where scenarios can quickly and easily be built up, varying aspects of agent personality (or emotions) and agent cognition (plans and beliefs). This will allow us to experiment with different combinations of agents in different worlds. We can then investigate the effect of various parameters on emergent behaviour of the agent system.

Some aspects of the system are described more fully in [PT97].

1 Motivation

The motivation for the system is to provide a testbed where we can easily build up scenarios of agents to investigate the effect of modeling emotion and personality as one of the important aspects of agents. Some researchers are interested in modelling emotion and personality as a way of creating agents that are engaging for the human user of a system (e.g. [Bat94a; HR95]). Others, such as Toda [Tod82] also believe that emotions play a functional role in the behaviour of humans and animals, particularly behaviour as part of complex social systems. Certainly the introduction of emotions, and their interaction with goals, at various levels, increases the complexity of the agents and social systems that can be modelled. Our hypothesis is that, just as modelling of *beliefs* and *goals* has facilitated the building of complex agent systems, so will the modelling of emotion and personality enable a further step forward in the level of robustness and complexity able to be developed. However significant work is needed before we can expect to understand the functional role of emotion sufficiently to successfully model it in our software agents. The PAC system will facilitate experimentation with groups of agents in varying worlds, and will allow us to observe the effect of representation of emotional characteristics and combinations of personality types under varying aspects of the simulated world.

Our aim in this system is not to develop a full psychological model of emotions, personality or cognition, but rather to use a simplified model to study how this might facilitate building of more complex systems, or more engaging and credible agents.

2 Emotional Model

We have based our model of cause of emotion on a simple version of two models found in the literature. The first is that explored in [OCC88], and used by both Dyer [Dye87] and Bates [Bat94b] in their systems. In this model emotional reactions are caused by goal success and failure events. For instance an agent which experiences a goal failure may feel unhappy, while one experiencing goal success may feel glad. Dyer [Dye87] develops a comprehensive lexicon of emotional states, based on goal success and failure. For example an agent which expects its goal to succeed will feel hopeful, an agent whose goal is achieved by the action of another agent will feel grateful, and an agent who expects its goal to be thwarted will feel apprehensive. Figure 2 shows examples of some emotions (modified from [Dye87]), indexed by contributing cause. We have currently implemented a simplified version of this model.

The second model we have used is based on what we call *motivational concerns.* These are long term concerns which differ from the usual goals of rational agent systems in that they are not things which the agent acts to achieve, but rather something which the agent is continually monitoring in the background. If a threat or opportunity related to a motivational concern arises, then an emotional reaction is triggered, which leads to behaviour.

Frijda and Swagerman [FS87] postulate emotions as processes which safeguard the long-term persistent goals

Figure 1: PAC allows setting of emotion threshold, base point and rate of decay

emotion(x)	to	goal-situation	by	mode
happy		achieved		
sad		thwarted		
grateful	y	achieved	y	
disappointed		thwarted		expect achieved
guilty	y	thwarted	x	

Figure 2: Some emotions and their causes

or concerns of the agent, such as survival, a desire for stimulation, or a wish to avoid cold and damp. According to them emotion is the process which instantiates appropriate sub-goals to deal with threat or take advantage of opportunity, associated with these concerns.

The earlier work of Toda [Tod82] also postulates emotions as processes which affect the rational system of the agent, and which are based on basic urges. He groups these urges into emergency urges, biological urges, cognitive urges and social urges. These are similar in principle to the long-term concerns of Frijda and Swagerrnan.

In our system it must be specified for a particular scenario what the motivational concerns of the various agents might be (e.g. survival, need for stimulation, concern for children, etc). The conditions which indicate threats to these concerns, or opportunities related to the concerns, must also be defined. Motivational concerns can then be selected for particular agents, resulting in the appropriate triggering events being recognised by the emotion module, leading to instantiation or re-prioritisation of goals.

Each emotion in our system is represented as a linear gauge, with a threshold and a base point which can be varied for each agent (see figure 1). As events happen which can cause a particular emotional reaction, the gauge moves up - at a rate which can be set for each agent. As time passes, if no further events trigger that emotion, the gauge moves back towards the base level, again at a rate which can be varied from agent to agent.

When a particular emotion crosses the threshold on the gauge for that emotion, the emotion is turned on/off. This causes a change to be noted by the cognitive system, where it may affect either the existence or prioritisation of goals.

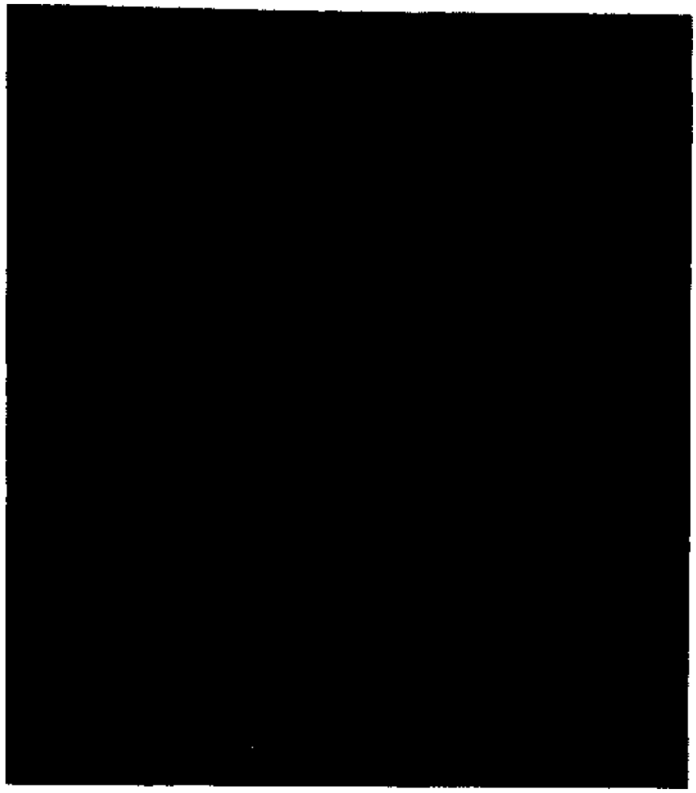

Figure 3: PAC allows editing of plans via access to the dMARS plan editor

3 Personality of Agents

The emotional model used provides three different mechanisms which can be used for modelling agent personality. These are the motivational concerns, the emotion thresholds, and the rate of change for an emotion.

One aspect of personality is a notion of what things are important to that person - a person who is always concerned about money and financial matters could be represented as a person having a motivational concern for financial well-being. This person will respond emotionally to events affecting this aspect of life, and will thus appear to have a different personality than an agent who ignores events associated with threats and opportunities related to this concern. Some motivational concerns will be more or less universal - such as that for stimulation or for survival. However others will be quite individual, and will be a significant aspect of the agent's personality.

The threshold at which an emotion is asserted is also an important aspect of personality. The individual who experiences many anger increasing events, before becoming "angry" has a different personality to the agent who becomes angry easily - perhaps after one or two such events. The change rate for each emotion is similarly important. Two agents with the same threshold for anger, and the same motivational concerns, may still exhibit differing personalities (with respect to this emotion), based on their decay rate for anger. An agent whose anger wears off very slowly has a different personality to the agent whose anger dissipates almost directly. The former personality trait could be described as "the sort of person who bears a grudge".

With these aspects of the emotional model at our disposal, the personality of an agent can then be said to depend on the motivational concerns of that agent plus the thresholds and change rates for each emotion, for that agent. These relatively simple mechanisms, give us a way to begin to represent varying personalities of agents.

4 Agent Architecture

Our underlying agent model is one of interacting aspects of emotion, cognition and physical behaviour. The PAC system allows us to build various scenarios, to set agent characteristics, and to observe the effects of those characteristics on agent "behaviour". Each agent is modeled as having a cognitive component, an emotional, or personality component, and a behavioural component. Each of these components can directly affect each of the other components by means of events or messages.

The cognitive component uses dMars[1], an advanced agent-oriented programming system based on the Belief-Desire-Intention paradigm [RG95], The declarative part of the cognitive component is a set of plans (figure 3 shows editing of one of these plans in PAC), representing the agent's knowledge of how to do things in the world, and a set of beliefs which represent the facts about the world, as the agent knows them.

dMars manages activation of plans to achieve agent goals, with choice of plan being dependent on beliefs regarding world state, as well as on goals and immediate events. A goal and plan together make up an intention - "the agent is intending to fulfill the goal using the particular set of procedures laid out in the plan" [AAI95]. Intentions are prioritised to allow the most important one to be executed first. Execution of the plan steps in an intention is interleaved with monitoring for significant events, to allow the system to remain reactive to the situation. Intentions can be either suspended or aborted if required by a change in circumstances. dMars also has maintenance conditions on plans, allowing representation of conditions that must remain true during the execution of the plan. The maintenance condition becoming false causes the plan to abort.

The emotional component consists of the various emotions represented, and their current state for the agent, as well as the motivational concerns which are active for the agent. Whenever a gauge for an emotion crosses a threshold an event is sent to the cognitive component to allow the fact that the agent now has (or no longer has), the emotion, to affect priorities or choice of plan. Whenever a threat or opportunity for a motivational concern is recognised, events are sent to the cognitive component to ensure that goals are instantiated and/or re-prioritised as appropriate.

The emotional component also directly affects the behavioural component by modifying parameters on particular actions - e.g. a happy agent moves faster, and jumps higher than a sad or depressed agent.

5 System Description

The PAC system provides a menu-based environment which allows users to build up a scenario consisting of objects and agents, to modify various aspects of the agents via menu choices and tools, and then to run the resulting scenario. There are a number of choices of physical agents available including a dog, cat, mouse and cartoon lamp. New physical agents can be added by developing the necessary models in OpenInventor. The system also includes a support tool which allows the user to bring in images from outside, and scale them appropriately with respect to the other objects in the system, before adding them to the object menu.

As agents are added to the scene, the user is able to specify aspects of the behavioural, cognitive and emotional components for each agent. Each graphical agent from the menu has fully specified defaults in each component, but these can also be modified by the user.

At the behavioural level the user can modify the sensing capabilities of the agent. The senses available currently are sight, smell and hearing. Sight is modelled as a cone, radiating from the dog's head. By modifying the distance and angle parameters, the user can modify the visual field of the agent. Smell is modelled as a circle, with modifiable radius, as is hearing. Agents can thus be modelled as having individually different sensing capabilities. It is also possible to modify the agent's basic movement speed.

Agent plans are specified using the high level graphical plan language which is a part of dMars. The menu option in our system, accesses the plan editor which is part of the dMars environment (see figure 3. The user can then edit or create the plans which will enable the agent to do things in the world. Agent beliefs are similarly specified, using dMars belief editor. Agents come with a default set of basic plans for doing things like wandering around, and moving from one point to another. It is likely that plans would need to be written for the particular scenario or set of scenarios.

The emotional and personality aspects of the agents can be modified using two mechanisms. The first is the choice of which motivational concerns the agent has. There are a limited number of these concerns. New motivational concerns can be added to the system in a modular way by writing the appropriate code to monitor for threats and oppportunities, and to set up appropriate goals or reprioritisation procedures when these arise.

The emotional representation for individual agents can be customised by modifying the various settings related to specific emotions such as anger, fear, aggression or pride. The user can set the base level of the emotion for the agent, indicating the stable point for the particular agent; the threshold at which the emotion is switched on/off; and the rate of decay towards the stable point (see figure 1).

We have built a number of sample scenarios using the system. One is a scenario involving two dogs, Max and Fido, who both want to eat from a bowl of food [Tay95]. Fido is customised to be have a more aggressive personality, while Max is more fearful. This results in an

[1] dMars - distributed multi-agent reasoning system, is available from the Australian Artificial Intelligence Institute.

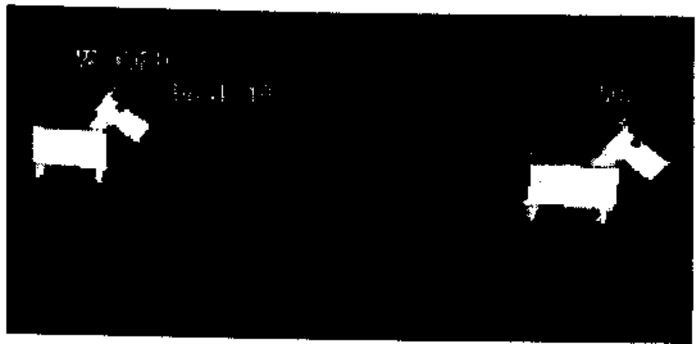

Figure 4: Max and Fido in Dog World

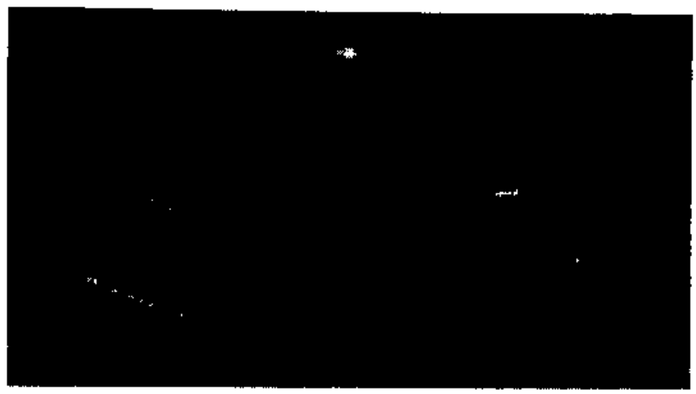

Figure 5: A scenario with parent and child cartoon lamps

interaction where Fido barks at Max, scaring him away from the food (figure 4). Later in the interaction as Max's hunger level builds up, the hunger takes precedence over his fear, and he eats from the bowl despite FidoV aggressive behaviour.

Another scenario we have developed [For96] was inspired by the computer animation, Luxo Jr [Las88] in which a parent and child cartoon lamp interact in a situation where the child plays with an inflatable ball while the parent watches (figure 5). This scenario was driven entirely by the emotional states of the agents, where emotions were caused by a combination of the motivational concerns, and the goal successes and failures. The child lamp had a motivational concern based on stimulation, which was threatened when nothing happened, and recognised an opportunity when a plaything was available. It also had a motivational concern of security which was threatened by being too far from the parent. The parent lamp had a motivational concern based on the child's safety, and one based on the child's happiness. With this very simple representation the system successfully generated some interesting, novel and engaging animations.[2]

Short excerpts of both these initial scenarios are shown on the video.

Acknowledgements

We thank the Australian Artificial Intelligence Institute for provision and technical support of the dMars system. We thank Greg Palstra and Ben Fortuna for programming work and contributions to ideas. We thank Terence Wong, Lawrence Cavedon, Murali Ramakrishna and Ghassan Al-Qaimari for discussions surrounding aspects of the project.

References

[AAI95] The Australian Artificial Intelligence Institute. *The dMARS VI.6.7 System Overview,*

[2]The scenario did not have (or attempt to achieve) anything resembling the graphical quality of Luxo Jr [Las88].

September 1995.

[Bat94a] J. Bates. The role of emotion in believable agents. *Communications of the ACM (Special Issue on Agents),* 1994.

[Bat94b] J. Bates. The role of emotion in believable agents. *Communications of the A CM, Special Issue on Agents,* 1994.

[Dye87] Michael G. Dyer. Emotions and their computations: Three computer models. *Cognition and Emotion,* 1(3), 1987.

[For96] Ben Fortuna. Representation and role of emotion in autonomous agents. Master's thesis, Royal Melbourne Institue of Technology, 1996.

[FS87] Nico H. Frijda and Jaap Swagerman. Can computers feel? theory and design of an emotional system. *Cognition and Emotion,* 1(3), 1987.

[HR95] Barbara Hayes-Roth. Agents on stage: Advancing the art of ai. In *IJCAV95,* Montreal, 1995.

[Las88] J. Lasseter. *Luxo Jr.* Pixar, Los Angeles, 1988.

[OCC88] A. Ortony, G. L. Clore, and A. Collins. *The Cognitive Structure of Emotions.* Cambridge University Press, 1988.

[PT97] Lin Padgham and Guy Taylor. A system for modelling agents having emotion and personality. In Wayne Wobcke Lawrence Cavedon, Anand Rao, editor, *Intelligent Agent Systems: Theoretical and Practical Issues, Lecture Notes in A.I. 1209,* pages 59-71. Springer-Verlag, 1997.

[RG95] A. Rao and M. Georgeff. BDI agents: From theory to practice. In *Proceedings of the First-International Conference on Multiagent Systems,* 1995.

[Tay95] G. Taylor. Animation of believable intelligent agents. Master's thesis, Royal Melbourne Institue of Technology, 1995.

[Tod82] Masanao Toda. *Man, Robot and Society.* Martinus Nijhoff Publishing, Boston, 1982.

DOCTORAL CONSORTIUM ABSTRACTS