

J. Ibanez · R. Aylett · R. Ruiz-Rodarte

Storytelling in virtual environments from a virtual guide perspective

Received: 31 January 2003 / Accepted: 21 May 2003
© Springer-Verlag London Limited 2003

Abstract This paper describes our proposal for storytelling in virtual environments from a virtual guide perspective, detailing the involved knowledge representation and algorithms. 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.

Keywords Digital storytelling · Narrative construction · Virtual environments · Virtual heritage

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 associate the environments that surround them with the stories they know about the places and objects in the environment. When we visit a new city, we learn about it through the stories our friends (or a tour guide) tell us, or through the stories we read about the places we visit, or through our own experience in these places (which we store as stories in our memory). Virtual environments lack a narrative layer. In order to obtain more human and useful environments, we need stories related to the places and objects in the world. Thus, we recognise the necessity of a virtual guide able to tell us these stories.

On the other hand, as pointed out in Tozzi [1], one of the most striking features of historical investigations is the coexistence of multiple interpretations of the same event. The same historical events can be narrated 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 if the virtual guide, who tells us stories about the virtual environment she¹ inhabits, could tell us these stories from her own perspective. In addition, such a guide would be 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, this paper describes the design and development of a novel proposal for storytelling in virtual environments from a virtual guide perspective. The structure of the paper is as follows. First, we survey related work. Then we expose our proposal detailing the designed architecture (including the

J. Ibanez (✉)
Departamento de Tecnología,
Universidad Pompeu Fabra,
Passeig de Circumval·lació, 8,
08003 Barcelona, Spain
E-mail: jesus.ibanez@upf.edu

R. Aylett
Centre for Virtual Environments,
University of Salford,
Business House, Salford,
Manchester M5 4WT, UK
E-mail: r.s.aylett@salford.ac.uk

R. Ruiz-Rodarte
Instituto Tecnológico de Estudios Superiores de Monterrey,
Campus Estado de Mexico,
52926 Mexico
E-mail: caruiz@itesm.mx

¹ 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.

processes and mechanisms involved) as well as its current implementation and the scenarios where we are using the system. Finally, we provide the conclusions.

Related work

In this section, we first describe literature on computer based storytelling pointing out general trends in this area. Next we expose works that especially inspire our proposal.

Computer-based storytelling

Computer-based storytelling research is a flourishing area. Different groups are working on different problems and approaching them from diverse 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 [2] (character-based approach). Little work tries to reconcile both approaches [3].

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 through one of these two ways: connecting these pieces in a hypertext-like network (and then navigating through this network to get a story) or using narrative theories, employing prescribed story structures to combine the story pieces. These works rely on a long tradition that believes oral and written narrative to have well-defined high-level structures. In this sense, the theories of Polti [4], Propp [5], Thompson [6] and Branigan [7] are especially interesting. The application of Propp's morphological approach to interactive storytelling was suggested in Murray [8]. For a recent example of the application of Propp's theory see Braun [9] and Spierling et al. [10]. For an example of the application of Branigan ideas, see Brooks [11].

The granularity of the information used to construct stories is an essential factor. In general, the larger the granules used, the easier it is to build. The smaller the granule, the more precise and smooth the building can be. Systems with large granules are easier to maintain and their related story generation process is simpler, but the generated stories are more 'predefined'. Systems with small granules are more difficult to maintain and their related story generation process is more complex (at least if the story is expected to show narrative devices as climax, tension, etc), but the generated stories are more surprising.

As pointed out in Brooks [11], 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. In this sense, the works trying to apply Propp's ideas to generate narrative are using a knowledge based approach, while works giving every character autonomy in order to get an emergent narrative [12, 13] from their interactions are using a behaviour based approach.

From a more linguistic point of view, as pointed out in Liu and Singh [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, as we show in Table 1, 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.

Inspiring our work

Some systems are not intended to generate complex stories nor simple presentations, but little stories in some point between both extremes. In this sense, Terminal Time [15] is a system that combines historical events, ideological rhetoric, familiar forms of TV documentary, consumer polls and artificial intelligence algorithms to create hybrid cinematic experiences for mass audiences that are different every time. Through an audience response measuring device connected to a computer, viewing audiences respond to periodic questions reminiscent of marketing polls. Their answers to these questions allow the computer program to create historical narratives that attempt to mirror and often exaggerate their biases and desires. The engine uses the past 1000 years of world history as 'fuel' for creating these custom-made historical documentaries.

Our proposal, as Terminal Time, tries to generate narratives from a particular perspective or ideology. However, while Terminal Time takes into account an ideology biased by the audience, we consider the virtual guide profile. While Terminal Time selects the story elements to be told according to its ability to show a plan predefined to support its ideology, we select the

Table 1 General approaches

More predefined	plot	story pieces	large granules	knowledge based	structuralism
More emergent	character	events	small granules	behaviour based	transformationalism

story elements according to several factors (as described below).

Our proposal deals, in particular, with the problem of guiding in a virtual environment. This problem, which has its own special characteristics, has been explored in Kyoto Tour Guide [16], a project aimed to develop an agent that is integrated in an on-line 3D website tour, a digital version of Kyoto. Kyoto Tour Guide explores ways to make an agent's narrative effectively adaptive to different groups who take the tour, like an expert human tour guide. The authors focused on developing storytelling strategies that produce an engaging experience for tour takers. They derived a list of target abilities for the agent by researching the behaviour of actual tour guides in Kyoto. They found that tour guides made use of illustrative stories frequently, supplementing the rich visual environment of the city with explanations of how Japanese people, both past and present, made use of these settings. Stories included things such as: descriptions of how a given site was constructed and how it was destructed and reconstructed; descriptions of peak historic events and customary activities that occur at the site. While visitors looked at the buildings, and took things in visually, the guide would create a narrative context for the site with these stories, providing visitors with stories that they could share with their fellow tour members, as well as with people back home.

In this sense, our proposal deals with the same problem as Kyoto Tour Guide does. However, while Kyoto Tour Guide uses a database of stories related to the sites our system constructs stories, step by step, from story elements (encoding historical facts) by means of processes that take into account several factors.

A virtual guide is expected to tell short stories about the places she visits. Therefore, our system is expected to generate this kind of short stories. MAKEBELIEVE [14] is a story generation agent that can generate short fictional texts of 5–20 lines when the user supplies the first line of the story. MAKEBELIEVE uses a subset of sentences extracted from Open Mind Commonsense (OMCS) [17] that describe simple commonsense causations. By inferring over these causations the system is able to generate a text like this: “John became very lazy at work. John lost his job. John decided to get drunk. He started to commit crimes. John went to prison. He experienced bruises. John cried. He looked at himself differently”. The authors reported that, although the system does not incorporate plot devices such as motifs, climax, tension, etc., many users involved in the evaluation nonetheless felt that these devices were present in the generated stories.

As shown below, we use this approach in our system. However, while MAKEBELIEVE generates complete short stories just by inferring on commonsense rules, our system uses commonsense rules just to enhance story elements previously selected and translated to show the virtual guide perspective. Note that MAKEBELIEVE can generate the complete story by inferring on commonsense rules because the stories it generates are

fictional, while we cannot do it because we construct stories based on historical facts.

Dautenhahn [18] explores how theories on autobiographical memory can be applied to a virtual environment. She argues that these environments have the potential to provide a rich storytelling environment, so that agents and objects within the virtual environment can remember and associate stories, and serve as editing, composing, and re-construction devices for autobiographical memories of human users of the environment. Such a concept only makes sense in environments in which persons are visiting on a regular basis over a large period of time. Dautenhahn classifies narrative agents in four types: Type 0 (an agent who is always telling a single story), Type I (an agent who can tell a great variety of stories, but the stories are not situated to the conversational context, i.e. the agent randomly selects a single story from its story-base and tells it in exactly the same way as it was stored), Type II (an agent who selects the story which fits the current context best), Type III (an agent who tells and listens to stories, ‘understands’ them in the sense that it is able to interpret the meaning and content of the story and finds in its own story-base the most similar story which is then adapted in order to produce an appropriate response) and Type IV (an autobiographic agent, whose story-telling ability is linked to a living, autonomous agent, a personality).

Our system does not construct an autobiographical memory, instead, the virtual guide is supposed to have her own autobiographical memory. In fact, the guide just has memories about historical facts related to different places and objects in the virtual environment. The guide seems to have an autobiographical memory by telling profile-biased stories based on general memories. Our guide can be considered an agent belonging to a point between Type II and Type III, as she selects the story which fits the current context best (location, guide profile, recent memories), but also adapts this story to produce a new one from her own view point.

Proposed system

In this section we describe our proposal to tell stories in virtual environments from a virtual guide perspective. We start by showing the global architecture, then we expose the structure of the information the system uses. Next we describe the narrative construction, detailing the processes and algorithms involved.

Global architecture

As shown in Fig. 1, the virtual environment runs on a virtual environment server and the guide is connected to this server through the Internet via sockets.




Fig. 1 Global architecture

The knowledge the system uses includes information, meta-information and guide dependant information. The meta-information establishes the logical structure of the data stored in the knowledge base as well as logical connections among the information items. These connections are vital to construct narratives with the coherence that makes them believable. The guide dependant information defines features that specify guide's characteristics and therefore allow generating

narratives from the guide's perspective. Note that the distinction between meta-information and information is just conceptual, therefore both could be stored in the same database.

At a conceptual level, the main processes are: the *broker* (controls the operation of the other modules), the *communications manager* (deals with the communication with the virtual environment server), the *KB manager* (controls the knowledge base operations), the *narrative constructor* (constructs the narrative and performs the play) and the *user interface* (allows configuring the guide).

Representation

In this section we describe the knowledge that the system uses. As shown in Fig. 1, this knowledge includes information, meta-information and guide dependant information. Note that, although the list of elements described in this section is not complete, it is enough to understand the processes and mechanisms the proposal involves.

Navigation information

According to Saretto [19], the process of navigating throughout an environment can be considered to be composed of two tasks and two tactics. The tasks are *searching* (looking for a known target) and *browsing* (seeing what is available in the world). The tactics are *querying* (submitting a description of the object being sought to a search engine which will return relevant content or information) and *navigation* (moving oneself sequentially around an environment, deciding at each step where to go next based on the task and the parts of the environment seen so far). A combination of a task with a tactic produces a navigation process.

Our guide is supposed to know the environment she inhabits. All the information she needs to navigate the virtual world (as a metaphor of her memories) is stored in the knowledge base and accessible to her. Therefore, in our case, the task is *searching* and the tactic *querying*, that is, our virtual guide will search a space by asking a

question about the location of a place or an object. In particular, the knowledge base stores the following data structures related to the navigation process:


- **Navigation network:** it is a network defined by a set of *accessible nodes* and a set of *connectivity segments*. Every accessible node is a reachable three-dimensional point in the virtual environment and every connectivity segment connects two accessible nodes, indicating the possibility of moving from one of them to the other one.
- **Minimal paths table:** it is a table storing the minimal path (a sequence of accessible nodes), in terms of distance, between every pair of accessible nodes. The minimal paths table is calculated just once and stored in the knowledge base in order to be used by the virtual guide.

Basic ontologies

The system defines three basic ontologies: an *object type ontology* in terms of which it is possible to describe the type of every relevant object in the world, a *guide profile ontology* which allows to describe guide profiles (in particular, our guide profile ontology defines roles and interests), and a *story element attributes ontology* which describes the attributes that can be used to annotate the story elements. Next we show a part of an attribute ontology for the case of the Palenque virtual world:

- society
- alimentation
- clothes
- ...
- science
- astronomy
- mathematics
- ...

Fig. 2 A sample guide profile



Guide profile

The guide profile is composed of roles and interests. Figure 2 shows an example of a guide profile.

Object information

Every relevant object located in the virtual world (i.e. every object susceptible of being referred to during the storytelling process) is described in the knowledge base in terms of a set of characteristics:

- **Name:** every object in the world is uniquely identified by its name.
- **Location:** this allows the guide to point or get close to a particular object during the storytelling process.
- **Width, Height and Depth:** this allows the guide, for example, to point more precisely to a particular object (by knowing its size, she can point to the object centre).
- **Type:** every object is annotated with its object type according to a particular ontology.
- **Description:** it is a textual description of the object.

Basic concepts

Some *basic concepts* (objects, characters or abstract concepts) are defined in terms of a textual description. These definitions are used to introduce the related concept the first time they appear in a particular storytelling process. These concepts also serve as clues to calculate the similarity between every pair of story elements, and therefore to calculate how much a story element is triggered off by the already told story elements, in a remembrance-like process.

Story elements

These are basic pieces of historical knowledge representing historical events expressed in a special format. These story elements, which are the basic pieces the virtual guide uses to construct stories, are, in principle, ideologically agnostic, that is, they do not reflect any particular ideology. Story elements are annotated with the following features:

- **Name:** every story element in a narrative space is uniquely identified by its name.
- **Type of event:** Each story element belongs to a particular kind of event. The kind of story element allows the narrative construction processes to guess what features the story element has, and what rules to apply in order to generate text sentences from the story element.
- **Location:** each story element is related to one or more locations (or objects in the world, and therefore related to these objects' locations). These relations are weighted. The higher the weight associated to a par-

ticular pair (story element, location), the stronger the relation between the story element and the location.

- **Attributes:** each story element is tagged according to several attributes. These attributes describe the nature of the story element (for example artistic, religious, political, social). These relations are weighted. The higher the weight, the stronger the relation between the story element and the attribute.
- **Basic concepts:** basic elements are specially tagged in the story elements so that the system knows that there is a definition for this concept in the knowledge base. The first time that the narrative constructor uses a particular basic concept, it will retrieve and use its related definition (description of the concept) in order to inform the user. The basic concepts are also used to simulate guide's memories.
- **Date:** each story element is related to a particular moment or period (date, month, year, temporal range, etc).
- **Special environment conditions:** weather, light, etc.
- **Granularity:** it indicates the 'size' of the story element.
- **Effects:** it is a list of events which are supposed to be caused by this story element.
- **Subject:** it is the subject of the historical fact encoded by the story element.
- **Object:** it is the object of the historical fact encoded by the story element.

Let us show some simple examples:

(mexican-independence type national-independence)
 (mexican-independence date year 1810)
 (mexican-independence location any)
 (mexican-independence attributes historical 1)
 (mexican-independence subject Mexican)
 (mexican-independence object Spanish)

(astronomic-observation type scientific-observation)
 (astronomic-observation date any-night)
 (astronomic-observation location Palace-Tower 1)
 (astronomic-observation attributes scientific 1)
 (astronomic-observation attributes cultural 0.5)
 (astronomic-observation subject Mayan)

Common sense knowledge

These are rules representing simple commonsense cautions. These rules are used to enhance the story elements previously selected to be told once these story elements have been translated from the guide perspective. Therefore, this knowledge allows the guide to extend the basic information represented as story elements by using her 'supposed' commonsense.

Relations between guide profile features and story element attributes

These are weighted relations that are established between interests and roles on the one hand, and attributes

on the other hand. These relations describe the importance that a particular attribute could present for a guide with a particular interest or role.

Meta-rules to translate story elements from the guide perspective

These meta-rules allow to translate story elements to show a particular guide perspective. For example, a story element describing a battle between Spaniards and Mexicans where Spaniards lose, will be translated to a new story element showing the guide profile implication (if the guide is Spanish the new story element will reflect a loser perspective, while if the guide is Mexican the new story element will reflect a winner perspective).

Rules for text and special effects generation

These are rules and templates that allow to generate the text and special effects corresponding to a particular story element, once it has been translated (from the virtual guide perspective) and enhanced.

Guide expression rules

These are rules which allow to translate guide attitudes into guide actions that allow her to express herself. These rules are guide dependant, that is, every guide has her own set of expression rules.

Narrative construction

In our model the guide begins at a particular location (manually fixed or randomly chosen) 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. Therefore, our guide tells stories by improvising, at each step, where to go and what to tell next, and this improvisation tries to emulate what occurs in the mind of a real guide.

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 (contained in the script provided by his company), 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. Figure 3 shows a diagram of the overall functioning. The processes and mechanisms involved in the algorithm can be separated in three global parts. These three global processes are carried out with every step. First, the guide should decide where to go and what to tell there. Actually, she only decides something to tell but not all the information that she will tell. That is, she decides that she wishes to tell something. Then, during the next process, she extends this information with additional story elements, translates these story elements from her own perspective, and again extends the selected information depending on her perspective, that is, with information showing her viewpoint. Finally, the system generates a story. The next three subsections detail these general phases.

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, that is, some concepts, objects or characters used in previously narrated story elements trigger off recollections in his mind. 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 fact, this could be a special case of the above commented remembrance by association, as the guide's interests and roles could be seen as persistent triggers of recollections in the guide's mind. 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




Fig. 3 Algorithm diagram

candidate pair. The weight of each score in the combination process can be tuned from the administrator interface, as shown in Fig. 4. Finally, the system chooses the pair with the highest overall score value.

Extending and contextualising the information

Figure 5(a) 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. Two story elements X and Y are connected by a cause-effect link if one of the following cases is fulfilled: X is cause of Y; X is effect of Y; X and Y are causes of the same third story element Z; X and Y are effects of the same third story element Z. A subject-object link between the story elements X and Y is established if any of the following cases is fulfilled: the subjects of X and Y are the same; the objects of X and Y are the same; the subject of X is the object of Y; the object of X is the subject of Y.

Figure 5(b) 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 criterion 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 5(c) shows the same part of the memory where three additional story elements have been selected by navigating from the original story element. The system allows the user to tune the relative weight of




Fig. 4 Administrator interface




Fig. 5 Storyboard construction

each criterion in this memory navigation process as well as the desired degree of granularity (see Fig. 4).

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 Fig. 5(d)). 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. We 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(Colonizer) =>
fact(colonizer Colonization, Colonizer, Colonized) and
guide attitude(contradictory)

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

meaning that a colonizerColonization fact and a contradictory guide's attitude should be inferred if a colonization fact is included in the story element and the guide profile matches the second argument of this fact,

that is, the guide is the Colonizer. In the same way, 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. Whatever rule is activated, 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 the guide to add some comments showing her perspective. The guide uses the new contextualised story elements (Fig. 5(d)) as input for the rules that codify commonsense (Fig. 5(e)). By inferring these rules the guide obtains consequences that are added to the contextualised story elements, obtaining a new data structure (Fig. 5(f)) which codifies the information that should be told. We continue with the previous example. Let us assume the following commonsense rules:

```
fact(colonizerColonization, Colonizer, Colonized) =>
    fact(culturalImprovement, Colonized) and
    fact(religiousSalvation, Colonized)
fact(colonizedColonization, Colonizer, Colonized) =>
    fact(culturalDestruction, Colonized) and
    fact(religionChange, Colonized)
```

meaning that the colonizer's view of a colonization process implies the cultural improvement and religious salvation of the colonized. The view of the colonized implies the destruction of their culture and the change of their 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.

Generating the story

As a result of the previous processes, the guide obtains a set of inter-related information items to tell (Fig. 5(f)). 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. As the information should be told in a sequential way, the first step is then to order the data elements. Several criteria can be considered. In particular, we use the three following concrete criteria (its weights can be tuned from the administrator interface, as shown in Fig. 4):
 - *Cause-effect*: the elements are ordered in such a way that 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.
 - *Classic climax*: the first selected story element (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). By this process, for example, a Mayan guide will show an anger gesture when talking about the Spanish colonization of the Mayan territories, while a Spanish guide will show a distress gesture when talking about the Mayan religious sacrifices.
 4. Every story element in the knowledge base can be associated to particular environment conditions or special effects. Thus, finally, a storyboard like the one shown in Fig. 5(g) is obtained.

Implementation

As pointed out in Laird [20], although the development of realistic virtual environments is an expensive and time-consuming enterprise requiring expertise in many areas, computer games provide us with a source of cheap, reliable, and flexible technology for developing our own virtual environment for research. During the last few years, artificial intelligence researchers have been using several computer game engines (Descent 3, Quake II, etc.) to check their algorithms.

The current trend is to use the Unreal Tournament (UT) engine as the development platform. UT is a cheap game which provides access to level editors for defining the environment, a scripting language (UnrealScript) for defining the physics of the world and the way objects interact, and the ability to import your own objects into the game. UT is currently being used by several groups to develop storytelling systems [21–23], and ai-based bots in general [20].

Gamebots [24] is a modification to UT that allows characters in the game to be controlled via network sockets connected to other programs. For a detailed description of Gamebots, see Kaminka et al. [25].

We have chosen UT as the platform on which our virtual worlds run. As we wished our system to be open and portable, we decided to use Gamebots to connect our virtual guide to UT. We have also chosen a second virtual environments platform to check that our guide is easily portable. In particular we have selected Deepmatrix [26], a multi-user VRML application that enables




Fig. 6 Implementation diagram

several users to share a single VRML world and interact with it. For a review of Internet-based virtual environments, see Hughes et al. [27].

We have developed the proposed system. Figure 6 shows a diagram of the current implementation. The core of the virtual guide is a Java application which is able to connect to both UT worlds (through Gamebots) and VRML worlds (running on a Deepmatrix Server). 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 [28] database to store the knowledge base containing the information and the meta-information. The Java application accesses this data through JDBC. The developed system uses Jess to carry out inferences on the information, that is, to reasoning. Jess [29] is a rule engine and scripting environment written entirely in Java language, and originally inspired by the CLIPS expert system shell.

Scenarios and preliminary results

We are using our system in two different virtual environments. First, we started to use it in a virtual representation of Palenque, the ancient Mayan city. In this virtual city the guide tells stories about Mayan culture, religion, science, architecture, etc. Figure 7 shows the guide in Palenque.

As far as we have tested the system, it seems that it works properly with small knowledge bases. Figure 8 exposes an extract from the text generated by an occidental female virtual guide in Palenque. In particular, the text was generated by the guide when pointing to a typical Maya home. We still have to check how the system behaves when dealing with large knowledge bases and evaluate different template systems for the generation of texts from the storyboard.

To evaluate the system in a more complex environment (in terms of the knowledge about the places and objects), we have defined another scenario. It is a home inhabited by three generations of the same family:

Fig. 7 Virtual guide in Palenque world



Fig. 8 Text generated by an occidental female guide in Palenque

Mesoamerican homes were not just safe places to eat and sleep, they were workplaces too. There were no refrigerators or household appliances, so women had to work hard preparing food for the day's meals or for winter storage. It is always the same story, women sacrificing themselves for the family! Vegetables were cleaned and chopped with stone knives, as there were no metal ones. Beans and chillies were spread out in the Sun to dry, and maize kernels were ground into flour.

Women and girls spent long hours spinning thread and weaving it into cloth, then sewing it into tunics and cloaks for the family. Some women wove cloth to sell or to give to the government as a tax payment. Women slaving away by sewing! Homes were also where most sick or elderly people were cared for.

grandparents, parents and children. Most of the rooms and objects in this house are associated with plenty of memories related to the members of the family. The children have listened to their parents and grandparents telling them stories about the pictures, books, etc. in their home. Therefore, all of the children are supposed to have almost the same shared memories. However, each child has a well-defined profile, different from these of her brothers and sisters. Then, when a user visits the place, he is welcomed by one of the children who tells him stories about her family history (in this home) from her own perspective.

Conclusions

In this paper, we have described 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 model of the Mayan City of Palenque, used in this work so far, and are working on a model of a further city, 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.

References

1. Tozzi V (2000) Past reality and multiple interpretations in historical investigation. *Stud Social Political Thought* 2
2. Mateas M (1999) An Oz-centric review of interactive drama and believable agents. In: Wooldridge M, Veloso M (eds) *AI Today: Recent Trends and Developments*. Springer, Berlin, NY
3. Mateas M, Stern A (2000) 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. AAAI Press
4. Polt G (1922) *The 36 Dramatic Situations*. The Writer, Boston
5. Propp V (1928) *Morphology of the Folktale*. University of Texas Press
6. Thompson S (1955–58) *Motif Index of Folk Literature*. Indiana University Press, Bloomington
7. Branigan E (1992) *Narrative comprehension and film*. Routledge, New York
8. Murray J (1998) *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*. MIT Press, Cambridge, MA
9. Braun N (2002) Automated narration – the path to interactive storytelling. 2nd International Workshop on Narrative and Interactive Learning Environments, Edinburgh, Scotland
10. Spierling U, Grasbon D, Braun N, Iurgel I (2002) Setting the scene: playing digital director in interactive storytelling and creation. *Comput & Graphics* 26:31–44
11. Brooks KM (1999) *Metalinear Cinematic Narrative: Theory, Process, and Tool*. PhD Thesis. School of Architecture and Planning
12. Aylett R (1999) Narrative in virtual environments – towards emergent narrative. Narrative Intelligence Symposium, AAAI 1999 Fall Symposium Series
13. Louchart S, Aylett R (2002) Narrative theory and emergent interactive narrative. 2nd International Workshop on Narrative and Interactive Learning Environments, Edinburgh, Scotland
14. Liu H, Singh P (2002) MAKEBELIEVE: using commonsense to generate stories. Proceedings 20th National Conference on Artificial Intelligence (AAAI-02), Edmonton, Canada
15. Domike S, Mateas M, Vanouse P (2002) The recombinant history apparatus presents: Terminal Time. In: Mateas M, Sengers P (eds) *Narrative Intelligence*. John Benjamins, Amsterdam
16. Isbister K, Doyle P (1999) Touring machines: guide agents for sharing stories about digital places. Working Notes AAAI Fall Symposium on Narrative Intelligence
17. Singh P (2002) The public acquisition of commonsense knowledge. Proceedings AAAI Spring Symposium on Acquiring (and Using) Linguistic (and World) Knowledge for Information Access, Palo Alto, CA
18. Dautenhahn K (1998) Story-telling in virtual environments. Working Notes: Intelligent Virtual Environments, Workshop at the 13th Biennial European Conference on Artificial Intelligence (ECAI-98), Brighton, UK
19. Saretto CJ (n.d.) A Survey of Interaction Issues Facing Navigation in Virtual Worlds. http://www.ndenial.com/docs/navigation_survey.pdf
20. Laird JE (2002) Research in human-level AI using computer games. *Comm ACM* 45: 32–35
21. Robertson J, Oberlander J (2002) Ghostwriter: educational drama and presence in a virtual environment. *J Comput Mediated Comm* 8. <http://www.ascusc.org/jcmc/vol8/issue1/robertson/robertson.html>
22. Young RM (2001) An overview of the Mimesis architecture: integrating intelligent narrative control into an existing gaming environment. Working Notes, AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment, Stanford, CA
23. Cavazza M, Charles F, Mead SJ (2001) Characters in search of an author: AI-based virtual storytelling. International Conference on Virtual Storytelling. Springer, France
24. Adobbati R, Marshall AN, Scholer A, Tejada S, Kaminka GA, Schaffer S, Sollitto C (2001) Gamebots: A 3D virtual world test-bed for multi-agent research. Proceedings Second International Workshop on Infrastructure for Agents, MAS, and Scalable MAS, Montreal, Canada
25. Kaminka GA, Veloso MM, Schaffer S, Sollitto C, Adobbati R, Marshall AN, Scholer A, Tejada S (2002) GameBots: a flexible test bed for multiagent team research. *Comm ACM* 45: 43–45
26. Reitmayr G, Carroll S, Reitemeyer A, Wagner MG (1999) DeepMatrix – an open technology based virtual environment system. *Visual Comput* 15:395–412
27. Hughes CE, Moshell JM, Reed D (2001) Internet-based virtual environments. In: Stanney KM (ed) *Handbook of Virtual Environments: Design, Implementation, and Applications*. Lawrence Erlbaum
28. MySQL: The World's Most Popular Open Source Database. <http://www.mysql.com/>
29. Jess the Rule Engine for the Java Platform. <http://herzberg.ca.sandia.gov/jess/>