Touring Machines: Guide Agents for Sharing Stories about Digital Places

Katherine Isbister
Open Laboratory
NTT Communication Science Laboratories
2-4, Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-0237
JAPAN
katherine@digitalcity.gr.jp

Patrick Doyle
Stanford University
Gates Computer Science Building 2A
Stanford, CA 94305-9020
pdoyle@cs.stanford.edu

Abstract

Human understanding of the physical world frequently derives from the narratives we construct about it. Analogously, we can incorporate narratives in virtual spaces to give visitors a sense of place and purpose. In this paper we describe two projects using interactive characters as a medium for story building and storytelling within and about the digital places they inhabit.

Introduction

According to (Schank and Abelson 1995), the bulk of human knowledge and memory is communicated and encoded in story form. Although there has been a long-standing interest in how interface agents might bring narrative perspective to information resources (Oren et al. 1990), there is currently little explicit use of adaptive storytelling in the repertoire of those agents. Most of the interface agents in widespread use today are only embodied alternatives to traditional menu- or prompt-driven mechanisms for performing simple tasks such as searching for files or providing context-sensitive help. Although they possess bodies, these utilitarian agents ordinarily lack intelligent believability (Ball et al. 1997; Rist et al. 1997). There has also been active research on pedagogical agents that do exhibit intelligent behavior, though they create narratives only in the sense that their pedagogical requirements impose a structured sequence of behaviors (e.g. (Rickel and Johnson 1997; Lester and Stone 1997)). Finally, there has been work on the creation of stories through the interaction between the user and agents that are characters in the story (Bates 1992; Hayes-Roth et al. 1997). In these systems, however, story is the purpose of the interaction.

In contrast, our interest is in the explicit use of narrative as a means to help users organize familiar ideas, learn new material, and more effectively engage them with this content. In particular, we are interested in situated stories—stories that are tied closely to the environment—that can illustrate it and bring it to life for the visitor. In this paper, we will discuss two projects that attempt to incorporate situated storytelling into the behavior of an agent.

Both of our projects involve creating interface agents that can act in the human role of "guide" to a digital location. In examining qualities a human guide must possess, we learned that interpretation of the location is one of the most important tasks. *Interpreting* is bringing a place to life using one's own perspective and artifacts at hand, and usually involves telling stories to visitors, to help make the experience more vivid and emotionally engaging to them (Pond, 1993).

We are each exploring different interpretation contexts and strategies. Isbister is developing an agent that is tightly integrated into an online 3-D website tour. Her project explores ways to make an agent's narrative effectively adaptive to different groups who take the tour, like an expert human tour guide. Doyle's work explores ways to enhance exploration of Web sites by enabling a user's persistent companion agent to behave intelligently in unfamiliar contexts. One of the intelligent behaviors he is examining is narrative guidance through these sites, driven by a scalable annotation mechanism that works within existing Web standards.

We will briefly summarize the structures of our projects and how we are attempting to build guides capable of these sorts of interpretations. We discuss the direction this research is taking, and close with a list of challenges in building agents as tour guides, challenges that we believe also serve as guidelines for their construction.

Kyoto Digital City Tour Guide

Isbister is currently part of a team that is building a digital version of Kyoto, Japan (http://www.digitalcity.gr.jp). This online city is meant to be an evolving resource both for outsiders and for Kyoto residents, with a rich store of useful local information. It currently consists of a database of geographically-arranged links to local websites, with 2-D and 3-D interfaces for browsing these links. The 3-D interface to the city is being constructed in a new development environment called 3DML. 3DML uses texture mapping and simple geometric forms to enable users to rapidly create photo-based 3-D worlds. The worlds can be explored using Doom-like navigation. Visitors can also click on buildings to open related Web links.

One of Isbister's tasks is to ensure that the Kyoto digital city will be an inviting place for real people to use—a truly useful and viable tool. As a visitor herself, she is focusing on making the site approachable and engaging for outsiders who might someday want to visit Kyoto. To help accomplish this goal, she is creating an agent-led group chat tour of the city. The tour will be a point of entry to the online resource and to Kyoto itself, ideally increasing visitor interest in and use of the digital city. Isbister hopes the group tour will also encourage dialogue and relationships among those who participate, as well as increase exposure to Kyoto's history among those who are friends and family of people who participate in the tour.

In creating the agent, Isbister is focusing on developing storytelling strategies that will produce an engaging experience for tour takers. She derived a list of target abilities for the agent by researching the behavior of actual tour guides in Kyoto. Isbister hopes that the process of developing the agent's storytelling abilities will lead to a contribution to the narrative intelligence/agent research domain, specifically involving timing and duration strategies for situated storytelling, especially to groups.

Researching the Tour Guide Role

In preparation for creating the tour guide agent, Isbister went on several guided tours of Kyoto, making notes about how the tour guides did their work. She 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 its history of destruction and reconstruction; descriptions of peak historic events that happened at a given site; and descriptions of seasonal events and customary activities that occur at the site. Tour guides also reincorporated material from previous tours, describing what other visitors said and did when they visited the sites.

The tour guides would introduce the stories upon arrival at the site, and at specific points in the site that were directly relevant to the particular story. 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.

The tour guide would time his or her storytelling to the visitors' absorption with the site at hand. That is, the guide would provide story context while the visitors were engaged with the site, and move on when it became clear that the visitors were through looking at the site and were ready for a new venue. In addition, the guide would provide follow-up stories to those that were met with particular interest by visitors.

Isbister once worked as an exhibit sign writer at a zoo, and observed another quality important in developing tour guide agents. At the zoo, parents often read the exhibit text and then retold the contents of the signs to their children on the spot. For example, children were most likely to be interested in stories about how animals behave or their unusual characteristics. So a parent would find a fact like this on the sign, and turn and tell the child this particular piece of information, omitting the rest. Allowing for this kind of content repurposing among audience members is another important goal to aim for in designing agent guides, and the stories that they tell. This allows the stories to serve a social as well as an informative purpose.

To summarize, Isbister observed the following qualities in human tour guide storytelling to groups:

1. Stories were told about particular locations, while visitors looked at those locations.
2. Some stories included material that tour guides reincorporated from previous tours.
3. Stories that the guide selected seemed to be those that would lead to easy retelling by visitors to others.
4. Guides adjusted the timing of storytelling as well as choice of follow-up stories based on visitor interest level.

Tour and Guide Design

Isbister used the storytelling criteria above to guide the development of a prototype group tour guide, which leads a tour of Kyoto's Nijo Castle on the Web.

The Kyoto Tour Guide project has four main components: 1) 3-D explorable tour sites; 2) a database of gesture-choreographed stories that relate to these sites, which are performed by the Microsoft Agent; 3) a commercial chat server (I-chat); and 4) an agent, written in perl, that drives the tour and the performance of the Microsoft Agent.

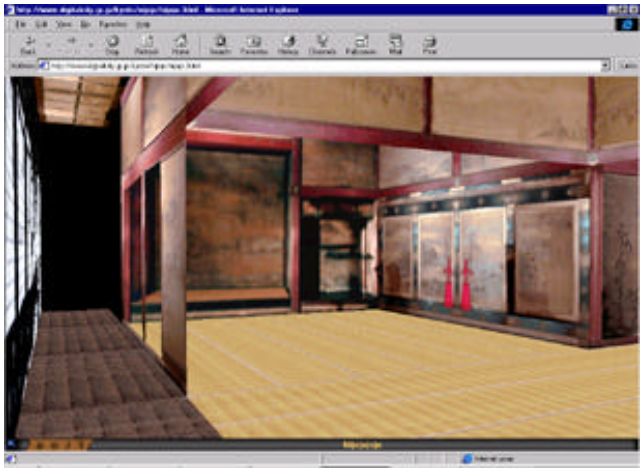

Figure 1. Screenshot of the Shogun's private quarters, one of the 3DML stops on the tour.

1) 3-D tour sites: The prototype tour currently consists of three 3-D sites, though we anticipate building two more sites before testing online. All of the sites are locations in Kyoto's Nijo Castle. Nijo Castle served as the Kyoto base for the Tokugawa Shoguns during the Edo period in Japan. We chose Nijo Castle for the prototype for several reasons: it was built in an important historical period in Japan and so it has great historic significance; it can be used to discuss Japanese lifestyle and politics of that period; and its architecture lends itself well to the 3-D tool that we used, which supports rectilinear forms especially well (see Figure 1).

2) Story database: We have created and are currently expanding a database of stories that are related to the sites on the tour. These stories have been crafted according to above criteria: anecdotes are selected that relate closely to the particular sites, and that reveal interesting and easily-retold information about Japan's history and culture.

Each story includes both dialogue and gestures that the Microsoft Agent character will perform. The gestures add emotional as well as informational value to the stories (Isbister and Ishida 1999). The agent's spoken delivery of the dialog allows visitors to devote more of their visual attention to the site that they are exploring. The stories are HTML files containing JavaScript commands that drive the MS Agent. The database currently includes three versions of each story—short, medium, and long—to be delivered by the guide depending upon the level of user interest and activity during the tour. The story database is implemented using PostgreSQL ver. 6.4.2, running on a Sun with the Solaris operating system. Currently, the database is a simple table that allows lookup and selection of stories based on length, title, and the tour stop to which they relate. In the future, we plan to categorize these stories by content type as well, to allow for adaptive story-telling based on tour-takers' specific topic interests. We would also like to develop a way to rank the stories based on user feedback as to which are the most enjoyable and interesting, to help keep the tour fresh and reflective of the latest visitor interests.

3) Chat server: The tour is hosted using a commercial chat server called I-chat. I-chat makes it easy to associate particular Web pages with chat rooms, and to push new pages to all chat room participants (see Figure 2). Tour takers log into our local I-chat server and are in the same chat room for the duration of the tour. I-chat's implementation also makes it easy to create an agent that can log into the chat environment in the same way that a user does. The perl agent that drives our guided tour is logged into the tour chat room, and can easily monitor and log user activity, in order to make the story selections that it will push as HTML pages to all users' Web browsers.

4) Tour agent: The tour agent, written in perl, makes decisions about what story to tell at each tour stop. This agent is logged into the chat room, and is able to monitor the group's conversation. The current implementation tracks

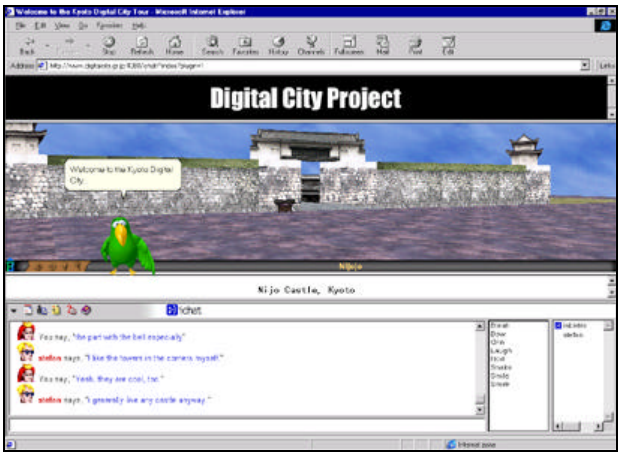

Figure 2. Prototype of the Kyoto Digital City Tour

the quantity of conversation, and looks for positive and negative keywords that indicate how visitors feel at the moment (negative words include words such as "boring, dull, too long"; positive words include words such as "wow, cool, neat, interesting"). The current prototype agent selects stories using a very simple decision rule (see Figure 3). We plan to adjust the agent's decision-making mechanism after examining its initial performance with tour-takers.

	Valence of Conversation Contents	
Quantity of Talk	***Negative***	***Positive***
Low	medium length	long length
High	short length	medium length

Figure 3. Decision Rule for Prototype Tour Guide Agent

To make sure the tour stops long enough (but not too long) at each tour location, the agent asks visitors (through dialog delivered by the Microsoft Agent) to provide an explicit verbal cue that they are ready to go on. The agent moves on to the next tour stop when it gets feedback from the majority of tour-takers that they want to move forward. We believe this explicit request for feedback serves two purposes: it allows the tour-takers to modulate the stop length far more subtly than the agent could, and it gives them a feeling that the tour (and guide) are adjusting to them—that they have an influence on events, and needn't feel hurried or bored.

There is also a post-tour questionnaire that asks visitors who have taken the tour to give feedback about the stories and the experience itself. As mentioned above, we hope to develop a way to feed this information back into the ranking and selection of stories by the tour guide agent.

Development and Evaluation Plans

We are currently completing the prototype tour, and plan to test the tour out on website visitors during September. We will use both the guide logs (which include decisions made by the agent about story length for each stop, as well as the conversation transcript) and the survey results to determine how effective the group tour agent was as an adaptive guide and storyteller. We will look for indications of visitor enjoyment, retelling and discussion of story topics, as well as extended user interaction in the chat room after the tour is done, to judge whether our agent was a success as cultural ambassador and dialogue-builder for the Kyoto Digital City. We also plan to use the transcripts to see if our ideas about group conversation tracking and interest monitoring have merit for developers of interactive agents for group settings like this one.

An Annotation-Sensitive Tour Guide

Doyle has arrived at this problem from a different direction. Given the growing size and popularity of the Internet, his concern has been to find mechanisms that enable interface agents of varying degrees of sophistication to act intelligently and believably across a wide array of sites. His approach has been to add *annotations* to these sites. Annotations are declarations and procedures embedded in the environment and made accessible to wandering agents. They provide useful information about available content and activities, as well as assisting agents in the choice and timing of their actions on these sites.

Previous explorations have taken place in text-based environments called MUDs, where agents have used annotations to solve puzzles, play games, and enhance their emotional intelligence with respect to the environment (Doyle and Hayes-Roth 1998a; Doyle and Hayes-Roth 1998b). While the World Wide Web does not offer as sophisticated or flexible a platform for agent behavior, its ubiquity and obvious amenability to simple forms of annotation make it a natural choice for exploring the annotation approach.

Web sites are organized collections of documents. While most don't explicitly tell stories, visitors can understand and make use of them by creating internal narratives about them – what their purposes are, how they're used, what they make available. The opportunity for narrative comes about in two ways. The first is the hyperlink topology of the site; the connections among a site's documents allow us to build corresponding connections in our minds, a narrative of navigation. A path of followed links forms a story of a sort.

The second source of narrative is the site's contents. If people understand the world in terms of stories, our tendency will be to create or associate stories with the text and images we encounter as we navigate the site. This suggests that the site's designer may be able to communicate more effectively by actively offering narratives to site visitors, rather than requiring them to create their own.

Doyle's work on annotations for enabling agents and enhancing sites seemed a natural mechanism for exploring this idea. By embedding details of the site's content as well as actual stories on the site in a form that the agents can make use of, we can produce an adaptive guide that has both knowledge of the user and an understanding of the site. This guide can entertain or edify according to the user's interests while at the same time furthering the annotation designer's goals for the site.

The agent operates in a split-frame window of a Web browser. The upper frame, occupying most of the window, contains the current site's content. The lower frame holds JavaScript code for interfacing between the agent, the visitor, and the site, and displays buttons the visitor can use to command the agent. The agent's architecture (its intelligence, behaviors, and annotation handling) is written in Java. Hooks to the Microsoft Agent API are used to provide the animated character and speech synthesis. The lower frame persists as the user navigates the Web, so while the agent is an extra layer of interface beyond the standard browser, it does not prevent the user from visiting any Web sites, whether or not they are annotated.

Web annotations are represented in the Extensible Markup Language (XML). Every page on an annotated site has a corresponding XML file containing its annotations. Whenever a user visits a new page on the site, the guide agent requests the annotation file from the server, parses it, and adds that information to its local store. If there are no annotations available, or if they are useful only to other kinds of annotation-sensitive agents with different abilities or goals, the agent can still fall back on its built-in repertoire of behaviors. The agent does not require that the annotations be there, but it will make use of them if they are.

A Virtual Art Museum

Our current testbed on the Web is an art museum. The museum consists of a set of galleries, each of which contains rooms filled with artworks. There are presently two galleries: the *Nativity Exhibit* houses medieval and Renaissance religious art revolving around the birth of Christ, while the *Pre-Raphaelite Exhibit* features 19th century artworks in the style of that movement.

Within these galleries, annotations provide a visiting agent with details about their artworks, including their historical contexts, the lives of their artists, their relationships to other artworks in the gallery, and details about specific characteristics of the art. Any agent familiar with the artwork ontology can immediately extract this information for whatever purpose.

The Web agent used in the virtual art museum is a character called Merlyn, aptly named after the forgetful old wizard of T.H. White's *The Once and Future King*. Figure 4 shows Merlyn in the Briar Rose Room of the Pre-

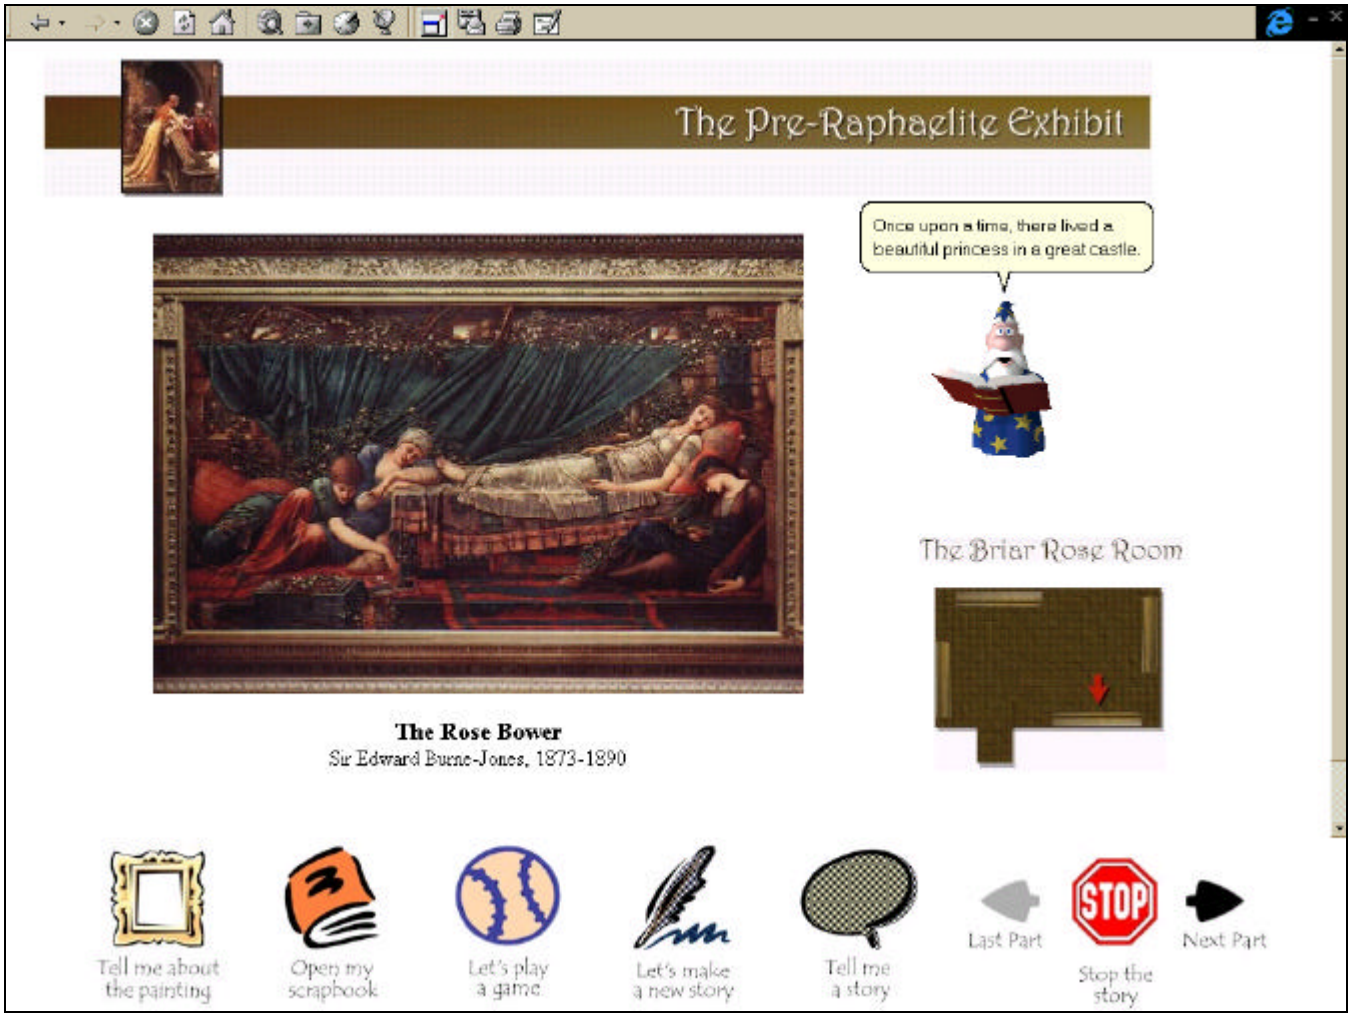

Figure 4. Merlyn in the Briar Rose Room

Raphaelite gallery, in the process of telling a story to the user.

Merlyn's purpose is to explore the art museum together with a child. To do this he uses the annotations in several ways. First, he can describe the paintings he and the child see as they travel through the museum; he can provide information about when they were made, by whom and how. This information is available both on demand and through his autonomous lecturing behaviors. Since he retains these annotations as he travels, he will also be able to refer back to them if they relate to current topics ("Remember the other picture of Tristan we saw?")

He can also use the annotations to play games with the user; "I Spy" was the first we implemented because of its simplicity[1]. More elaborate games involving names, dates, and connections between rooms are in the works.

[1] "I Spy" is a children's game in which the participants take turns reciting the rhyme, "I spy, with my little eye, something beginning with (some letter of the alphabet)." The challenge is to find the desired object.

The agent's third major function is storytelling. Merlyn uses his internal database or the annotations in the museum as his sources. One of the implemented tales in the Pre-Raphaelite gallery, for example, is the story of Sleeping Beauty, which is told using Edward Burne-Jones' four "Briar Rose" paintings as illustrations. This is the story Merlyn is reading to the user in Figure 4. He will autonomously offer to tell stories when he encounters new ones, and the user can direct him to stop or move forward or backward through the story through the buttons on the control panel. Merlyn uses the combination of an internal timer and a sensor that monitors the user's actions and will offer to proceed if the user appears to be bored or finished with the page.

In the art museum ontology, a story is a kind of *tour*, and consists of a sequence of artworks on Web pages, each of which has one or more pieces of dialog associated with it. Each artwork may be associated with one or more tours, and within each tour different pieces of dialog may be associated with different categories of reader (so the rendition of the Sleeping Beauty story might be different for a

```
<TOURS>
 <TOUR NAME="Sleeping Beauty" TYPE="story">
 <TOPIC>the legend of Sleeping Beauty</TOPIC>
  <TEXT TARGET="ALL">
    Once upon a time, there lived a
    beautiful princess in a great castle…
    . . .
  </TEXT>
  <STOP>
   <NAME>Briar Rose Room South Wall</NAME>
   <URL>gallery/briar/briar4.html</URL>
  </STOP>
  <NEXT-STOP>
   <NAME>Briar Rose Room North Wall</NAME>
   <URL>gallery/briar/briar2.html</URL>
  </NEXT-STOP>
 </TOUR>
</TOURS>
```

Figure 5. Fragment of an XML Tour Annotation

child and adult). A tour can be any sequential narrative, and not simply a story; in addition to hearing about Sleeping Beauty, a visitor might also opt to hear the history behind the painting of the Briar Rose series. In either case the active process of narration, we believe, results in a more attentive, more engaged audience, and there is reason to believe it improves recall of the material as well (Lester et al. 1997).

Unlike other annotations that are highly structured text fragments, tour dialog is stored in the annotation files in full English text. While Merlyn can use those fragments to build natural-language descriptions of the art and artists, the tour text quickly moves outside the range of the ontology. Instead, he takes the English text and alters it through simple syntactic manipulations, such as pauses, stutters, or interjections. This allows him to adapt it, albeit in a limited way, to his personality, his emotional state, and to the style of interaction he is using with his audience. In this case we are sacrificing flexibility for the sake of a well-scripted tale. Ultimately we would prefer an agent and a markup language sophisticated enough that we could build English text from a highly structured description of a story, but the magnitude of such a problem is well beyond the scope of our work.

Since the agent stores the annotations it encounters, it retains a memory of every story it sees, as well as where these stories came from. He can easily track requests for stories to be retold (or requests to stop telling others), so he can infer from their topics and keywords what other known and annotated locations might be of interest to his interactors. This sort of simple adaptation to the user's preferences is a necessary first step in making these embodied agents relevant and entertaining for their audiences, which is itself necessary if they are ever going to have real value as a new kind of interface.

Extensions

Merlyn's repertoire of tales comes strictly from the environment in the current implementation. We are in the process of allowing users to create their own stories, which Merlyn learns and stores to tell again later. Using the *Scrapbook*, which records snapshot images and basic facts about every artwork the user sees, it is possible for the user to build new stories by specifying sequences of artworks and appropriate corresponding text. Merlyn will remember these tales and offer them to their author or any new children with whom he interacts. Thus the user can "teach" the agent and extend its abilities over time in the same kinds of ways that the environment does.

Currently all annotations on our Web sites are in XML data files. This is fine for representing static information, but more adaptivity on the part of the site may be desirable. A first step in achieving this is to make the XML output script-generated, so agents may provide some details about themselves and their abilities, and the site can take those details into consideration when producing its output. An agent specializing in art technique, for example, doesn't need to be given historical or biographical information. An agent specializing in storytelling may want more stops and longer text to offer a visitor.

This is still very limited. A more ambitious approach is to provide the agent with code or binaries to execute on the client side as a means to expand its abilities. Since our agent is written in Java, we can send it modules it may dynamically plug into itself so, for example, a site that allows children to play games, draw or write might provide visiting agents with the means to detect and participate in these activities. The idea of a plug-in is commonplace for Web browsers, but when we consider plug-ins as a way to give our characters new abilities, their possibilities become much more exciting. Future work here in our Web environment as well as our MUD-based world will explore this issue in greater detail.

Challenges

While interactive tour guides offer the possibility of increased user attentiveness, retention, and enjoyment (Lester et al. 1997), these benefits require we make our agents believable to their users, both as living entities and as expert guides. Our research suggests that the following four traits are critical for creating believable and compelling guides. The first three points have emerged from Doyle's work on annotation-based tour guide storytelling; the last has emerged from Isbister's investigation of human tour guides leading group tours.

- *Intelligent reincorporation*. Reincorporation in this context means the reintroduction of ideas or entities that have been seen earlier in the tour, with some reference to how they relate to the topic under discussion. In improvisational theater (Johnstone 1992) it is well

understood that reincorporation is a key to building a story satisfying to the audience; a sequence of unrelated events does not make a story. Similarly, we suspect that a key to creating an effective tour narrative is reincorporation of earlier material. This requires that the agent not only track what material has been seen but also when connections are either pedagogically or dramatically effective.

- *Empathy with the content*. (Elliott et al. 1997; Elliott 1998) has argued that understanding of and reasoning about narrative is strongly associated with reasoning about emotion. If we think of stories as descriptions of sequences of events that have emotional associations for the listener, then we can enhance the significance of our tour content to the user by infusing emotion into the presentation. Beyond incorporating emotions in the presented material, we can build our agents to react emotionally to what they present. This not only heightens the significance of these narratives, it also enhances the believability of the tour guide.

- *Presentation through personality*. The most obvious flaw of an intelligent character that relies upon prewritten text is that the text must either be designed for that particular personality, rendering it dangerous for other personalities to use, or so devoid of character that the recitation seems stilted and unbelievable. Ultimately, one would desire an annotation language and an agent powerful enough that the concepts could be explained, reinterpreted, and formed into dialog by the agent. Unfortunately, this is still a hard unsolved problem. Syntactic sugar is a simple approximation that frequently provides good results, as users come to associate idiosyncratic verbal behaviors with the agent's personality. Nonetheless, a more intelligent mechanism for integrating content and presentation will ultimately be required for widespread use.

- *Artful timing/delivery*. As noted in the section on human tour guides, successful storytelling involves knowing when to begin a story, how long it should last, and whether one should elaborate with related stories. Making good decisions about timing and delivery requires the ability to detect user interest, and react appropriately to it. Human beings use many subtle cues to indicate low or high engagement, and detection of user interest is currently very primitive. This will continue to be an important area of research, both for development of characters and for development of satisfactory and subtle interactions with interfaces, in general.

Future Plans

Both Isbister and Doyle are interested in incorporating visitor reactions and stories into the repertoire of tour guide agents. Isbister is interested in allowing Web site hosts (such as shop owners or temple leaders) to provide stories that can be retold by the tour guide agent to visitors during the tour. She is also experimenting with ways to solicit direct visitor input about their interest in, and engagement with, stories during the tour, without interrupting the flow of the experience.

Doyle is expanding the virtual art museum with additional areas in order to more fully explore the problems associated with annotated environments as those environments grow in size. Merlyn is being rewritten as part of an effort to build a unified basic annotation-sensitive believable agent architecture, one that can support a variety of interfaces and types of environmental annotations. He hopes to employ adaptive modeling of the user both to enhance believability and improve the agent's effectiveness as a guide.

Conclusions

Storytelling is an important part of a human tour guide's social role. By telling stories, guides put what might be otherwise dry or overwhelming information into structures that give it meaning and social value. Using reincorporation, emotion and personality, and artful timing, human guides can present information about tour sites in appealing and engaging ways. We anticipate that the lessons we are learning designing interface agents according to these principles will allow us to produce sophisticated storytelling agents, and will be useful to others designing characters to interact in social settings and tasks.

Acknowledgments

Patrick Doyle's work has been done as part of the Virtual Theater Project at Stanford University, under the direction of Barbara Hayes-Roth. He was supported by NSF Grant IRI-9424261, gifts from Microsoft and Intel, and an Intel Foundation Fellowship. Katherine Isbister would like to thank NTT Open Lab and Omron Software researchers who contributed to this project, as well as the students at the Stanford Japan Center.

References

Ball, G.; Kurlander, D.; Miller, J.; Pugh, D.; Skelly, A.; Stankosky, A.; Thiel, D.; van Dantzich, M.; and Wax, T. 1997. Lifelike computer characters: the Persona Project at Microsoft Research. In J. Bradshaw, ed., *Software Agents*. Menlo Park, CA: AAAI Press.

Bates, J. 1992. The nature of characters in interactive worlds and the Oz project. Technical Report CMU-CS-92-

200, School of Computer Science, Carnegie Mellon University.

Doyle, P., and Hayes-Roth, B. 1998a. Guided exploration of virtual worlds. In F. Sudweeks, M. McLaughlin and S. Rafaeli, eds., *Network and Netplay: Virtual Groups on the Internet*. New York: MIT Press.

Doyle, P., and Hayes-Roth, B. 1998b. Agents in annotated worlds. In *Proceedings of the Second International Conference on Autonomous Agents*. New York: ACM Press.

Elliott, C.; Lester, J.; and Rickel, J. 1997. Integrating affective computing into animated tutoring agents. In *Notes of the IJCAI '97 Workshop on Animated Interface Agents: Making Them Intelligent*, 113-121.

Elliott, C.; Brzezinski, J.; Sheth, S.; and Salvatoriello, R. 1998. Story-morphing in the Affective Reasoning paradigm: Generating stories automatically for use with "emotionally intelligent" multimedia agents. In *Proceedings of the Second International Conference on Autonomous Agents*. New York: ACM Press.

Hayes-Roth, B.; van Gent, R.; and Huber, D. 1997. Acting in character. In R. Trappl and P. Petta, eds., *Creating Personalities for Synthetic Actors*. Berlin: Springer-Verlag.

Isbister, K. and Ishida, T. 1999. Designing for Social Interaction in Cyberspace. In *IPSJ (Information Processing Society of Japan) Magazine 40*(6), 569-574.

Johnstone, K. 1992. *IMPRO: Improvisation and the Theater*. New York: Routledge.

Lester, J.; Converse, S.; Kahler, S.; Barlow, T.; Stone, B.; and Bhogal, R. 1997. The persona effect: affective impact of animated pedagogical agents. In *Proceedings of CHI '97*. New York: ACM Press.

Lester, J., and Stone, B. 1997. Increasing believability in animated pedagogical agents. In *Proceedings of the First International Conference on Autonomous Agents*, 16-21. New York: ACM Press.

Oren, T., Salomon, G., Kreitman, K., and Don, A. 1990. Guides: Characterizing the interface. In B. Laurel (Ed.), *The Art of Human-Computer Interface Design*. Reading, MA: Addison-Wesley.

Pond, K. 1993. *The Professional Guide: Dynamics of Tour Guiding*. New York: Van Nostrand Reinhold.

Rickel, J., and Johnson, W. L. 1997. Integrating pedagogical capabilities in a virtual environment agent. In *Proceedings of the First International Conference on Autonomous Agents*, 30-38. New York: ACM Press.

Rist, T.; André, E.; and Müller, J. 1997. Adding animated presentation agents to the interface. In *Proceedings of the 1997 International Conference on Intelligent User Interfaces*, 21-28. New York: ACM Press.

Schank, R. and Abelson, R. 1995. Knowledge and memory: the real story. In R. Wyer, ed., *Advances in Social Cognition*, vol. VII. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 1-86.