Integration of Psychological Models in the Design of Artificial Creatures

Catherine M. Bryant[⋆‡]
[⋆]Department of Computer Science
University of Exeter
Exeter EX4 4QF, U.K.
catherinebryant@care2.com

Gareth J. F. Jones[⋆†]
[†]School of Computing
Dublin City Universtiy
Glasnevin, Dublin 9, Ireland
gjones@computing.dcu.ie

A. J. Wills[‡]
[‡]School of Psychology
University of Exeter
Exeter EX4 4QG, U.K.
A.J.Wills@exeter.ac.uk

Abstract

Artificial creatures form an increasingly important component of interactive computer games. Examples of such creatures exist which can interact with each other and the game player and learn from their experiences. However, we argue, the design of the underlying architecture and algorithms has to a large extent overlooked knowledge from psychology and cognitive sciences. We explore the integration of observations from studies of motivational systems and emotional behaviour into the design of artificial creatures. An initial implementation of our ideas using the "sim_agent" toolkit illustrates that physiological models can be used as the basis for creatures with animal like behaviour attributes. The current aim of this research is to increase the "realism" of artificial creatures in interactive game-play, but it may have wider implications for the development of AI.

1 Introduction

Over the last few decades Artificial Intelligence (AI) has become more than a philosophical consideration or science fiction plot device. With hardware advances it has become possible to incorporate more powerful AI into games as well as increasingly complex graphics and environments. A recent poll of developers showed a sevenfold increase in CPU time used for AI in the average game since 1997 (Johnson, 2002). A large proportion of this interest in AI is in improving the behaviour of NPCs (non-player characters), making them more believable and engaging. It is important to stress the difference between this 'character-based' AI and that in strategic or turn-based games. Isla and Blumberg (2002) elucidate this in a recent paper:

"These latter categories might be considered attempts to codify and emulate high-level logical human thinking. Character-based AI, on the other hand, is an exercise in creating complete brains. Strategic and logical thinking in this type of work usually takes a back seat to issues of low-level perception, reactive behaviour and motor control....work is often rendered with an eye towards recreating life-like behaviour, and emotion modelling and robustness are often also central issues." (2002, p.1)

Essentially 'character-based' AI is a move away from programming an artificial opponent capable of playing against the human mind in intellectual or strategic games such as chess. Rather than refining specific high-level logical thinking, the aim is to capture life-like behaviour and move towards modelling a complete mind. Thus it aims to populate the game environment with agents who act in a realistic and capable manner. Enemy 'bots' in games such as "Quake" or "Half-life" do not need to understand chess or engage in complex reasoning, but they do need to navigate their environment and know when to attack the player. These virtual 'creatures' should be able to perceive and learn about the environment on their own, make decisions, and in some instances interact with other 'creatures' in a limited way.

The applications for this type of AI are becoming increasingly popular in commercial games, and fairly sophisticated designs are emerging. For example Peter Molyneux's game 'Black and White' included creatures with impressive learning and the potential to develop interesting 'personalities' depending on how the player interacted with them. 'Bots' in games such as the "Quake" series need to navigate a 3D environment realistically as well as try to kill the player without being shot in the process. In later incarnations of similar games, for example "Return to Castle Wolfenstein", the bots also interact with each other and can develop limited team-based plans. However at present knowledge from psychology and cognitive sciences about the processes of the mind appears to a large extent to be under used or overlooked in the design of game AI.

This is clearly an interesting area not just in terms of making better games, but in the development of new AI techniques and algorithms. Laird (2002) argues that computer games provide challenging environments and offer many isolated research problems. As the worlds become more realistic, so too must the behaviour from their characters become more complex. Psychologists, in particular those who have worked on animal cognition, have been studying and detailing the behaviours of autonomous creatures in complex environments far longer than AI researchers have been attempting to model them. Yet many designers of 'virtual creatures' seem unaware of recent developments in psychology and how these might be applied. Emotion provides a good example of one such area of research.

Laird mentions that "emotion may be critical to creating the illusion of human behaviour", but seems at a loss how to go about incorporating this - "Unfortunately, there are no comprehensive computational models of how emotions impact with behaviour. What are the triggers for anger? How does anger impact other behaviours?" (Laird (2002), p.4).

Isla and Blumberg (2002) also discuss the modelling of emotions in character-based AI. They point out that much of the work done so far uses emotion as a "diagnostic channel"; a convenient indicator which can be routed from an internal "emotion" value straight to a facial-expression or visual animation. This value is usually derived from a series of expressions to calculate how 'happy', 'sad' or 'angry' the character is feeling. Isla & Blumberg assert that "emotions clearly play a far larger role in our behaviour ... (they) influence the way that we make decisions, the way we think about and plan for the future and even the way we perceive the world" (2002, p. 4). The general approach of Blumberg and other members of the MIT 'synthetic character research group' is that Game AI should be inspired by work from animal learning and psychology. For example they discuss how the Pavlovian conditioning paradigm can be used, and the importance of the character being able to form predictions about the world. With regard to emotions, they discuss their possible application in "action-selection functions", and making exploratory decisions through a "curiosity emotion". However, they make no reference in this case to work done in psychology.

Emotion is certainly very subjective and personal, and at first seems quite inaccessible to the manipulations and measurements of science. However psychologists have been theorising about emotion for over a century. Since William James first tried to define emotion in his 1884 thesis, research has been done to investigate what emotion is, and more importantly if and how it interacts with the rest of our cognitive system. James himself contended that emotions were nothing more than the feelings which accompany bodily responses to a stimuli. Recent work in cognitive neuroscience provides evidence to the contrary: emotions are linked to brain function, to the point that neural systems of emotion and other mental behaviour are interdependent (Gazzaniga, Ivry and Mangun, 2002). The implications of these results are now finding interest in current work in AI. In this work it is important to focus away from the subjective, conscious 'feelings' of emotion and study the underlying systems which give rise to them and their impact on behaviour. Generally, it seems that these systems are heavily involved in reactive mechanisms and learning, and possibly also decision making and attention.

This paper describes our work towards the development of a basic agent architecture which incorporates motivational and emotional elements derived using ideas and findings from psychology to inform the design. In particular this aims to incorporate some emotional mechanisms that have a deep effect on the decision making process.

The remainder of this paper is organised as follows: Section 2 reviews literature on the psychology of animal motivation, Section 3 outlines work from current developments in artificial intelligence, Section 4 describes our working environment, Section 5 introduces the architecture of our artificial creature agents, Section 6 gives some initial results and finally Section 7 draws conclusions from our current study and considers how the work might be extended.

2 Animal Motivation Theories

In this section we explore some key observations from animal motivation theories and their implications for the design of our model for an artificial creature.

2.1 Miller's equilibrium model and the approach-avoid conflict

Generally speaking, animals react to signals they receive from environmental stimuli. Depending on the nature of the stimulus itself and knowledge of past experience with this type of object, the animal will either approach or avoid it. An approach-avoidance conflict occurs when these signals impel an animal towards these two incompatible forms of action.

Gray (1987) notes that conflict of this kind is extremely common. For animals, it is particularly apparent in their behaviour towards a novel object. Novelty is an important stimulus for both eliciting fear (avoidance) and encouraging exploration (approach). In general, animals appear to avoid extremely novel stimuli, but be attracted to ones which are mildly novel.

Experimental psychologist Neal Miller performed a series of studies on the approach-avoid behaviour of rats. The resulting findings allowed him to develop a model which incorporates the various factors involved.

In Miller's basic experimental situation, a rat is trained to run down an alley to get a food reward. However, every time it reaches the goal, it receives a shock. This sets up a conflict situation. Miller observed that the rat ended up oscillating round an equilibrium 'stopping point' a certain distance from the goalbox. The distance of this point from the goal is defined by the strength of the tendencies to approach and avoid the food. The diagram below shows the factors that affect these tendencies and the resulting decision. Miller's model is represented in Figure 1.

Note that the factors include both internal states of the rat as well as external information from the environment and previous experience. Increasing the hunger or de-

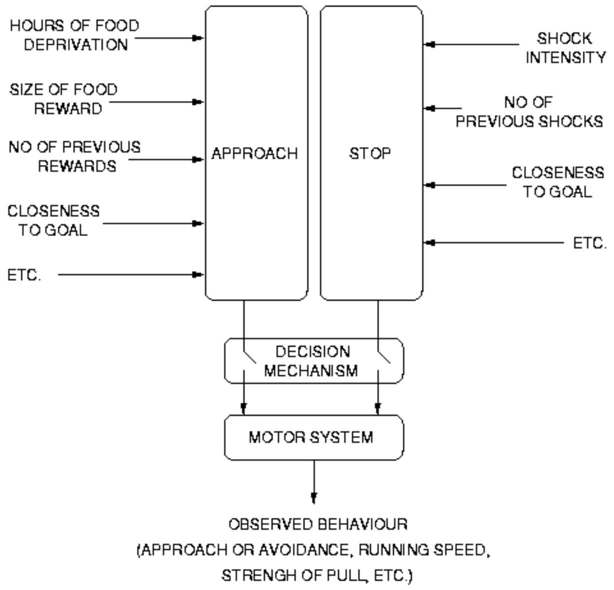

Figure 1: Miller's equilibrium model. (Adapted from Gray (1987), p.142.)

creasing the shock intensity will in turn affect the approach/stop tendencies, and move the equilibrium point closer to the goal. If the approach tendency is much larger than the stop one, you would expect the rat to actually reach the food.

Another point is that 'distance to the goal' is a critical factor in both 'approach' and 'stop' tendencies. However distance cannot affect them in an identical manner: if this was the case then whichever was stronger at the start point would be stronger at the end, resulting in a behaviour where the animal either stops as far as possible from the goal, or completely approaches it.

Work in Miller's (1951, 1959 as cited in Gray 1987) laboratory demonstrated that the strength of the avoidance tendency increases more rapidly with nearness to the goal than that of approach.

Miller noted that there are two main forces behind the tendencies: those that are internal to the animal (such as hunger or other 'drives'), and those relating to the environment and the stimulus itself. They pointed out that there are no internal sources of motivation for the avoidance tendency, and hence it is more purely dependent on environmental factors than the approach tendency. This helps explain why distance has a greater effect on the avoid tendency, especially when near to the goal.

It is clear then that the action towards a certain object is not clear-cut. It is not a simple case of approaching food and avoiding negative objects. Where an animal has learnt to associate pain with an otherwise positive stimulus it may avoid it; conversely if it is hungry enough it will still approach food even if this means receiving a shock.

In terms of programming design, this means that it is wrong to divide the world up into 'good' and 'bad' objects. Instead, every object has the potential to be an overall positive-approach stimulus or a negative-avoid one. It depends not just on the properties of the object, but also what it is associated with and the current internal condition of the animal. This notion of approach-avoid conflicts forms the core of our system design.

2.2 Motivation systems

It is difficult to find one all-inclusive definition of motivation, instead there are various different features which are important to consider.

Firstly, a motivated action differs from a reflex because it is not simply a reaction to an external stimulus. It is also in someway 'driven' by internal states. Teitelbaum (1977, as cited in Toates 1986) argues that "To infer motivation we must break the fixed reflex connection between stimulus and response." Teitelbaum feels that motivation is always directed towards obtaining a certain goal.

Epstein (1982, as cited in Toates 1986) also argues that motivations are complex properties that arise from both external and internal factors. He also considers a third factor: what the animal remembers from past encounters with an incentive object, and the consequence of this encounter.

There are a variety of different models of motivation, of which the simplest is a homeostatic model. Essentially, a homeostatic model is about maintaining essential parameters (e.g. energy level, fluid level) at a near constant 'normal' level. If there is a disturbance then corrective action is taken. Homeostatic mechanisms are driven by 'negative feedback', which can 'switch off' motivation once the deficit has been recovered. The homeostatic model is represented in Figure 2.

According to Grossman (1967, as cited in Toates 1986), there are two types of motivation systems: one which is homeostatic and includes hunger, thirst and other internal factors, while the other is only driven by external factors and includes sex, exploration and aggression.

This dichotomy, however, is too simple, and models developed later do not separate out motivations into these two different types. Homeostatic mechanisms may play a part in explaining the negative-feedback aspects of hunger and thirst, but by themselves are not sufficient as a model. There are other factors to take account of, such as the availability or 'cost' of food - when access to food is made difficult and more energetically costly, animals eat less Toates (1986).

Homeostatic models which look at correcting an energy depletion also do not explain why animals (or indeed people) will overeat if provided with sweet or tasty foods. A final problem is that they do not adequately explain how having a water deficit can then steer an animal towards a water-related goal: in other words they miss the link between the internal state of the animal, and acting towards the external incentives available.

In Bindra's theory (1976, 1978, as cited in Toates 1986), the emphasis is on the role of 'incentive stimuli'

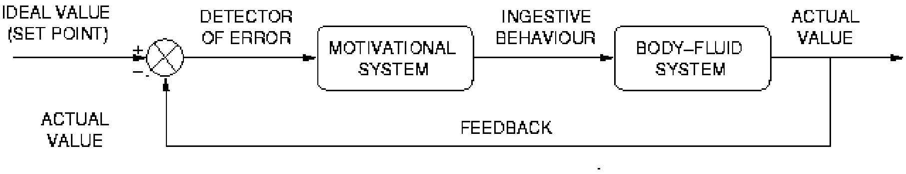

Figure 2: Homeostatic model of motivation. (Adapted from Toates (1986), p.37.)

as well as internal states in the motivation of behaviour. An incentive stimulus is an object or event judged as 'hedonically potent' - one which is affectively positive or negative. This is similar to Miller's approach/avoid tendencies; an animal will react in an appetitive way to hedonically positive incentives, and in an aversive way to negative ones.

Whether a stimulus is seen as hedonically potent depends on various factors, including previous experience with that stimulus as well as physiological states. An animal may assimilate information about a stimulus which it sees as 'neutral'; later on, if the physiological state of the animal changes, that same object could become a positive incentive. For example, an item of food may appear as neutral while the animal is satiated, but once it becomes hungrier that same piece of food becomes a positive incentive which elicits an appetitive reaction.

Bindra develops these ideas into a concept of a 'central motivational state' (c.m.s), which he defines as "a hypothetical set of neural processes that promotes goal-directed actions in relation of particular classes of incentive stimuli" (Bindra, 1974 as cited in Toates 1986).

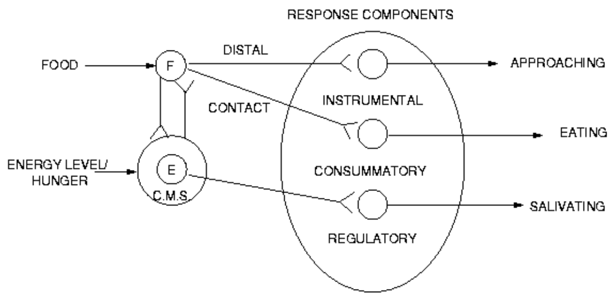

Figure 3: Bindra's model of motivation. The food acts as an incentive stimuli in the feeding motivation system. (Adapted from Toates (1986), p. 43.)

A c.m.s arises from an interaction of 'organismic states' (e.g energy level, testosterone) and the presence of incentive stimuli, see Figure 3. If there are no relevant stimuli present, for example no food when the animal is hungry, then a depletion of energy will not cause systematic goal-directed behaviour. Instead, an increase in general activity may be observed. Also, Toates (1986) notes that novel hedonically neutral stimuli may still arouse some exploration.

In contrast to the homeostatic model, where the internal state drives behaviour, the existence of an incentive stimulus is key. In feeding c.m.s, energy depletion only serves to accentuate the food representation. This explains why tasty and palatable food is sufficient to motivate consumatory behaviour without any kind of energy deprivation.

Thus we can conclude that a homeostatic model is too simplistic for understanding how animals are motivated. All the theories outlined here emphasise a complex interplay between the internal states of the animals with the properties of objects in their external environment. In Bindra's model, an animal cannot just feel motivated to eat because its energy level is depleted - it is only motivated to act in the presence of hedonically potent stimuli. These ideas counter the notion than an animal, once at a certain 'level' of hunger, then sticks rigidly to an explicit goal of 'find food' until its hunger is reduced.

Thus our system needs to include a motivation system which is more flexible than is perhaps usual in existing artificial creatures. The motivation system is a key aspect in that it affects the decision of how the creature should act at each turn in a game.

2.2.1 Toates System theory model of motivation

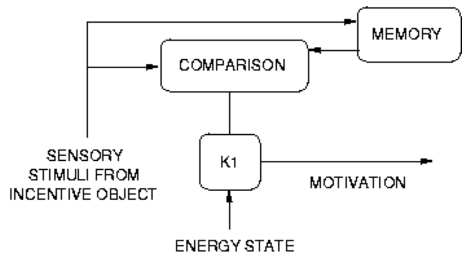

Figure 4: Toates' system theory model. $K1$ represents the energy 'gain' of the system, which determines the level and type of motivation. (Adapted from Toates (1986), p.49.)

Toates (1986) describes his own 'systems theory' model which draws together ideas on motivation similar to Bindra's work. Toates' model is shown in Figure 4. This type of model makes a good bridge from psychological models to computation ones. Toates' model takes account of the three important factors:

- the need for a sensory stimulus to arouse a motivated response.
- the role of the energy level or internal states of the animal in adjusting the 'sensitivity' of the system.

- information from past experiences.

$K1$ represents the 'gain' or sensitivity of the nervous system, and subsequent motivation. If the sensory stimulus 'revives' negative memories of a past experience with this object, it will reduce the value of $K1$. If $K1$ drops to negative numbers this will result in an active avoidance response at the motivation level.

The $K1$ parameter in Toates' model provides a convenient mechanism to encapsulate all the factors involved in motivation in a single number, making the programming of subsequent processes neater. However, it seems likely that there is more to animal motivation systems than described by Toates. Specifically, there is probably a role for emotions, such as fear or pleasure, in motivation and related decision making.

2.3 Emotions

In game AI where emotions have appeared at all, it is generally at a cosmetic level - giving the character the appearance of showing a certain emotion. Here we are concerned not with the subjective feeling or visual appearance of emotions, but rather the underlying mechanisms which give rise to these states.

In this section we review three examples from neuroscience and animal behaviour providing emotional mechanisms that could play a part in the motivation system of our artificial creatures.

2.3.1 Neuroscience and Fear Conditioning

Joseph LeDoux (LeDoux, 1999) identifies two neural routes - one cortical and one subcortical - involved in emotional learning (such as that involved in fear conditioning). The amygdala is a major part of the subcortical route, and removing it prevents fear conditioning from occurring at all. LeDoux suggests that the role of this subcortical route is as a quick-and-dirty reaction mechanism; emotional responses such as fear begin in the amygdala before we even recognise completely what it is we are reacting to.

LeDoux maintains that "Emotion is not just unconscious memory: it exerts a powerful influence on declarative memory and other thought processes." According to Antonio Damasio, one such thought process is that of decision making. He argues that the idea of a totally rational decision maker is not appropriate when quick decisions must be made, and affective memories are invaluable in these cases (Damasio, 1994).

Damasio proposes a "somatic marker hypothesis" which suggests that certain structures in the prefrontal cortex create associations between somatic responses triggered by the amygdala and complex stimuli processed in the cortex. The idea is that both positive and negative associations can be created. Somatic markers help limit the number of possibilities to sort through when making a decision by directing the person away from those associated with negative feelings.

These ideas suggest that not only do affective associations play a part in decision-making, but that there is a physically different route in the brain which processes basic emotional information. In terms of the design of an artificial creature, it would seem sensible to have a similar route, whereby fearful reactions can override more complex processing and steer the animal away from danger.

How do these findings relate to the design of synthetic characters? Firstly, as asserted by LeDoux, whilst consciousness is needed for the subjective feeling of emotion, the basic function of emotional processing and response can be found even in a fruit fly. Thus it seems a possible and useful task to incorporate emotional learning into an AI agent in some way. Since fear conditioning has been extensively studied, it would seem to make a good choice as a place to start. Damasio's hypothesis of 'somatic markers' suggests ways that emotion is important in decision making as well as aspects of learning. It would be interesting to see if basing algorithms around his hypothesis could make for a more 'emotional agent'; one that makes more than completely rational, logical decisions as is generally the case in current game AI. Could this make for a more believable character?

2.3.2 Learning

Toates (1986) notes that when it comes to motivation systems, animals respond to 'primary incentives' (such as food) and 'cues predictive of primary incentives'. In fear-conditioning, animals learn to associate a particular stimulus (e.g. the sound of a bell) with an aversive stimulus such as shock. Once this has occurred, the initial stimulus alone is enough to rouse the animal into a state of fear.

In this way, fear plays a role in animal learning. If a stimulus puts the animal in a state of fear, then its aversive reaction to a subsequent powerful or noisy stimulus is enhanced Toates (1986).

Combined with Damasio's theory, this means that any stimuli occurring while the animal is in a state of fear will be associated more strongly with a negative somatic marker. To replicate this idea, the design of an AI architecture could include a process whereby being in a state of fear affects the strength and type of associations formed by the program.

An advantage of reacting fearfully to cues which predict pain is that the animal will take an appropriate avoidance response before the pain actually occurs.

Gray (1987) explains that rats respond differently in two conditions - receiving a shock, and being exposed to a stimulus that they have learnt predicts a shock occurring. In the first condition, there is a great increase in activity, frantic scampering, or attacking some feature of the environment. In contrast, encountering a stimuli which predicts shock results in the rat freezing. Gray suggests this is an adaptive response that occurs when a rat spots a predator - it freezes in an attempt to avoid detection. He also adds that the response is affected by distance - if the stimulus (or predator) gets too close, the rat shows a strong aversive reaction.

By incorporating fear appropriately into learning and decision mechanisms, an approach to AI could be developed that responds pre-emptively rather than just reactively to pain. Also, the priming effect of fear on forming associations may result in a program which learns to avoid painful situations more efficiently than one with no fear.

2.3.3 The Role of Pleasure

Emotions can also impact animal behaviour to support positive behaviour. For example, there is the concept of a 'positive feedback' priming mechanism that helps to sustain certain activities. Evidence for this was found by McFarland and McFarland (1968, in (Toates, 1986)). They noticed that interrupting doves while they were drinking caused them to 'lose momentum'. This implies that there was something about drinking itself that increased the motivational state of the dove. Toates (1986, p. 116) explains that an animal needs such a positive feedback effect, particularly in situations where simultaneous feeding and drinking tendencies exist of almost identical strength. If it decides to eat and only negative feedback exists, then after the first couple of mouthfuls the feeding motivation will drop, in turn making the drinking tendency stronger. The animal would end up oscillating between food and water, which is costly in terms of time and energy. It would be more advantageous to stick with one activity for a longer period of time before switching.

It would seem vital to have some kind of positive feedback mechanism to reduce the chance of the AI oscillating, and hence to look more believable as well as being more efficient. While the animal motivation literature does not discuss pleasure as such, this concept makes at least a good metaphor for the 'positive feedback' concept. It would make sense that the animal would feel something good when it starts eating or drinking. Essentially, pleasure can be thought of as a reward from an internal, rather than external, origin. Finally, in the same way that the fear emotion might enhance learning about dangerous objects, it would seem a good idea to have a similar 'emotion' which affects the learning about really positive objects or encounters.

3 Artificial Intelligence

In order to make use of the ideas from the previous section, we need to consider what sort of design and framework would be conducive to the incorporation of emotional processes. Despite the lack of sophisticated emotional agents in modern computer games, emotions in general are not a new topic for AI. For example, Simon (1967) had already explored the need to account for 'alarm mechanisms' in artificial systems.

Since the 1980s, many different programs have been specified and sometimes implemented. One of the most notable examples in this area is the work of Sloman (Sloman, 1999)(Sloman, 2000) (Sloman, 2001). He argues for more sophisticated theories of affect and emotion, and has suggested an architecture-based approach to the design of affective agents. This means starting with specifications of architectures for complete agents, and then finding out what sorts of states and processes are supported by those architectures. Sloman himself specifies a multi-level 'CogAff Architecture Schema' (Sloman, 2001) in which 'affective' states and processes "can be defined in terms of the various types of information processing and control states supported by different variants of the architecture, in which different subsets of the architecture are present."

It interesting to note that Sloman has severe objections to Damasio's hypothesis and does not believe that "emotions somehow contribute to intelligence: rather they are a side-effect of mechanisms that are required for other reasons." Despite the debate over emotions and intelligence, Sloman's work is still consistent with that of LeDoux and neuroscience in general. For example, the 'reactive layer' in his architecture which monitors automatic responses is similar to the direct activation of the amygdala from the sensory thalamus e.g in fear conditioning. His 'deliberative', reasoning layer is equivalent to the slower reasoning performed in the cortex. The 'meta management' layer, for monitoring internal states and processes is a little more tricky to pinpoint, however LeDoux (1999) identifies neural systems which may support the awareness of the activity of bodily responses.

Work done by Moffat (2001) 'on the positive value of affect' also draws on psychology to improve AI performance, and provides more inspiration for the relevance of emotion. Moffat feels that cognitive psychologists tend to focus on the function of negative emotions (such as fear), but positive emotions are also important, particularly in learning. On the other hand, machine 'learning classifier systems' (LCSs) model reward and not punishment. 'EMMA', the model resulting from attempts to combine positive and negative affect, was found to learn certain behaviours better than the LCSs. More importantly, Moffat found that the 'emotions' provided a way of signifying importance to EMMA:

"LCSs do not distinguish between stimuli of varying priorities.... EMMA devotes her attention and all her resources to the most important aspect of her current situation. In this respect, emotion is a kind of biological optimiser that could be put to good use in artificial agents too; especially learning ones" (Moffat, 2001), p.61.

Moffat's work suggests the importance of incorporating negative and positive affect. Our work adopts an archictecture-based model as advocated by Sloman. This means rather than trying to code specific behaviours and abilities as they are needed, the starting point is to specify an architecture for a complete agent, and investigate which processes are supported by that architecture.

4 Agent and Game Design

In this section we outline the "sim_agent" toolkit used to implement our prototype agent system, and the design of a simple game framework to explore agent behaviour.

4.1 Programming Environment

The "sim_agent" toolkit developed by the 'Cognition and Affect project' at University of Birmingham, is designed with the specific intention of enabling the building of agent architectures[1]. It runs using the Pop-11 language within the POPLOG environment, on both Linux and Windows systems. Sim_agent was chosen for our work since it allows a wide range of programming techniques, and for the possibility of hybrid systems, for example incorporating neural networks.

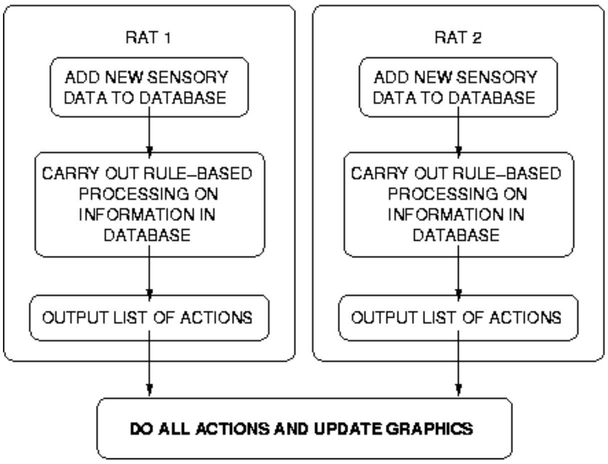

Figure 5: For each 'time-slice', the sim_agent Scheduler runs through processes for each agent. After this is complete, the Scheduler executes any actions, such as moving the agents to a new location, and updates the graphics accordingly.

Figure 5 shows the operation of the sim_agent toolkit. Time is simulated in discrete 'time-slices', which effectively act as a counter. This means that time is not truly continuous, and that the agents all act in a synchronous way. During each time slice, the agent does the following:

- New sensory data is added to the agent's personal database.
- Next, its rulesystem runs, acting on the information available in the agent's database. Unless the agent is going to do nothing during this time-slice, the rulesystem will output one or more 'do X' items into the agent's database.
- The scheduler moves on to any other agents or objects that exist in the environment, and repeats the procedure. When this is finished, it goes back and 'picks up' all the 'do' actions, and executes them.

[1]Details available from: `http://www.cs.bham.ac.uk/~axs/cog\_affect/sim\_agent.html`

4.2 Game Design

A simple game was designed to explore our approach to programming artificial creatures for computer games. This incorporates a set of 'Rat' agents, two sets of 'Rat' agents were designed, one with 'emotional mechanisms' involving fear and pleasure, and the other without. The aim then is to ask participants to play two different versions of the game, taking objective measures of the Rat's performance and a subjective measure of which version the participant thought was more believable.

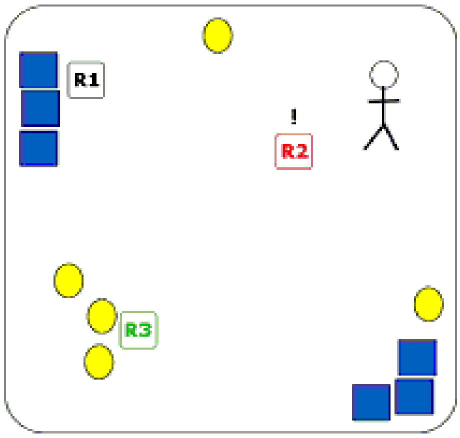

Figure 6: Concept diagram showing typical graphics for the game.

Rats will be implemented in sim_agent, and consist of a 'hunger level', 'thirst level', 'speed', and a 'heading' (direction). The Rat also has a value expressing its current emotional state (fear, pleasure or neutral), and a flag for being in pain or not. 'Food', 'water' and 'person' are all created as objects, of which the game player can move only the food and person.

The idea of the game from the player's point of view is to score points by shocking Rat agents. It uses a turn-based system, whereby the Rats all make an action choice and move, then the player takes a turn.

The aim of the Rats is to basically stay alive, by keeping their hunger and thirst levels relatively low. They have a simple learning system whereby they can form associations between objects which occur together in space, and events that occur together in time. They start off knowing nothing about the player. In other words they have no 'instinctive fear'. Also, the Rats do not immediately understand that a received shock is related to the person - this is something they should learn to associate over time.

Shocking a Rat puts it into a state of pain. In Rat agents with emotions, it also puts them into a state of fear. Both these affect the processing of the Rat during its subsequent turn.

Each turn, the player can move the person within a certain distance, then has the option to shock up to one Rat, if that Rat is 'in range' of the shocking device which the person carries. The Rat cannot discern the direction that the shock came from; instead it decides which object is the most likely 'cause' of the pain, based upon the associations stored in its memory. Note that the range of the shocking device is greater than the visual range of the Rat. This means it is possible to shock the Rat without it seeing the person at all. If the Rat cannot decide where the shock came from it will react differently; perhaps running in a random direction as opposed to freezing or actively avoiding the object it links with causing pain. We hope that this feature will make the Rat appear more believable.

The player also has the option of moving one piece of food around, within a certain distance. This ensures not only a more dynamic environment, but opens up a few more strategies to the player, such as piling all the food together in one place and standing the person next to it.

Rats that feel fear should learn more quickly that the person is associated with pain. This is because being in a state of fear enhances the memory updating and associations involved with pain and objects that might be causing it. Secondly, it is possible for Rats to feel fear at certain objects before they are actually in pain. This should help them pre-empt the shock and hopefully avoid the feared object before it causes pain.

The role of the pleasure emotion is slightly more subtle. It occurs when the Rat starts eating or drinking; to a greater extent the more hungry or thirsty it is. It provides a positive feedback mechanism, which will encourage the Rat to continue consuming until its hunger/thirst level drops quite low. This aim here is to prevent the Rat from 'oscillating' between food and water objects if its hunger and thirst levels are at similar values.

Both emotions are continuous, occurring at certain levels rather than being simply on or off. This allows for some more complex possibilities, such as a situation where the Rat feels a little bit fearful but very hungry; so it approaches the food despite being slightly afraid of it.

While we have a complete design of the architecture for the game, its implementation is incomplete. The system currently does not incorporate interaction with a user, and the memory and emotion systems are not yet functional.

5 Architecture Overview

5.1 Basic Framework

Figure 7 shows the architecture of the Rat agent. The currently implemented basic design is shaded grey. This includes the core decision-making aspect, and the motivation systems. Running from top to the bottom is roughly equivalent to the order of the sim_agent rulesystem run by each agent during the cycle.

Perceptual system This identifies what the object is, along with other properties such as how far away it is, how much there is, and in the case of food/drink a 'hedonic' value representing how 'tasty' or desirable it is. Any information about objects recognised as food will be passed on to the feeding motivation system, and the details of drink objects filtered to the drinking motivation system. At this stage any other objects, such as Rats or perhaps the human player are not processed further.

Motivation systems Here a value for each object is calculated. The value represents an overall 'weight' of importance. It takes account of the properties of the individual item, and how far away it is, along with specific information on the internal condition of the Rat. The Feeding motivation system uses the Rat's hunger value, while the drinking systems uses the thirst value. (Hunger does not affect the drinking motivation system.) An equation for this is as follows:

$$\text{Weight} = \frac{a \times \text{Hunger} + b \times \text{Amount}}{c \times \text{Distance}} + d \times \text{Hedonic Value}$$

A weight value is computed for each object, along with an appropriate action. If the weight value is positive, then the action will be to approach the object; if it is negative then the suggestion will be to move away from it (particularly unpleasant food i.e. with a large negative hedonic value, might be aversive). If the Rat is currently consuming the object, the weight will represent how important it is to carry on doing so.

Finally, if there are no food objects going into the feeding system, it will output an 'explore' action, with a weight evaluated using the Rat's current hunger level as the main variable.

Decision The decision mechanism simply chooses whichever action has the highest 'weight' associated with it. However, it could be more complex than this - taking account of what other objects lie in the same direction. So a good decision might be to go towards a mediocre item of food if there also happens to be some water nearby. Conversely, if a great item of food is very close to a dangerous object it might be better to avoid that direction.

Motor system The processes here figure out how far the Rat can move in the chosen direction, and evaluates the new co-ordinates to be put out as a movement action.

5.2 Full Version

The Full Architecture design shown in Figure 7 includes two important additions to the basic version: memory and emotion systems.

Memory This stores locations of objects which the Rat encounters, and includes a simple learning mechanism

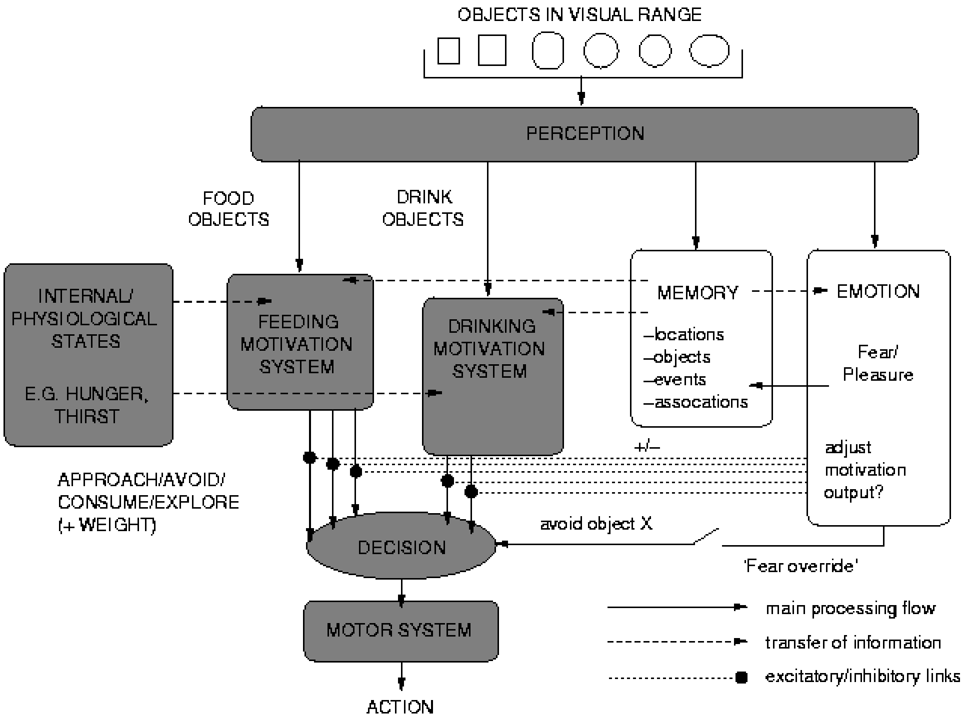

Figure 7: Rat Agent architecture design. The implemented base design is shaded grey.

which can develop conditional associations between objects and events which occur together in space or time, in additional to unconditional ones arising from the unconditional stimulus of the object. It provides extra detail to the motivation systems, so their evaluation equation can also take account of any past experience with the object.

The object memory does not remember food items as 'specific' e.g 'food item one', but instead stores food by location e.g 'food at (x,y)'.

Emotion This does several things, but all the actions essentially involve fear and pleasure. Firstly, it cross references incoming visual information with details in the memory to see if any objects should elicit a state of fear, and if so then what level of fear. The level relates directly to the strength of the association between that object and being in pain.

In terms of pleasure, at the moment it only produces this state if the Rat is actually consuming, however this could be extended to an anticipative pleasure. The level of pleasure is determined by how hungry the Rat is. So if it is really hungry before it starts eating, the level of pleasure will be high. In a sense 'pleasure' here can also be thought of as 'relief'.

The emotion system can adjust the weight values produced by the motivation systems to enhance or reduce particular signals. As an example, if one of the food objects is associated with something nasty the feeding motivation system may output a negative 'avoid' signal for that object. If it is particularly nasty - enough to cause some degree of fear - the emotion system will enhance the signal, making it particularly aversive, while decreasing the strength of all the other signals.

It is important to note that in this situation the emotion system does not necessarily get the last word - if the rat is especially thirsty, one of the 'approach water object' signals might still be greater than the avoidance one. However, if recognising an object pushes fear above a certain threshold, an override happens; the rat will run from that object despite how hungry or thirsty it might be. This route is approximately similar to the 'quick and dirty' fear reaction mechanism discussed by neuroscientists.

If the Rat is feeling pleasure at consuming an object, the emotion system will also adjust the weights, increasing the consume signal while decreasing the others. The amount that the signals are altered will relate directly to the level of emotion - a higher level resulting in a greater signal adjustment.

Feeling either emotion to any level will also feed back into the memory system, enhancing specific associations formed or reinforced during that cycle. In particular if the Rat was in a state of fear because it could see the player, and then subsequently experienced a painful shock, the association between the player and pain would be strengthened to a greater level than if the Rat was in a neutral emotional condition.

5.3 System Implementation Details

While it is often comparatively easy to specify the desired features and behaviour of a system, actually encoding these into a working agent is often much more difficult. In this section we discuss our current implementation of the systems within the Rat and how these might be extended.

5.3.1 Hunger/Thirst Systems

After some consideration the following relationship was used to calculate the hunger and thirst values in each cycle.

$$Y = \frac{2x - 1}{(2x - 1)^2 + 1} + 0.5$$

where Y is the new Hunger or Thirst value and x represents a counter which increments each cycle. It is a fairly arbitary choice, and could be replaced with an equation (or indeed series of equations) which more accurately reflect how hunger changes in a real animal.

This function was chosen since it increases slowly, indicating that the Rat's hunger/thirst level rises slowly at first, but then increases rapidly to a point where it is 'very hungry', with the limiting value $Y = 1.0$ leading to death of the Rat from starvation. This function is not taken from any particular animal psychology literature, but is based on intuition of the relationship between hunger/thirst and time. During each run of the Rat agent the hunger and thirst levels are updated.

It would be good if food of a higher 'quality' actually reduced their hunger by more - in other words there would be some real benefit in going for these type of objects. This is one of the many ideas which could relatively easily be added into the program in the future.

5.3.2 Motivation Systems

The feeding and drinking systems are identical, and we describe only the feeding system here.

The purpose of the motivation system is to process the relevant visual information and output a database entry for each object determining the most appropriate action.

'Food Weight' FW is calculated using the following equation,

$$\text{Food Weight} \propto \frac{H^2}{D} + FQ$$

where H is the hunger, D is the distance to the food, and FQ the food quality. The hunger value is squared so that the resulting weight is exponentially greater at high levels of hunger. FW is proportional to $1/D$, resulting in lower weights with greater distances between the Rat the the food. Food quality is added to the end to provide a final adjustment. If it is negative, it may push the resulting weight to negative values and a subsequent 'avoid' action. The constants in the equation were derived from trial-and-error testing until the Rats behaved in a reasonably balanced way.

At the end of the day, the motivation equation is key to the decisions made by the Rat, and behaviour may be further improved by use of alternative functions. Another option would be to use a genetic algorithm approach to try to 'evolve' an optimal equation that produces the most 'fit' Rats. Fitness could be simply a survival rate, or relate to how well the Rat maintains a balanced level of hunger and thirst.

5.3.3 Explore System

If there is no visual data on food objects available, the system outputs an explore action.

The 'Explore Weight' EW is calculated using the following method,

$$EW \propto \text{need}^2$$

where need is the current hunger level of the Rat.

Again, this is another equation that could benefit from being 'evolved' by genetic algorithms. At the moment it is roughly balanced so as to become more urgent to find food the hungrier the Rat becomes, but at lower levels of hunger it's still better to carry on drinking if drink is available.

Again the exploratory mechanism is not based on psychological literature, but in its current intuitive form merely ensures that the Rat moves to locate sources of food and drink. There is considerable existing work on animal foraging patterns that could be applied here.

The following exploration method was developed using trial and error experimentation. The explore action has the potential to span up to 6 turns, during which the Rat does the following:

Turn	Action	Count
1	Choose random direction X, move that way.	1
2	Continue to move in X direction	2
3	Continue to move in X direction	3
4	Reverse direction X, move that way	4
5	Continue to move in (reversed) X direction	5
6	Continue to move in (reversed) X direction	6
7	Back in starting position, choose direction Y	1

This means that the Rat spends 3 turns moving in one direction, at which point it turns round and goes back to the starting position. If in any of these turns it encounters food/drink then it reacts to those objects: in other words it is not 'committed' to completing the exploration sequence.

When it comes to step 6, a new angle for exploration is chosen. Essentially the new angle cannot be anywhere within the range of the old one, plus or minus 45 degrees. This makes sure that after an unsuccessful exploration in one direction, the Rat chooses a significantly different direction to explore in next.

6 Results

Figure 8 shows a series of images showing the progress of Rat agents. The frame number refers to the time-slice at which the snapshot was taken. The rats are the square

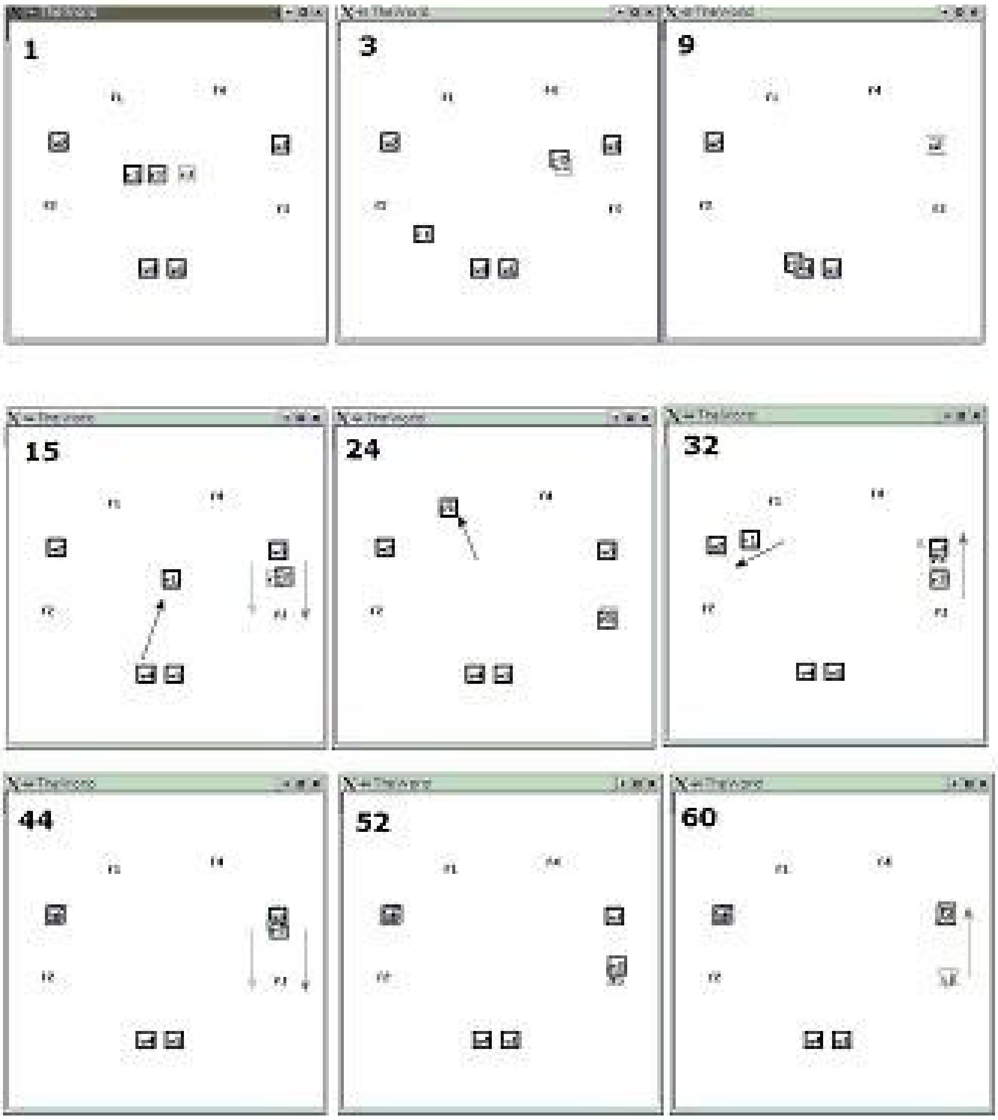

Figure 8: Images showing an example of Rat agent progress.

boxes in the centre of frame (1). They all start off with hunger and thirst at the same low level in all of the results discussed. This is probably why r2 and r3 head towards the same water object (the square boxes) at the beginning.

At cycle 3 R1 can be seen to be fairly close to both a food item (the circles) and a water item. The water item it heads towards has the highest 'quality' value of the objects in the environment, so this move makes sense. After drinking for a bit, the food motivation system pushes him to explore (about cycle 12/13). He finds food and consumes this for about 5 cycles then makes his way to the nearby water object. At this point though he hits a bug whereby no matter how much he drinks the thirst does not go down. By cycle 52 he is dead from hunger.

Rats R2 and R3 essentially oscillate between the food and water objects on the far right.

While the motivation equations could do with some adjusting, it is still good to see that the agents make some attempt to keep their hunger and thirst levels low. Also it is good to see the inefficient oscillating behaviour occurring as predicted. Including the 'pleasure emotion' may really help to reduce this.

7 Discussion

Although the implementation of the Rat agent architecture is not complete, some conclusions can be drawn at this stage. The work completed so far is very promising, and we are confident that developing it further would result in some very interesting results.

Firstly, there is a wealth of psychology literature which makes for good source material and inspiration. The work described here focuses on motivation and emotion systems. However, there is much more information and theory available than has been incorporated in this design.

Many of the ideas described in this literature are not ones that are typically explored when considering problems purely from an AI point of view.

The biggest advantage of considering animal motivation studies is that the researchers spent a lot of time observing and testing the animals, and really getting to grips with the basic systems that drive and affect behaviour. Regardless of whether their findings accurately explain how the animal mind really works, their descriptions still relate strongly to real observable actions. The resulting models and diagrams make it fairly simple to port the ideas over to a computing environment.

It is encouraging to see that even the basic version resulted in agents that made appropriate decisions to reduce their hunger and thirst levels. Their behaviour was always slightly unpredictable (and hence, perhaps more believable?) since they never followed a 'set path' or 'set procedure'. It was not a case of 'when hunger is X, find food'.

The architecture-based approach lends itself well to the approach taken in this work. Once the basic architecture design was in place, other aspects could integrated in quite a natural way. For example, once the base motivation system was in place, it was fairly straightforward to see how the fear emotion could be incorporated and affect the decision making.

In the games industry, it is becoming more common to use pre-designed 'engines' to cover whole aspects of the coding. These engines tend to be specialised, for example it is possible to get physics engines that deal specifically with car crashes. Considering how complex just designing the motivation, or learning, or perceptual system can be it would seem a good idea to put them together as an AI creature 'engine'. This could be the basic all-purpose agent which could then be tweaked and adapted by the specific game designers to suit their needs. Using a system like sim_agent would be perfect for this, since it is easy to adjust old rulesets, add in new ones, or simply change the base variables for the agent instance. The architecture and design ideas presented here could form a component of such an engine.

To really achieve this effectively, it may be necessary to bring together a hybrid of AI techniques. In this investigation we saw how difficult it is to know what functions to use to provide the most efficient and realistic behaviour. This is exactly the type of problem that genetic algorithms could help with. Neural networks, or at least a connectionist approach, seem like the best strategy for implementing learning systems. However, without being implemented in a way that makes them useable inside the symbolic environment of game code they are not too practical. Both these areas would provide good grounds for further study.

Overall, we are encouraged by our results. They demonstrate that psychology literature is a very fruitful resource. If a complete AI engine which has been inspired by psychology in is developed, we feel that it would indeed create more believable agents and much more immersive game play.

References

G. Colombetti. The Somato-Cognitive System. Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing, University of York, pp21-28, 2001.

A. Damasio. *Descartes' Error: Emotion, Reason, and the Human Brain*, Putnam, 1994.

M. Gazzaniga, R. Ivry and G. Mangun *Cognitive Neuroscience*, W.W. Norton & Company, Inc., 2002.

J. Gray. *The Prychology of Fear and Stress*, (2nd edition) Cambridge, 1987.

D. Isla and B. Blumberg. New Challenges for Character-Based AI for Games. Proceedings of the AAAI Spring Symposium on AI and Interactive Entertainment, Stanford, CA, 2002.

S. Johnson. Wild Things. *Wired* (issue 10.03), 2002.

J.Laird. Design Goals for Autonomous Synthetic Characters. Retrieved from `http://ai.eecs.umich.edu/people/laird/papers/AAAI-SS00.pdf`.

J. LeDoux. *The Emotional Brain*, Phoenix, 1999.

D. Moffat. On the Positive Value of Affect. Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing, University of York, pp58-62, 2001.

H. A. Simon. Motivational and Emotional Controls of Cognition. In *Models of Thought*, pp29-38. Yale University Press, New Haven, 1967.

A. Sloman. What Sort of Architecture is Required for a Human-Like Agent? In M. Wooldridge and A. Rao, editors, *Foundations of Rational Agency*, pp35-52, Kluwer Academic, 1999.

A. Sloman. Architectural Requirments for Human-Like Agents Both Natural and Artificial (what sorts of machines can love?). In K. Dautenhahn, editor, *Human Cognition And Social Agent Technology, Advances in Consciousnes Research*, pp163-195, John Benjamins, 2000.

A. Sloman. Varieties of Affect and the CogAff Architecture Schema. Proceedings of the AISB'01 Symposium on Emotion, Cognition and Affective Computing, University of York, pp39-38, 2001.

F. Toates. *Motivational Systems*, Cambridge University Press, 1986.