Don't you escape! I'll tell you my story

Jesus Ibanez[1], Carlos Delgado-Mata[2], Ruth Aylett[2], and Rocio Ruiz-Rodarte[3]

[1] Departamento de Tecnología, Universidad Pompeu Fabra,
Passeig de Circumvallació, 8, 08003 Barcelona, Spain
jesus.ibanez@tecn.upf.es
[2] Centre for Virtual Environments, The University of Salford,
Business House, University Road, Salford, M5 4WT, Manchester, UK
C.Delgado@pgr.salford.ac.uk, R.S.Aylett@salford.ac.uk
[3] Instituto Tecnológico de Estudios Superiores de Monterrey,
Campus Estado de México, 52926 México
caruiz@itesm.mx

Abstract. This paper makes two contributions to increasing the engagement of users in virtual heritage environments by adding virtual living creatures. This work is carried out on the context of models of the Mayan cities of Palenque and Calakmul. Firstly, it proposes a virtual guide system. The guide navigates a virtual world and tells stories about the locations within it, bringing to them its personality and role, so that differently characterised guides will produce different stories. Secondly, it develops an architecture for adding autonomous animals to virtual heritage. It develops an affective component for such animal agents in order to increase the realism of their flocking behaviour and adds a mechanism for transmitting emotion between animals via virtual pheromones, modelled as particles in a free expansion gas.

1 Introduction

Nowadays, virtual environments are becoming a widely-used technology as the price of the hardware necessary to run them decreases. Current video games show 3D environments unimaginable some years ago. Many recently developed virtual environments recreate real spaces with an impressive degree of realism. In such contexts, however, a lack of information for the user is frequently perceived, which makes him lose his interest in these environments. In the real world, people relate the environments that surround them to the stories they know about the places and objects in the environment. Therefore, in order to obtain more human and useful virtual environments, we need to add a narrative layer to them. We need stories related to the places and objects in the world. And finally, we need a virtual guide able to tell us these stories.

Furthermore, as pointed out in [30], one of the most striking features of historical investigations is the coexistence of multiple interpretations of the same event or process. The same historical events can be told as different stories depending on the storyteller's point of view. The story of the same battle between two cities, for example, will be different depending on the origin of the storyteller. It would be interesting that the virtual guide which tells us stories about the virtual environment she [4] inhabits could tell us these stories from her own perspective. Such a guide would be, in addition, very useful for educational purposes. Children would be more open-minded if they could listen to different versions of the same historical events depending on the profile of the storyteller. In this sense, the first part of this paper describes the design and development of a novel proposal for storytelling in virtual environments from a virtual guide perspective.

On the other hand, in order to obtain more believable virtual environments, some research groups are trying to model and simulate virtual animals. This objective is usually tackled from a multidisciplinary point of view. Theories and techniques from ethology, virtual reality, graphic simulation, artificial intelligence, artificial life, robotics, and software agents (among others) are used and sometimes combined. Although modelling the behaviour of autonomous animals is really a useful task, the fact is that in nature usually animals behave as part of social groups. Therefore, if we want to populate virtual environments with believable virtual animals, we need to emulate the behaviour of animal groups.

Furthermore, our assumption of life and smartness in real animals derives from our perception of their reaction to the environment, and in particular, to their reaction to perceived actions. For example, if one runs through a flock of deer, we expect them to move in order to avoid us. This reaction seems driven by emotional stimulus (most likely fear), and is communicated amongst conspecifics at an emotional level. Thus, believable virtual animals should display this kind of emotional response and communication. The second part of this paper proposes an architecture to model and simulate animals that not only "feel" emotions, that affect their decision making, but are also able to communicate them through virtual pheromones. The virtual animals also show a group behaviour, in particular a flocking behaviour based on the *boids* algorithm [26] modified so that it takes into account the animals' emotions. In this sense, the emotional communication "drives" the flocking emergence.

This paper describes the design and development of a novel proposal for storytelling in virtual environments from a virtual guide perspective and how it can be used together with animal group behaviour to add relevance to the virtual heritage instalations of Palenque and Calakmul. The structure of the paper is as follows. First we review related work. Next we describe the proposed architecture for storytelling. Then we detail the architecture for the virtual deer. Next we expose the implementation and preliminary results. Finally we show the conclusions and point out future work.

2 Related Work

This section describes related work on storytelling, emotion and group animation.

[4] In order to avoid confusion, in this paper, the virtual guide is supposed to be female, while the human guide is supposed to be male

2.1 Storytelling

Computer based storytelling research is a flourishing area. Different groups are working on different problems and approaching them from different perspectives. Some groups are trying to obtain methods and mechanisms to automatically generate narratives (this is the plot based approach), while others are more concentrated on trying to obtain believable characters [16] (character based approach). A little work tries to reconcile both approaches [17]. From an information point of view, another distinction can be pointed out. Some researchers work with story pieces which are small parts of stories, while other researchers work with simple events. People working with story pieces usually approach the problem by using narrative theories, using prescribed story structures to combine the story pieces. In this sense, the theories of Polti [23], Propp [24], Thompson [28] and Branigan [3] are especially interesting. The granularity of the information used to construct stories is a determinant factor. In general, the larger the granules used to build with, the easier it is to build. The smaller the granule, the more precise and smooth the building can be. As pointed out in [4], from an artificial intelligence point of view, two different approaches have been considered to try to solve the narrative generation problem. The knowledge based approach makes an a priori attempt to capture the rules for successfully solving or navigating a domain, the behaviour based approach instead relies on a set of lower level competences which are each "experts" at solving one small part of the larger problem domain. From a more linguistic point of view, as pointed out in [14], previous work on story generation has generally taken one of two approaches: structuralist and transformationalist. Structuralists use real-world story structures such as canned story sequences and story grammars to generate stories, while transformationalists believe that story-telling expertise can be encoded by rules, or narrative goals that are applied to story elements such as settings and characters.

In general, more predefined stories are obtained by using: plot based approach, story pieces, large granules, knowledge based approach or structuralism. More surprising and emergent stories are obtained by using: character based approach, events, small granules, behaviour based approach or transformationalism.

2.2 Emotions

Until recently the field of Artificial Intelligence (AI) had largely ignored the use of emotion and intuition to guide reasoning and decision-making. Minsky [18] was one of the first to emphasise the importance of emotion for Artificial Intelligence. Other models of emotion have been proposed; Picard focuses on recognising emotions [22]; Velásquez [32] synthesised emotions and some of their influences on behaviour and learning, using a similar approach to the one proposed in this work, that is a model based on Izard's Four Types of Emotion Elicitors [11].

Based on the description given above we decided to define emotions as reflective autonomic responses, that is primary emotions [5], triggered by particular stimuli. In [7] an extended review is provided.

2.3 Group animation

As pointed out in [21], when animating groups of objects, three general classes of techniques are usually considered, depending on the number of elements being controlled and the sophistication of the control strategy. These general types of techniques are particle systems, flocking, and (autonomous) behavioural animation.

A *particle system* is a large collection of objects (particles) which, taken together, represent a fuzzy object [25]. Particle systems are based on physics, and particles themselves are supposed not to have intelligence, therefore this kind of technique is usually employed to animate groups of objects that represent non-living objects. *Flocking systems* typically have a medium number of objects, each one of them controlled by a relatively simple set of rules that operate locally. The objects exhibit some limited intelligence and are governed by some relatively simple physics. This type of technique is usually employed to animate objects that represent living objects with simple behaviour. Particle systems and flocking systems produce the so-called *emergent behaviour*, that is, a global effect generated by local rules. For example, one can speak of a flock of birds having a certain motion, even though there is not any control specific to the flock as an entity. *Autonomous behaviour* is usually employed to animate one or a few objects which are supposed to be "intelligent". Therefore this kind of technique is usually employed to animate groups of objects that represent living objects. Generally each object is controlled by a sophisticated set of rules.

As we are trying to simulate the behaviour of medium-size animal flocks, we have decided to base the group animation algorithm on flocking techniques. However, we also want every animal to show a sophisticated behaviour based, among other things, on his emotional state. Therefore, our system is hybrid, in that it shows a global flocking behaviour but at the same time each object is controlled by a complicated autonomous behaviour.

Boids The collective behaviour of flocks, herds and schools has been studied in artificial life [13] and complex systems [8] areas. This kind of behaviour seems complex, however as Reynolds suggested [26] it can be modelled by applying few simple rules to every individual. In his model of so-called boids, every individual (boid) tries to fulfil three conditions: cohesion or flock centring (attempt to stay close to nearby flockmates), alignment or velocity matching (attempt to match velocity with nearby flockmates), and separation or collision avoidance (avoid collisions with nearby flockmates). Boids has been successfully used to animate the behaviour of flocks in several famous films.

3 Narrative Construction

In our model the guide begins at a particular location and starts to navigate the world telling the user stories related to the places she visits. Our guide tries to emulate a real guide's behaviour in such a situation. In particular, she behaves as a spontaneous real guide who knows stories about the places in the virtual world but has not prepared an exhaustive tour nor a storyline.

Furthermore, our guide tells stories from her own perspective, that is, she narrates historical facts and events taking into account her own interests and roles. In fact, she extends the stories she tells with comments that show her own point of view. This mixture of neutral information and personal comments is what we can expect from a real guide who, on the one hand, has to tell the information he has learnt, but on the other hand, cannot hide his feelings, opinions, etc about the information he is telling. We have designed a hybrid algorithm that models a virtual guide behaviour taking into account all the aspects described above. The mechanisms involved in the algorithm can be separated in three global processes which are carried out with every step. The next three subsections detail these general phases.

3.1 Finding a Spot in the Guide's Memory

Given a particular step in the navigation-storytelling process (that is, the virtual guide is at a particular location and she has previously narrated a series of story pieces), the guide should decide where to go and what to tell there. To emulate a real guide's behaviour, the virtual guide evaluates every candidate pair (story element, location) taking into account three different factors: the distance from the current location to location, the already told story elements at the current moment and the affinity between story element and the guide's profile.

A real guide will usually prefer nearer locations, as further away locations involve long displacements which lead to unnatural and boring delays among the narrated story elements. In this sense, our guide prefers nearer locations too, and therefore shorter displacements. When a real guide is telling stories in an improvisational way, the already narrated story elements make him recall, by association, related story elements. In a spontaneous way, a real guide tends to tell these recently remembered stories. In this sense, our guide prefers story elements related (metaphorically remembered) to the ones previously narrated. Finally, a real guide tends to tell stories related to his own interests (hobbies, preferences, etc) or roles (gender, job, religion, etc). In this sense, our guide prefers story elements related to her own profile.

The system evaluates every candidate pair (storyelement, location) such that there is an entry in the knowledge base that relates storyelement to location (note that this means that storyelement can be narrated in location) and such that storyelement has not been narrated yet. In particular three scores corresponding to the previously commented factors are calculated. These three scores are then combined to calculate an overall score for every candidate pair. Finally the system chooses the pair with the highest overall score value.

3.2 Extending and Contextualising the Information

Figure 1a represents a part of the general memory the guide uses. This memory contains story elements that are interconnected with one another in terms of different relations. In particular, in our case, cause-effect and subject-object relations interconnect the story elements. Figure 1b shows the same part of the memory, where a story element has been selected by obtaining the best overall score described in the previous section. If the granularity provided by the selected story element is not considered to be large enough to generate a little story, then more story elements are selected. The additional story elements are chosen according to a particular criteria or a combination of several criteria (cause-effect and subject-object in our case). This process can be considered as navigating the memory from the original story element. Figure 1c shows the same part of the memory, where three additional story elements have been selected by navigating from the original story element.

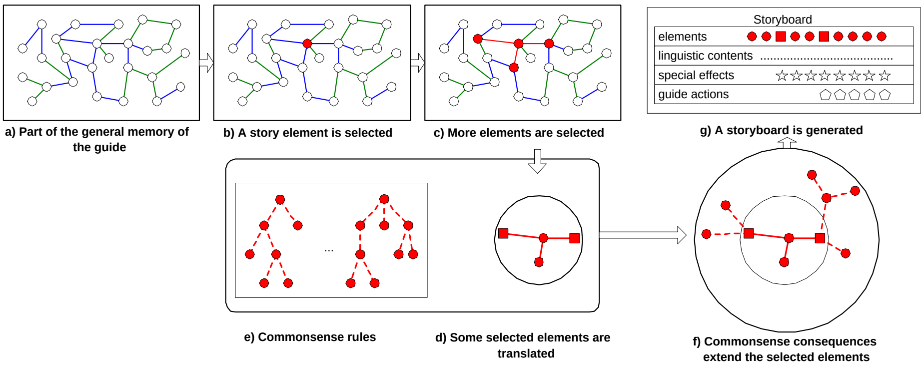

Fig. 1. Storyboard construction

Once the granularity provided by the selected story elements is considered to be large enough, the selected story elements are translated, if possible, from the virtual guide perspective (see figure 1d). For this task the system takes into account the guide profile and the meta-rules stored in the knowledge base that are intended to situate the guide perspective. The translation process also generates guide attitudes that reflect the emotional impact that these story elements cause her. Lets demonstrate this by a simple example. Let us assume the following information extracted from a selected story element

```
fact(colonization, spanish, mayan)
```

meaning that the Spanish people colonized the Mayan. And let us assume the following meta-rules included in the knowledge base, aimed to situate the guide perspective

```
fact(colonization, Colonizer, Colonized) and profile(Colonized) =>
        fact(colonizedColonizacion, Colonizer, Colonized) and
        guideattitude(anger)
```

meaning that a colonizedColonization fact and anger as the guide's attitude should be inferred if a colonization fact is included in the story element and the guide profile matches the third argument of this fact, that is, the guide is the Colonized. In this example that will happen if the guide is Mayan. The new inferred fact represents the original one but from the guide's perspective.

In addition, the new translated story elements are enhanced by means of new information items generated by inferring simple commonsense rules allowing to add some comments showing her perspective. The guide uses the new contextualised story elements (figure 1d) as input for the rules that codify commonsense (figure 1e). By inferring these rules the guide obtains consequences that are added to the contextualised story elements (figure 1f), obtaining a new data structure which codifies the information that should be told. Let us continue with the previous example. Let us assume the following commonsense rule

```
fact(colonizedColonizacion, Colonizer, Colonized) =>
        fact(culturalDestruction, Colonized) and
        fact(religionChange, Colonized)
```

meaning that the colonized's view implies the destruction of the colonized's culture and the change of the colonized's religion. Therefore, if in our example the guide were Mayan, the story element to be told would be enhanced with the facts culturalDestruction and religionChange.

3.3 Generating the Story

As a result of the previous processes, the guide obtains a set of inter-related information items to tell (figure 1f). These elements are stored as a structure that reflects the relations among them, as well as the reasons why each one was selected. Some elements are also related to particular guide attitudes. Now the system generates the text to tell (expressing these elements) as well as special effects and guide's actions to show while telling the story. The phases of this story generation process are as follows:

1. The first step is to order the data elements. To do so we consider three criteria: *cause-effect* (if an element Y was caused by another element X, then X should precede Y), *subject-object* (the elements whose subject/object are similar should be grouped together) and *classic climax* (the first selected story element, i.e. the one that obtained the best overall score, is supposed to be the climax of the narration, and therefore all the rest of the elements are arranged taking it into account).
2. The text corresponding to the ordered set of elements is generated. The complexity of this process depends on the particular generation mechanism (we use a template system) and the degree of granularity employed (we use a sentence per every story element).
3. A process that relies on the guide expression rules (the set of rules that translate abstract guide's attitudes in particular guide's actions) generates a set of guide actions (each one related to a particular story element).
4. Every story element is associated to particular environment conditions or special effects. Thus, finally, a storyboard like the one shown in figure 1g is obtained.

4 Communicating Emotions and Group Behaviour

4.1 Overall architecture

The basic task of an animal brain has often been split into three sub-tasks. Our model adds a fourth sub-task, emotions. The four sub-tasks in our system are therefore: perception (sensing of the environment and interpretation of the sensory signals to provide a high-level description of the environment), emotions (which affect the behaviour of the animals, exemplified by the conspecifics flight-flocking), action selection (using the perceptual and emotional inputs to decide which of the animal's repertoire of actions is most suitable at that moment) and motor control (transforming the chosen action into a pattern of "physical" actions to produce the animation of the animal). Figure 2 shows a detailed diagram of the designed architecture, and the next sections describe its components.

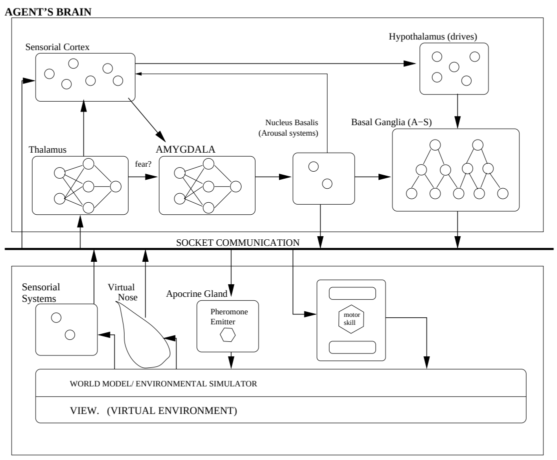

Fig. 2. Detailed architecture

4.2 Communicating emotions

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 re-usability and in matching real-world behaviour (that is in real deer 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 being particularly difficult to model. Thus we have limited ourselves for now to a single mode, and the one we have chosen is pheromones, to be perceived by a virtual olfaction sensor.

The nose has been linked with emotional responses and intelligence. Recent experiments [10] have shown that mammals, including humans, emit pheromones through apocrine glands as an emotional response, and as means to communicate that state to conspecifics, who can adapt their behaviour accordingly; research has found that odours produce a range of emotion responses in animals, including humans [11], which is adaptively advantageous because olfaction is part of the old smell-brain which can generate fast emotional-responses, that is without the need of cognitive processes.

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 [20]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.

In real animals chemoreceptors (exteroceptors and interoceptors) 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 this work, to illustrate the use of emotion and drives to influence behaviours, deer had been selected as the exemplar creature. To support the communication of emotions, an environmental simulator has been developed, its tasks include changing the temperature and other environmental variables depending on the time of day and on the season, which depends on statistical historical data. An alarmed animal sends virtual pheromones to the environmental simulator and they are simulated using the free expansion gas formula in which the volume depends on the temperature and altitude (both simulated environmental variables). To compute the distribution of the pheromones a set of particles has been simulated using Boltzmann distribution formula 1, which is shown and described next.

$$n\left(y\right) = n_o e^{\frac{mgy}{k_b}T} \tag{1}$$

Where m is the pheromone's mass; g is the gravity; y is the altitude; k_b is the Boltzmann number; T is the temperature; n_o is N/V; N is number of molecules exhuded by the apocrine gland, which is related to the intensity of the emotional signal; and V is the Volume.

The virtual animal includes a virtual nose used to detect pheromones, if any, that are near the creature. To smell a pheromone the threshold set in the current experiment is $200x10^{-16}$ because [15] has shown that animals have 1000 to 10000 more sensitivity than humans and Wyatt [33] claims that the threshold in humans to detect certain "emotional response odours", like those exuded from the armpits, is 200 per trillion parts, that is $200x10^{-12}$.

4.3 Action selection

The problem of action selection is that of choosing at each moment in time the most appropriate action out of a repertoire of possible actions. The process of making this decision takes into account many stimuli, including (in our case) the animal's emotional state.

Action selection algorithms have been proposed by both ethologists and computer scientists. The models suggested by the ethologists are usually at a conceptual level, while the ones proposed by computer scientists (with some exceptions as [31] and [2]) generally do not take into account classical ethologic theories.

According to Dawkins [6], a hierarchical structure represents an essential organising principle of complex behaviours. This view is shared by many ethologists [1] [29], and some action selection models follow this approach.

Our action selection mechanism is based on Tyrrell's model [31]. This model is a development of Rosenblatt & Payton's original idea [27] (basically a connectionist, hierarchical, feed-forward network), to which temporal and uncertainty penalties were added, and for which a more specific rule for combination of preferences was produced. Note that among other stimuli, our action selection mechanism takes the emotional states (outputs of the emotional devices) of the virtual animal.

4.4 The flocking behaviour

The flocking behaviour in our system is based on *boids* [26], although we have extended it with an additional rule (escape), and, most importantly, the flocking behaviour itself is parameterised by the emotional devices output, that is, by the values of the emotions the boids feel. The escape rule is used to influence the behaviour of each boid in such a way that it escapes from potential danger (essentially predators) in its vicinity. Therefore, in our model each virtual animal moves itself along a vector, which is the resultant of four component vectors, one for each of the behavioural rules, which are:

Cohesion - attempt to stay close to nearby flockmates.
Alignment - attempt to match velocity with nearby flockmates.
Separation - avoid collisions with nearby flockmates.
Escape - escape from potential danger (predators for example).

The calculation of the resultant vector, $Velocity$, for a virtual animal A is as follows:

$$V_A = \underbrace{(Cf \cdot Cef \cdot Cv)}_{Cohesion} + \underbrace{(Af \cdot Aef \cdot Av)}_{Alignment} + \underbrace{(Sf \cdot Sef \cdot Sv)}_{Separation} + \underbrace{(Ef \cdot Eef \cdot Ev)}_{Escape} \tag{2}$$

$$Velocity_A = limit(V_A, (MVef \cdot MaxVelocity)) \tag{3}$$

where Cv, Av, Sv and Ev are the component vectors corresponding to the cohesion, alignment, separation and escape rules respectively. Cf, Af, Sf and Ef are factors representing the importance of the component vectors Cv, Av,

Sv and Ev respectively. These factors allow to weight each component vector independently. In our current implementation they can be varied, in real time, from a user interface. Cef, Aef, Sef and Eef are factors representing the importance of the component vectors Cv, Av, Sv and Ev respectively, given the current emotional state of the virtual animal. That is, each of this factors is a function that take the current values of the animals emotions and generate a weight for its related component vector. $MaxVelocity$ is the maximum velocity allowed to the animal. In the current implementation it can be varied from a user interface. $MVef$ is a factor whose value is calculated as a function of the current values of the animals emotions. It allows to increase and decrease the animal's $MaxVelocity$ depending on its emotional state. *limit* is a function whose value is equal to its first parameter if this is not greater than its second one, otherwise the function value is equal to its second parameter.

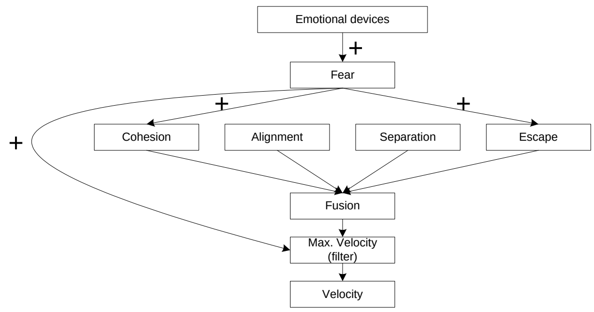

Fig. 3. How *fear* parameterises the flocking algorithm

The emotional factors (Cef, Aef, Sef, Eef, and $MVef$) reflects ethologic heuristic rules. Figure 3 shows an example of how emotions parameterise the flocking behaviour. In particular it shows how *fear* affects the component vectors of the animals' behaviour. The greater the fear an animal feels, the greater the weight of both its cohesion vector (the animal try to stay closer to nearby flockmates) and its escape vector (the boid try to stay farther from the potential danger). The resultant vector obtained by adding the four basic vectors is then scaled to not exceed the maximum speed. This maximum velocity is parameterised by the fear as well. The greater the fear an animal feels, the greater the speed it is able to reach.

5 Implementation and preliminary results

We have chosen Unreal Tournament (UT) engine as the platform on which our virtual guide run. As we wished our system to be open and portable, we decided to use Gamebots to connect our virtual guide to UT. Gamebots [12] is a modification to UT that allows characters in the environment to be controlled via network sockets connected to other programs. The core of the virtual guide is a Java application which is able to connect to UT worlds through Gamebots. This Java application controls the movement and animations of the guide in the world as well as the presentation of special effects and texts which show the generated narratives. The current version uses a MySQL [19] database to store the knowledge base. The Java application accesses these data through JDBC. The developed system uses Jess [9] to carry out inferences on the information.

The described system has been developed and it is working properly with small and medium size knowledge bases. Figure 4 shows both the user and administrator interfaces of the current implementation. We still have to check how the system behaves when dealing with large knowledge bases and evaluate different template systems for the generation of text from the storyboard.

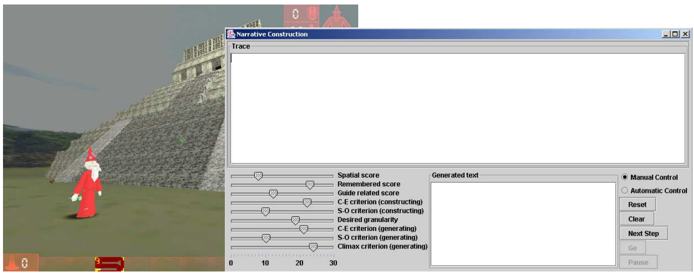

Fig. 4. User and administrator interfaces

On the other hand, the implementation of the emotional virtual animals architecture is three layered. Namely the agent's brain, the world model and the virtual environment. As seen in figure 2 the agents' brains are processes that runs independently on a Linux workstation and each agent's brain receives the sensorial data via network sockets and sends (through the network) the selected action to the world's model which contains the agents' bodies and the environmental simulation. The changes to this model are reflected on each frame in the Palenque virtual environment (see figure 5b) which was developed using OpenGL Performer. This mechanism allows modularity and extensibility to add/modify the behaviour of the virtual animals. Furthermore, the behaviour of the deer in the Calakmul environment (see figure 5a) was prototyped in VRML and Java.

The tests carried out so far, although preliminary results, show that the proposed architecture is able to cope with the problem of simulating flocks of virtual animals driven by its emotional state. For example, the emergent behaviour observed when the virtual deer are attacked by predators seems more real than the one obtained by using classical flocking algorithms. While in models without emotional communication animals only react to a predator when it is in their field of view, in our system animals without a predator in its field of view start to react after its neighbours do so. This behaviour, which is due to the pheromones communication/perception [33], emulates nature more believably.

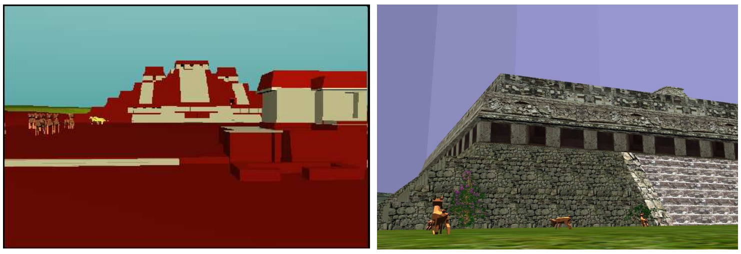

Fig. 5. Snapshot of the Calakmul (a) and Palenque (b) implemented systems

6 Conclusions and Future Work

In this paper we have described work which aims to put 'life' into virtual heritage environments in two distinct ways. In this way the feeling of 'lovely visuals, but so what ?' which can affect users will be replaced by a greater engagement with the heritage both conceptually and in terms of a feeling of presence. Firstly, we discussed work towards the creation of an "intelligent guide with attitude", who tells stories in a virtual heritage environment from her distinct point of view. This application requires story-telling rather than any other type of narrative, and must link the memory and interests of the guide to her spatial location so that her stories are relevant to what can be immediately seen. We do not at this stage incorporate user interaction, as this poses a number of new problems and in any case has been more widely studied. Work will continue in conjunction with groups in Mexico who have already produced a virtual models of the Mayan Cities of Palenque and Calakmul (whose location in the middle of a jungle makes it particularly inaccessible in the real world).

Our generic architecture means that such a guide can be ported to other virtual heritage models, though one should not underestimate the amount of content that has to be created for a particular site. Further work will be carried out in order to support the authoring process required. We believe that the growing popularity of virtual heritage produces a growing need for intelligent guides and that this work will therefore find many potential applications.

Secondly, this work has shown that it is feasible to bring together the simple rule-based architectures of flocking with the more complex architectures of autonomous agents. Initial results suggest that populating a virtual heritage site with virtual animals can improve the experience particularly if the animals are engaged in some autonomous activity. This produces more believable and more specific flocking behaviour in the presence of predators. Further work will be carried out to more accurately characterise the changes in flocking behaviour obtained by this extended architecture. Transitions between grazing and reacting to predators will be carefully analysed. We further plan to validate this work by modelling a different flocking animal, for example the musk ox, which responds to predators by forming a horns-out circle of adult animals with the young inside.

We argue that our extended architecture can support a much wider and more realistic spread of flocking behaviour and therefore represents a contribution to believable virtual animals.

References

1. Gerald P Baerends, *The functional organization of behaviour*, Animal Behaviour **24** (1976), 726–738.
2. Bruce Blumberg, *Action-selection in hamsterdam: Lessons from ethology*, Proceedings of the Third International Conference on Simulation of Adaptive Behavior (Brighton, England), 1994.
3. Edward Branigan, *Narrative comprehension and film*, New York: Routledge, 1992.
4. Kevin Michael Brooks, *Metalinear cinematic narrative: Theory, process, and tool*, Ph.D. thesis.
5. Antonio Damasio, *Descartes' error:emotion, reason and the human brain*, Papermac, Oxford, United Kingdom, 1996.
6. Richard Dawkins, *Hierarchical organisation: A candidate principle for ethology*, Growing Points in Ethology (Bateson & Hinde, ed.), Cambridge University Press, 1976.
7. Carlos Delgado-Mata, Rocio Ruiz-Rodarte, Ruth S Aylett, M A Perez-Cisneros, and P A Cook, *Illusion of life in palenque: Communicating emotions inside a virtual environment for maya cultural heritage*, PRICAI: Seventh Pacific Rim International Conference on Artificial Intelligence (Tokyo, Japan) (Helmut Prendinger, ed.), PRICAI, 2002.
8. G. W. Flake, *The computational beauty of nature : Computer explorations of fractals, chaos, complex systems, and adaptation*, MIT Press, 1999.
9. Ernest Friedman-Hill, *Jess the rule engine for the java platform*, Available at http://herzberg.ca.sandia.gov/jess/.
10. Karl Grammer, *5-alpha-androst-16en-3alpha-one: a male pheromone ?*, Ethology and Sociobiology **14** (1993), no. 3, 201–207.
11. Carol E Izard, *Four systems for emotion activation: Cognitive and noncognitive processes*, Psychological Review **100** (1993), no. 1, 68–90.
12. G. A. Kaminka, M. M. Veloso, S. Schaffer, C. Sollitto, R. Adobbati, A. N. Marshall, A. Scholer, and S. Tejada, *Gamebots: a flexible test bed for multiagent team research*, Communications of the ACM **45** (2002), no. 1.
13. C. Langton, *Artificial life: An overview*, MIT Press, Cambridge, MA, 1995.
14. H. Liu and P. Singh, *Makebelieve: Using commonsense to generate stories*, Proceedings of the 20th National Conference on Artificial Intelligence, (AAAI-02), Student Abstracts, (Edmonton, Canada), 2002, pp. 957–958.
15. D A Marshall, L Blumer, and D G Moulton, *Odor detection curves for n-pentanoic acid in dogs and humans*, Chemical Senses **6** (1981), 445–453.
16. M. Mateas, *An oz-centric review of interactive drama and believable agents*, AI Today: Recent Trends and Developments (M. Wooldridge and M. Veloso, eds.), Lecture Notes in AI, vol. 1600, Springer, Berlin, NY, 1999.
17. M. Mateas and A. Stern, *Towards integrating plot and character for interactive drama*, Working notes of the Social Intelligent Agents: The Human in the Loop Symposium. AAAI Fall Symposium Series. Menlo Park, CA: AAAI Press, 2000.
18. Marvin Minsky, *The society of mind*, Simon & Schuster, New York, USA, 1985.
19. MySQL-AB, *Mysql*, Available at http://www.mysql.com/.

20. Mike Neary, *Sheep sense*, The Working Border Collie (2001).
21. Rick Parent, *Computer animation: Algorithms and techniques*, Morgan Kaufmann, San Fransisco, CA, 2002.
22. Rosalind Picard, *Affective computing*, MIT Press, Cambridge, MA, USA, 1997.
23. Georges Polti, *The 36 dramatic situations*, Boston, The Writer, Inc., 1922.
24. Vladimir Propp, *Morphology of the folktale*, University of Texas Press, 1968, First edition 1928.
25. W. Reeves, *Particle systems - a technique for modeling a class of fuzzy objects*, ACM Transactions on Graphics **2** (1983), no. 2, 91–108.
26. Craig W Reynolds, *Flocks, herds, and schools: A distributed behavioral model*, Computer Graphics **21** (1987), no. 4, 25–34.
27. J.K. Rosenblatt and D.W. Payton, *A fine-grained alternative to the subsumption architecture for mobile robot control*, Proceedings of the IEEE/INNS International Joint Conference on Neural Networks (Washington DC), vol. 2, June 1989, pp. 317–324.
28. Stith Thompson, *Motif index of folk literature*, Indiana University Press, Bloomington, 1955-58.
29. Nikolaas Tinbergen, *The study of instinct*, Oxford University Press, United Kingdom, 1969.
30. Veronica Tozzi, *Past reality and multiple interpretations in historical investigation*, Studies in Social and Political Thought **2** (2000).
31. Toby Tyrrell, *Computational mechanisms for action selection*, Ph.D. thesis, University of Edinburgh, Edinburgh, Scotland, 1993.
32. Juan Velásquez, *Modeling emotions and other motivations in synthetic agents*, Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97) (Providence, RI), MIT/AAAI Press, 1997.
33. Tristam D Wyatt, *Pheromones and animal behaviour*, Cambridge University Press, Cambrigde, U.K., 2003.