Do Virtual Sheep Smell Emotion ?

Carlos Delgado, Ruth Aylett
Centre for Virtual Environments
The University of Salford
Business House, University Road
Salford, M5 4WT, United Kingdom
Tel: +44 161 295 2918
Fax: +44 161 295 2925
C.Delgado@pgr.salford.ac.uk
R.S.Aylett@salford.ac.uk

Abstract

In this paper we describe an emotional- behavioural architecture. The emotional engine is a higher layer than the behaviour system, and it can alter behaviour patterns, the engines is designed to simulate Emotionally-Intelligent Agents in a Virtual Environment, where each agent senses its own emotions, and other creature emotions through a virtual smell sensor; senses obstacles and other moving creatures in the environment and reacts to them. The architecture consists of an emotion engine, behaviour synthesis system, a motor layer and a library of sensors.

KEYWORDS: autonomous agents, multiple agents, emotion, virtual environment, behavioural architecture, virtual sensors, virtual smell.

1 Introduction

This paper discusses work-in-progress in the use of emotion as a form of interaction between (pseudo-)embodied agents in a virtual environment (VE). It attempts to integrate emotions at the pseudo-physiological level into an existing behavioural architecture which will be briefly described. It then seeks mechanisms for the transmission of emotion between agents, and for the perceived emotion of one agent to influence the emotion and thence the behaviour of another. One branch of this work confines itself to the behavioural level, taking sheep as exemplar agents, and it this we will discuss here. The other branch considers how a planning system might be integrated into the architecture and takes the idea of the Holodeck, a soap-opera type drama played out by virtual pseudo-human agents, as its test-bed.

While emotional interaction between agents naturally requires an internal emotional architecture, this architecture has to be linked both to transmission and reception mechanisms. Thus the focus is somewhat different from much existing work which either concentrates on the internal agent architecture or considers the agent-human user relationship. As discussed below, transmission and reception are asymmetric: transmission is the sole step in communicating emotion to another agent, but reception may also involve an interaction between the sensed emotion and the existing emotional state of the agent before feeding into its behaviour. Note that we do not include intentionality in this system - emotional interaction is being modelled as essentially involuntary.

2 Behavioural architecture - previous work

Previous work had taken a behavioural architecture developed for multiple cooperating robots - the Behavioural Synthesis Architecture or BSA [4] - and reapplied it to agents in a virtual environment (VE) in the Virtual Teletubbies project [2]. We here give an overview of this architecture. The BSA incorporated three structures at increasing levels of abstraction: behaviour patterns, behaviour packets, and behaviour scripts.

2.1 Behaviour patterns

At the most primitive level, a behaviour pattern, (**bp**) as illustrated in figure 1, was defined as a pair of functional mappings, one from incoming sensory stimulus to outgoing desired motor response, and the other from sensory stimulus to utility, a mapping defining the importance of the motor response for the given level of stimulus. An agent possesses a repertoire of behaviour patterns, with each active pattern at any given time proposing its desired motor response according to its current sensory input. These responses were weighted by their utility values and synthesised together to produce an emergent response, which was the actual behaviour of the agent. Thus second-to-second variation in emergent behaviour was dealt with via weighted synthesis on a continuous basis, unlike the time-sliced Brooksian architecture.

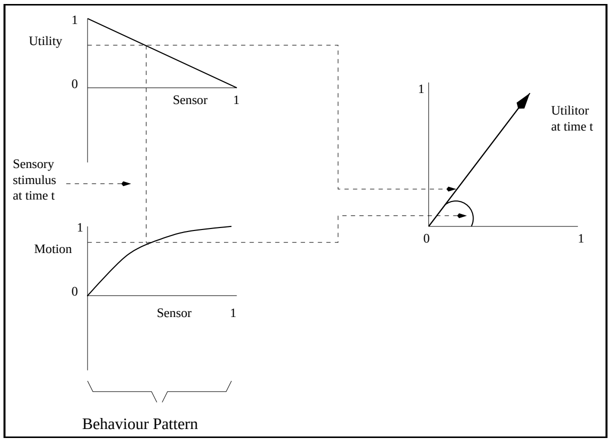

Figure 1: Behaviour pattern example

Consider the situation where the sensory stimulus relates to an agent's forward facing distance-to-obstacle measuring sensor and the associated motion response relates to the forward translate velocity for that agent. From Figure 1 it can be seen that as the agent gets nearer to the object then its forward translate velocity will be reduced to zero. At the same time, the associated utility for this motion response increases. A similar pair of functions for rotation produces an increasing rotation away from the obstacle, also with increasing utility, as its distance from it deminishes. Thus as the agent gets nearer to an object in its path, it becomes more important for the agent to slow down and turn away from it.

In the BSA four behaviour levels (often known as strategy levels) were identified, originally for purely conceptual convenience [3]. A Self level contains those **bp**s concerned with the maximisation and replenishment of internal resources, e.g. making sure an agent does not go hungry, or not walking up hills when energy is low An Environment level contains **bp**s associated with activities involving objects within the agent's environment, e.g. collision avoidance, collision detection, playing with a toy.

A Species level contains those **bp**s associated with cooperant activities e.g. two agents carrying an object together, two agents hugging each other. A Universe level contains **bp**s specific to a particular task, e.g. navigating to the initial location of an object to be relocated (such as the location of a slice of bread for example), then subsequent navigation to the desired goal location (the toaster in this case).

These strategy levels became an implementational feature as discussed below when the architecture was moved from robots to virtual agents.

2.2 Behaviour packets

If all the **bp**s in an agent's repertoire were active at the same time then the overall emergent behaviour of the agent might be of little value. For example patterns designed to produce obstacle avoidance as described above are not useful if you want an agent to sit down on a chair or hug another one of its species. The **bp** designer must always bear in mind that the low-level architecture is sensor-driven, and not task or even sub-task dependent. What is needed in this case is an automatic mechanism for deactivating the 'obstacle avoidance' **bp**s when the 'sit' **bp**s or 'hugging' **bp**s are active. Associated therefore with every **bp** within an agent is an 'active flag', which enables or disables it. Thus obstacle avoidance **bp**s for example can be turned off and on when required. A **bp** is 'deactivated' in the BSA by forcing the respective utility to zero. The action effectively produces a **bp** of zero importance and hence one which does not contribute to the overall emergent behaviour of the agent.

This mechanism is applied by grouping together **bp**s in goal-achieving sets known as behaviour packets. A behaviour packet is a small data structure which includes a sensory pre-condition for activating the **bp**s it references, and a sensory post-condition which controls deactivation of the named **bp**s. Behaviour packets show some similarity with AI production rules [7], though they work at the sub-symbolic level and are driven by incoming sensor data rather than by an inferencing system. They support behavioural sequencing for agents performing at a task (universe) behaviour level. Thus a sensory pre-condition of 'being near the chair' could be used to move from a beahviour packet in which obstacle avoidance **bp**s were active to one in which they are not.

Thus behaviour packets provide a mechanism for contextually sensitive behaviour switching, which is seen as a more flexible mechanism than the finit-state machine definition of inhibition and exitation between behaviours of the subsumption architecture.

2.3 Behaviour Script: high-level sequencing and agent drives

A behaviour script is simply a set of behaviour packets assembled for the achievement of a particular task, using the sensory rep-and post-conditions of Figure 2. The original approach was to generate behaviour scripts on the fly using a reflective agent incorporating a symbolic AI planner, and then send the individual scripts to behavioural-based agents. This hybrid approach was taken with the co-operative robots in MACTA [1] and is appropriate where the domain is predominantly task-based. However, while the lower levels of the architecture were moved from robots to virtual agents with no change, virtual agents are less likely to live in a task-oriented environment. It was at this point that the issue of incorporating emotion into the architecture came to the fore since if behaviour is not task-directed then drives or emotions are required instead. Broadly, if an agent is not being told what to do, then it does 'what it likes'. Thus a sequencing engine linking behaviour to internal motivations or drives was developed as seen in Figure 3. This was applied in the Virtual Teletubbies project [2].

A set of drives were developed for the virtual agents - hunger, excitability, happiness, curiosity, and sleep - comparable to the homeostatic variables found in work of Bruce Blumberg [5]. These drives play the same role of contextually-driven behaviour switching that had previously been played by an AI planner. The framework developed contains four queues, one for each of the conceptual categories self, species, environment and universe discussed earlier. The entries in this queue consist of groups containing one or more behaviour packets, effectively sub-scripts known as behaviour scriptlets, each with an attached priority. The priority is generated automatically and is typically related to a predetermined threshold level of a drive, so the more hungry an agent becomes, the greater the priority. The scriptlet with the highest priority is then selected for packet execution. Thus if one adds fear to the set of emotions/drives, a certain level of fear may trigger panic behaviour in sheep which up until then have been quietly eating grass.

The behaviour-sequencing engine although always processing stimuli might not be executing a scriptlet and therefore has a default script at the environment level. The default script executes a single packet containing **bp**s that effectively lets the low-level module handle wandering in the environment while avoiding obstacles. In the case of sheep this might be modified to grazing behaviour. The default script is changed when another sensory precondition from another set of packets is met. This is typically at another strategy level so if a agent sensed the presence of another agent this could trigger the behaviour at the species level.

It should be clear that to sequence behaviour, it is enough to model emotion/drives as a meter with a threshold above which switching occurs, though this is of course not at all biologically plausible. The level of the drive is itself determined by either perception and actuation in most cases: curiousity rises when a new object is seen, hunger rises more quickly if an agent is more active.

3 Emotion as a behaviour pattern or pattern modifier

The use of emotion/drives to switch behavioural mode is however only one way of dealing with the relationship between emotion and behaviour. The architecture also supports a much lower level relationship at the level of the synthesis mechanism. Two approaches are possible here. One is to consider a drive or emotion as a behaviour pattern in its own right, driven by a particular sensor input. Its output is an actuator response which is synthesised into that of the other **bp**s, thus having an effect upon the global emergent behaviour. Thus if an unfamiliar smell is sensed on the prevailing wind, a grazing sheep might slowly graze in the opposite direction to it so as to move away from it, with its fear behaviour . Or alternatively it might graze closer to other sheep, producing a more compact flock.

Alternatively, one could view it as an extra input into the stimulus-utility function component of a behaviour pattern. By affecting the utility component of more than one pattern, an emotion/drive in effect couples utilities together, reintroducing the idea of inhibition/excitation between patterns but in a flexible functionally defined manner. This would not only allow new behaviour to be added, but would allow existing behaviour to be suppressed via a flattening of its utility function. It would also allow for a emotional behaviour to gradually become more dominant - for example a grazing sheep becoming more and more jumpy in the face of a threatening stimulus. These approaches are not mutually exclusive, since a free-standing emotional behaviour could be different from other behaviour patterns precisely through its ability to cross-couple with utilities. Analysis of real sheep behaviour is needed to constrain the design of this emotional system.

A sheep subject to an alarming sensory stimulus may become 'restless' - fidgeting while it grazes,

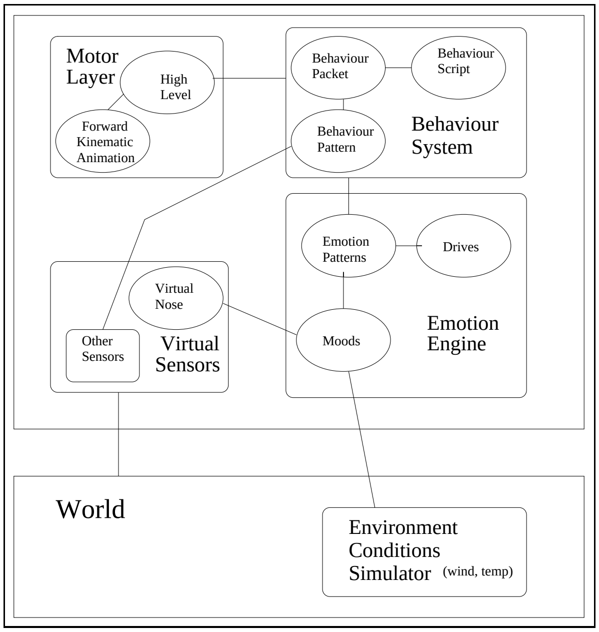

Figure 2: Architecture Block diagram

twitching its ears. As its level of fear rises, it may break off grazing to look around, before resuming; at a higher level still it may stop grazing and enter a state of alert, before finally panicking and running. With several drives/emotions acting concurrently, possibly with some modelled as a temporal cycle, it ought to be possible to model the characteristics put forward by Neary [12]. He notes that sheep will move more sluggishly on hot days and during the middle part of the day. This is simulated in our architecture by affecting the moods, as illustrated in figure 2. The mood can act as a filter for the activation of emotions patterns. Thus if the sheep is in a bad mood it could get angier fast, or if the sheep is sluggish it could react (the sensing sensitivity could be affected) slowly and move awkwardly.

The sheep are creatures of habit and they usually graze, drink and are more active in the morning and during the evening. Thus, if you are running (i.e. herding them) in the middle part of the day, the sheep may flock tighter, move slower, and spend more time looking at and fighting a dog. This suggests that careful analysis is needed in order to establish the most appropriate set of drives/emotions - Neary also points out that for sheep to be sheep, there needs to be a minimum number of five in the group. Less than five and their behaviour is not as predictable, suggesting a drive/emotion one might describe as 'group security' acting as a behavioural modifier.

3.1 Communicating emotion

Because agents exist in a VE and not in the real world, in principle the transmission of emotion between agents could just be carried out by 'cheating', that is by allowing agents to read each other's internal state directly. We choose not to do this however, since we see advantages in reusability and in matching real-world behaviour (that is in real sheep for example) by trying to model emotional interaction in a slightly more principled way. In the real-world however, emotional transmission may well be multi-modal, with certain modes such as the perception of motion (or in the general case 'body language') being particularly difficult to model. Thus we have limited ourselves for now to a single mode, and the one we have chosen is scent, to be perceived by a virtual nose sensor.

In our emotional architecture we use a virtual nose sensor, because the nose has been linked with emotional responses and intelligence. Goleman [10] states "The most ancient root of our emotional life is the sense of smell, or, more precisely, in the olfactory lobe, the cells that take in and analyse smell. Every living entity, be it nutritious, poisonous, sexual partner, predator or prey, has a distinctive molecular signature that can be carried in the wind." Neary [12]points out that sheep, particularly range sheep, will usually move more readily into the wind than with the wind, allowing them to utilise their sense of smell.

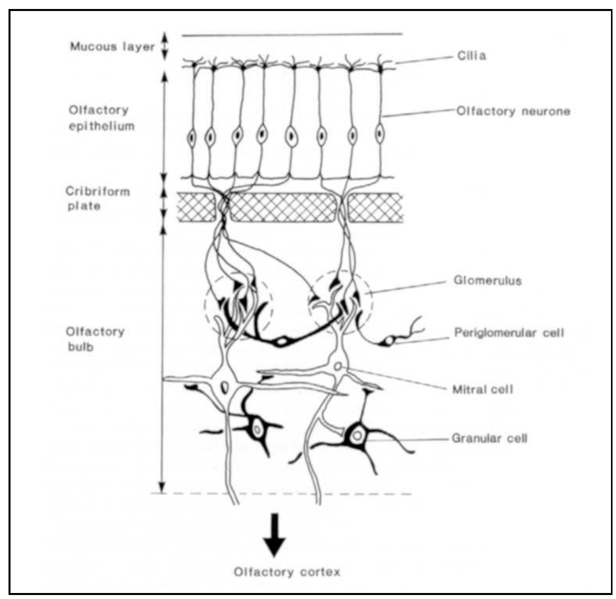

Figure 3: Mammalian olfactory system (from Gardner 1999)

There is ongoing research in the field of electronic noses, Gardner [9] but we are not aware of the use of virtual noses in a virtual environment, so important issues arise like modelling molecules and wind flow, and how the wind moves them in the environment. The structure of a particular molecule is important in determining odour. This requires sensors which are responsive to the shapes or structural features of the organic molecules being modelled.

In real animals chemoreceptors (exteroceptors and interocetors) are used to identify chemical substances and detect their concentration. Smell exists even among very primitive forms of life . In our architecture we intend to model the exteroceptors which detect the presence of chemicals in the external environment.

In vertebrates the olfactory receptors are primary sensor neurons with dendrites that extend as cilia into a mucous layer, as illustrated in Figure 3. These axons from the olfactory nuerons synapse with the dentrites of the mitral and tufted cells in the olfactory bulb. At the level of the glomeruli the quality of the odorous stimulus is encoded in the form of activated glomeruli, Axel [15]suggests that"The brain is essentially saying something like, I'm seeing activity in positions 1, 15, and 54 of the olfactory bulb, which correspond to odorant receptors (glomeruli)1, 15, and 54, so that must be jasmine, ". Most odours consist of mixtures of odorant molecules. Therefore, other odours would be identified by different combinations.

In the electronic noses industry an empirical approach is used, making use of available sensor types and attempting to modify sensor designs to meet the requirement of the electronic nose [9].

To sense smell in the virtual environment, the molecules must be distributed in the environment; this is modelled by setting the density of each of the molecules within the environment represented as a grid. To simplify the computation the current grid is 2D, but we plan to use a voxel-based grid in the future.

The smell sensor and the virtual molecules used to represent the signature of different smells, is an important feature of our architecture, because it can be used to communicate emotions between agents through the environment. For example if a sheep panics it could exude a distinctive odour, or molecular signature, to the environment using a density function, and through time the molecules would disperse depending on several artificial environment factors like wind, rain, season and time of day. Other sheep will sense the panic smell and they will panic themselves and also exude the distinctive odour for panic so that the odour will propagate through the flock in few simulation steps.

Primary odour	Familiar substance
Camphoraceous	Moth repellant
Musky	Angelica root oil
Floral	Roses
Pepperminty	Mint candy
Ethereal	Dry-cleaning fluid
Pungent	Vinegar
Putrid	Bad egg

Table 1: Primary odours with more familiar examples (Adapted from McFarland)

To simulate volatile particles, i.e. scent which can be detected with the virtual ofactory system, we use particle animation theory, in which each particle represents a concentration of molecules with a different primary odour as in Table 1 or other chemical signal, like panicking sheep. The particles behaviour depend on the emitter propeties:

1. Emitting directions
2. Initial position.
3. Initial velocity.
4. Initial size.
5. Initial transparency.
6. Shape and colour
7. Lifetime

The number or particles at a particular time can be derived from the following formula.

$$N(t) = M(t) + I(t)$$

Where M(t) is the mean number of particles perturbed by the intensity of the scent I(t) in the case of sheep panic.

The position of the particles will be influenced by weather condition, in particular wind and time of day.

$$P_i(t) = P_i(t-1) + W(t)$$

Where $P_i(t)$ is the Position of the particle at time t and W(t) is the weather vector at time t.

Molecule properties are encoded in shape and color of each particle, we are going to use John Amoore's classification, illustrated in Table 1, which is one of the most widely accepted[11].

Recent scientific studies [15] show that many mammals use a separate set of sensory receptor cells in their nose to receive social and sexual information from members of their own species, and there is growing suspicion that we (humans) do, too. We will encode this virtual "scents" as a different colour to trigger emotions, such a panic in sheep.

The reaction time of the receptor is given by:

$$t_{reaction} = t_{min} + \frac{1}{kS^n}$$

4 The cognitive dimension

In sheep it seems that the reception of emotion is rather straightforward - a panicking sheep causes others around it to panic. In humans though, it seems that emotion reception is often mediated by a base emotional stance towards an individual as well as by the general sensitivity of reception that forms one of the components of empathy. Again, for sheep, emotion can be modelled as a purely behavioural phenomenum, while it clear that in a human it also enters into cognitive functions.

It is not intended to carry out a comprehensive discussion in this paper of how a cognitive capability in the form of an AI planner can be added into the behavioural architecture discussed above in such a way as to take input from the emotional system. However the extra complexity required in the system can be illustrated from the following scenario considered in the Holodeck context.

In this scenario, one agent, representing a daughter, enters crying, into a room holding an agent representing her mother. The mother's reaction is to hug the daughter, and say "What's wrong?". The daughter says "I've split up with Dave', where Dave is her boyfriend. After a certain amount of hugging, the mother says "Sit down and let me make you a cuppa" (this refers to a cup of tea, a culturally-conditioned comforting reflex in the UK).

While it is true that the daughter's grief is in real-life detected by sound and vision, there seems to be no reason why it could not be modelled as a smell just like panic in sheep. However the outcome is not that the mother receives the emotion of grief and bursts into tears too. Her base emotional state towards the daughter might be described as 'protective', and the interaction of perceived grief with this base state produces a 'comforting' behaviour. There seems to be no reason why this cannot be modelled at the behavioural level. The idea of a modified emotional input triggering a 'comfort' behaviour seems at least as plausible as a cognitive account which would require a symbolic-level categorisation of the incoming emotion ('if a person cries then they are unhappy') followed by some reasoning about an appropriate response.

5 Virtual Sheep

To simulate the "emotional" virtual sheep in a virtual environment we are using the architecture described in Delgado [8], with the difference we have substituted Performer for Maverik, because the former offers better performance and reliability for applications developed for the CAVE see figure 4.

Figure 4: Users in the CAVE

6 Conclusion

In this brief paper we have tried to explain the behavioural architecture being used for virtual agents and to look at how drives/emotions can be integrated into it in a more sophisticated manner than the existing use of a meter representation. We believe that by modelling transmitted emotion as a smell and using a virtual nose for reception, we are employing a biologically plausible mechanism. Work will continue both to produce hopefully accurate collective sheep behaviour as well as to investigate the more varied interaction between emotion received and induced in more human-like agents.

7 Acknowledgments

The authors thank Carlos Delgado Flores, AgroIndustrias CEJA, Universidad Bonaterra and The University of Salford for their support.

References

[1] Aylett, R.S. *Communicating goals to behavioural agents - A hybrid approach*, Proceedings of the

Third World Congress on Expert Systems, Seoul, Feb 1996

[2] Aylett, R.S; Horrobin, A; O'Hare, J.J; Osman, A.& Polyak, M. *Virtual teletubbies: reapplying a robot architecture to virtual agents* Proceedings, 3rd International Conference on Autonomous Agents .

[3] Barnes, D.P. *A behaviour synthesis architecture for cooperant mobile robots.* Advanced Robotics and Intelligent Machines, eds. J.O.Gray and D.G.Caldwell, IEE Control Engineering Series 51.295-314, 1996

[4] Barnes, D.P., Ghanea-Hercock, R.A, Aylett, R.S. and Coddington, A.M. *Many hands make light work? An investigation into behaviourally controlled cooperant autonomous mobile robots*, in: Proceedings of the first International Conference on Autonomous Agents, 413-20, Johnson, L. eds., Marina del-Rey, Calif.: ACM Press

[5] Blumberg, B.. *Old Tricks, New Dogs: Ethology and Interactive Creatures.* Ph.D. diss. Massachusetts Institute of Technology

[6] Cañamero, D. *Modeling Motivations and Emotions as a Basis for Intelligent Behaviour*, Proceeding of the First International Conference on Autonomous Agents, 1997.

[7] Davis, R., and King, J.*An overview of Production Systems, in Machine Intelligence*, E.W. Elcock and D. Michie (Eds.), Wiley and Sons, 300-332, 1977

[8] Delgado, C. and Aylett, R. *From Robots to Creatures: Extending a Robot Behaviour Based Architecture to Create Virtual Creatures.* In ISRA 2000, Monterrey Mexico, 2000.

[9] Gardner, J. and Bartlett, P. *Electronic Noses:Principles and Practice.*Oxford University Press. United Kingdom. 1999

[10] Goleman, D. *Emotional Intelligence.* Bloosmbury.United Kingdom. 1996

[11] McFarland, D., *Animal Behaviour*, Addison Wesley Longman, 1999

[12] Neary, M.*Sheep Sense.* http://www.working-border-collie.com/article3.html

[13] Ortony, A., Clore, G, Collins, A. *The cognitive Structure of Emotions*, Cambdrigde University Press, 1988

[14] Picard, R., *Affective Computing*, The MIT Press, 1997

[15] Pines, M., *Seeing, Hearing and Smelling the World: New findings help Scientists make sense of our senses*, Report from the Howard Hughes Medical Institute, 1997

[16] Reilly, S. and Bates, J., *Building Emotional Agents*, Technical Report CMU-CS-92-143, School of Computer Science, Carnegie Mellon University, May 1992.

[17] Velásquez, J., *When Robots Weep: Emotional Memories and Decision-Making*