Representing the Semantics of Virtual Spaces

Chaomei Chen, Linda Thomas, Janet Cole, and Chiladda Chennawasin
Brunel University

In the StarWalker virtual environment, users explore a shared semantic space and collaborate with concurrent visitors. StarWalker's design illustrates the potential of unifying spatial models, semantic structures, and social navigation metaphors in the development of multiuser virtual environments.

The structure of social interaction and the role of context have been traditional topics of research in linguistics, sociology, and anthropology. For example, linguists talk about *contextualization cues*, which, by indicating significant aspects of the social context, enable people to make inferences about one another's communicative intentions and goals. In sociology, *frame analysis* focuses on the perception and interpretation that participants give to their current social activity—what's going on, what the situation is, and what roles the participants adopt within it.

Our research addresses similar ideas within the realm of multiuser virtual environments. One of our primary concerns is the role that spatial models play in mediating social interaction and in establishing a mutually intelligible context, a common ground.

Over the past few years, spatial metaphors—ranging from text-based descriptions to 3D virtual-reality models—have become increasingly popular in the design of collaborative systems. For example, MASSIVE,[1] a virtual space for distributed multi-user interaction, uses a spatial model of interaction. MASSIVE supports participants' awareness of one another through spatial extensions of each user's presence (*aura*), attention (*focus*), and influence (*nimbus*). In design, these mechanisms employ real-world awareness and interaction models.

Advances in information visualization technologies and pioneering information visualization systems such as SemNet[2] demonstrate spatial models' great potential for a wider range of applications. In particular, the proliferation of information visualization models can play a significant role in extending and enriching the conventional design of inhabited, multiuser virtual environments.

However, spatial models alone may not suffice to foster and stimulate social activity in a virtual world. A natural, stimulating, and collaborative environment needs persistent resources from which users can derive a sense of context that lets them adapt and organize their courses of action and conversations.

To strengthen the contextual function of a spatial model, we introduce a unifying approach to the development of multiuser virtual environments. In particular, we exploit the potential of using a tightly coupled spatial-semantic model as a focal point in a multiuser virtual environment. We argue that this approach turns a spatial model into a useful resource that helps people understand and interpret social interaction in this context. Furthermore, we expect that virtual environments based on this approach would provide an intuitive starting point for people exploring a shared virtual space collaboratively.

In this article, we examine semantically organized virtual spaces, focusing on how such spatial organization shapes users' cognition, interpretation, and interaction. In particular, we describe the design of StarWalker, our own multiuser virtual environment. StarWalker illustrates our design principles for unifying spatial models, semantic structures, and social interaction patterns within multiuser environments.

Social interaction in MUDs

MUDs and MOOs (multiple-user domains and object-oriented MUDs) have a unique place in the history of computer-mediated communication and online communities. Traditionally, MUDs were designed for adventure games played by distributed users. Social use of MUDs subsequently developed, and this remains the predominant use of MOOs today. In parallel to the recreational and social uses, MOOs began to appeal to specialized user groups, especially researchers from related scientific disciplines and associated communities. In early-generation MOOs, users typically chatted through text-based dialogues; advanced systems allowed the use of video and audio channels.

More recently, a new generation of 3D, VR-empowered, chat-room environments have emerged on the Internet—notably Active Worlds (see Figure 1), Blaxxun Online Communities (see Figure 2), and Sony's Virtual Place. Although spa-

1070-986X/99/$10.00 © 1999 IEEE

tial models play an increasingly predominant role in the construction of virtual environments, researchers know relatively little about the role they play in mediating communications and social interaction.

A substantial body of research addresses the use of MOOs and, to a lesser extent, 3D virtual worlds. Table 1 lists some of the relevant studies; most of them examine the social behavior of participants in nonrecreational MOOs.

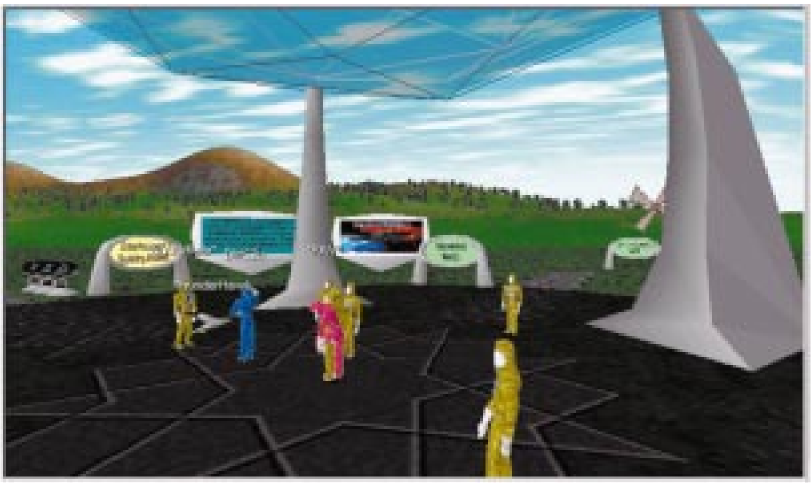

Figure 1. A virtual world in Active Worlds.

Figure 2. A virtual world in Blaxxun Online Communities.

Table 1. Studies of MOOs and shared virtual worlds.

Subjects	MOOs or Online Communities Examined
Nonrecreational use[3]	LambdaMOO (Astro-VR, Jupiter)
Social navigation[4]	Juggler
Turn-taking in virtual meetings[5]	Massive
Media spaces and MUDs[6]	Jupiter, Pueblo
Social positioning[7]	Active Worlds
Social activity in game MUDs[8]	Illusion
Patterns of social interaction[9]	LambdaMOO
Meetings in virtual spaces[10]	The Palace

LambdaMOO is a landmark in the history of MUDs. Much of the interest in social MUDs originated from the use of LambdaMOO or its direct descendants. A recently published survey[9] provides new insights into life in LambdaMOO. The survey shows that social interaction is the major activity in LambdaMOO, with exploring and building in the second and third positions. The survey further reveals that expert users tend to spend a considerable amount of time on social interaction and building their own chat rooms, whereas inexperienced users tend to explore the environment. Users regard public spaces primarily as places to meet people, especially initially. They tend to spend considerable time in private places, particularly their own virtual homes in LambdaMOO.

The survey suggests that the purely text-based LambdaMOO does provide a sense of place and space for many users. However, whether and how to promote spatiality in the design of online communities remains a complex and challenging issue.[9] A potentially fruitful perspective would exploit the interaction between sociality and spatiality in an online community.

The appropriateness of social behavior in a particular multiuser virtual environment depends on the interpretation of individual participants. We can regard such an interpretation as a type of social construction of knowledge. Harrison and Dourish[11] pointed out that it is such collective sense-making activity and mutual intelligibility that distinguishes a place from a space.

Research in nonrecreational virtual environments approaches the spirit of a relatively new type of collaborative environment, the collaboratory.[12] Collaboratories combine the interests of the scientific community at large with those of the computer science and engineering community to create integrated, tool-oriented computing and communication systems to support scientific collaboration. Collaboratory system design involves information navigation, distributed storage and computation, and advanced collaborative multimedia environments. Existing collaboratories include systems that allow telescope observations from a distance, remote participation in upper atmospheric research, and use of distributed tools for environmental molecular sciences.

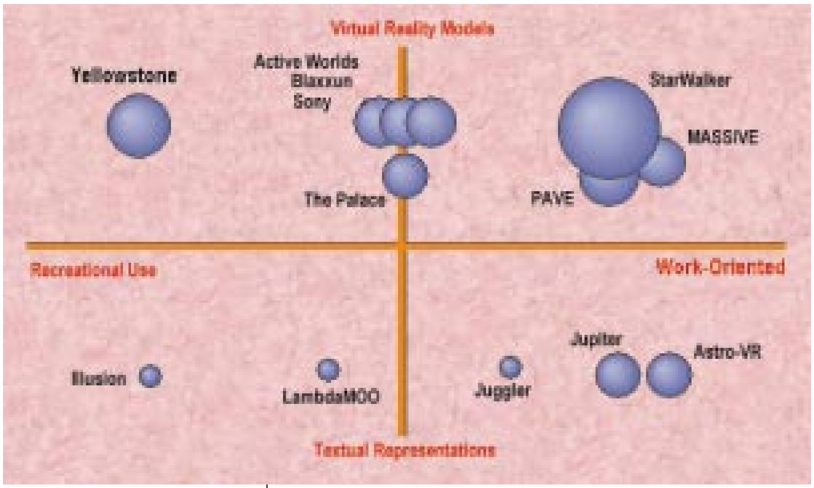

Figure 3. Virtual environments in context: application domains (horizontal position), visual representations of spatial models (vertical position), and strength of semantic cues (bubble size).

To put our work in a useful context, we've organized related work in online communities within a conceptual framework, shown in Figure 3. Three factors determine the position of a particular online community environment. The horizontal scale ranges from recreational to work-oriented use. The vertical scale indicates the extent to which the community provides visual representations of spatial metaphors. The third dimension, represented by bubble radius, indicates the strength of semantic cues that the virtual environment provides. (The visual representation in Figure 3 is illustrative of and created for this particular article. We would want more quantitative rating schemes for classifying a large number of computer systems and associated studies.)

Yellowstone in Active Worlds (see Figure 3) is a virtual world based on Yellowstone National Park. It uses 3D geographical models of the park to provide a relaxed social environment. In contrast, Jupiter and Astro-VR are selective in terms of audience and appropriate social behavior. These environments cater to people strongly associated with particular disciplines. We've placed StarWalker, our own environment, in the upper right quadrant because of its domain-specific focus, its VR-based spatial model, and the predominant visualization of the semantic structure.

StarWalker design

Now let's turn to the design of StarWalker, our 3D multiuser virtual environment. Unlike many chat-room virtual environments, StarWalker is a persistent virtual place: even if you encounter no other users to chat with, you can always explore the semantic structure on your own. The organizational principle of StarWalker is an enriched spatial model. Our intention is to provide a spatial model that can stimulate participants to question, interpret, and discuss the meaning of the underlying semantic space.

Building the semantic space

A crucial component of StarWalker is the network visualization of a semantic structure specific to a particular source collection—for example, a collection of articles in a subject domain or a database of information maintained by a company or an organization. We have developed an integrated set of tools, collectively called Generalized Similarity Analysis (GSA), for extracting a variety of intrinsic structures from the source collection and visualizing them as a special family of associative networks.[13]

The semantic structure of the virtual world in StarWalker's current prototype uses three years of proceedings from the ACM Computer Human Interaction conference (CHI, 1995–1997). You can access this environment on the Web at http://www.brunel.ac.uk/~cssrccc2/vrml2/starwalker/. (We are also testing StarWalker with a virtual world of the semantic space derived from the entire ACM Hypertext Conference series. See http://www.brunel.ac.uk/~cssrccc2/hypertextbib.)

We derived this spatial model automatically using a unique integration of latent semantic indexing (LSI)[14] and Pathfinder network-scaling techniques.[15] We used LSI to generate interdocument similarities, and we used the Pathfinder network-scaling technique to extract salient structural patterns.

LSI, designed to overcome the vocabulary mismatch problem of many information retrieval systems,[14] has attracted much attention in recent years in experimental information retrieval, psychology, and education. LSI allows us to approximate human judgments of overall semantic similarity, a task often necessary in psychological scaling and knowledge representation. LSI aims to generate similarities that can capture the underlying semantics by considering not only term co-occurrences but also wider contexts of such co-occurrences.

LSI essentially provides a dimension-reduction tool for the semantic space based on a mathematical matrix decomposition technique called singular value decomposition (SVD). In LSI, a semantic space builds on a large matrix of term-document association observations. You can approximate a very large term $\times$ document matrix with a truncated SVD matrix. Proper truncation can remove noise from the original data and also improve the

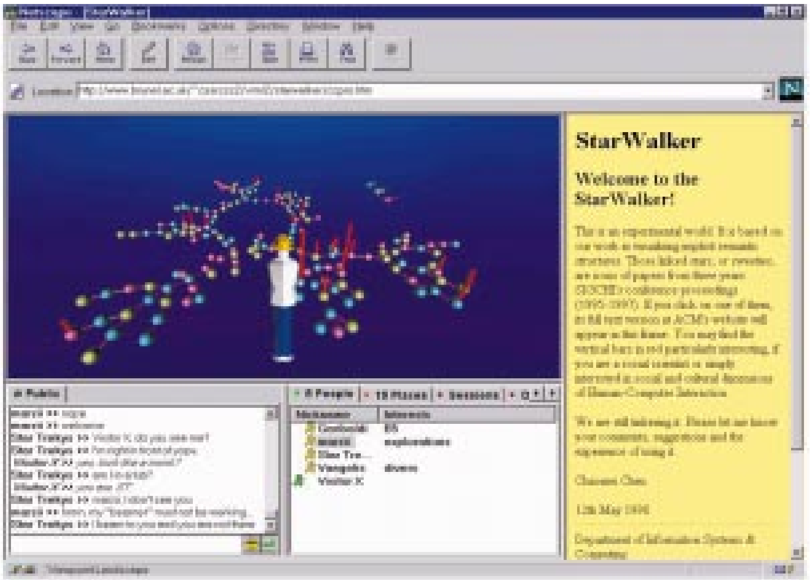

Figure 4. The user interface of StarWalker consists of a virtual-world window, a content window, a chat window, and a session window.

recall and precision of information retrieval.

Pathfinder network scaling, a structural and procedural modeling technique, extracts underlying patterns in proximity data and represents them spatially in a class of networks called Pathfinder Networks.[15] This technique gives us the ability to extract salient structural patterns, a feature we adapt in our approach. The scaling mechanism relies on a criterion called triangle inequality to eliminate counterintuitive links. Triangle inequality ensures that any two nodes in the resulting network connect only through the shortest paths. As a result, the network is a set union of all the possible minimal-spanning trees.[15] This desirable feature lets us extract a concise network representation without losing the network's original integrity. We use this property in our design to improve the clarity and quality of information visualization.

Once we obtain the semantic space and its Pathfinder Network representation, we render the semantic space into a VR model in VRML 2.0. This model can function within a multiuser virtual environment as well as on a normal Web page for individual access. We can also use different data sources to generate semantic spaces that define virtual environments for different user groups and online communities.

User interface

StarWalker currently employs Blaxxun's Online Community client-sever architecture, accessible through the Blaxxun Contact 4.0 client browser plug-in. Figure 4 shows a screenshot of the StarWalker client-user interface. The semantic model in the background represents the ACM CHI proceedings. The virtual world includes several viewpoints that let users explore the entire model from different perspectives. Users enter the virtual world via a landscape viewpoint (see Figure 4). The semantic model is rendered as a star constellation, with an overall metaphor of deep space, implying the infinite depth of the virtual world. Users can go anywhere they like.

In StarWalker, users most frequently ask the nature of the semantic model, the meaning of the links, and why they occupy a shared virtual environment, especially when visiting StarWalker for the first time. These questions often lead to more specific questions about the spatial model and, more importantly, about the underlying subject domain.

StarWalker's user interface consists of four windows: virtual world, content, chat, and session (see Figure 4). The environment lets users perform the following tasks:

- Explore the virtual world
- Change to predefined viewpoints in the virtual world
- Beam to an avatar in the virtual world
- Teleport to other virtual worlds
- Load and display the content of a document from the Web into the content window
- Talk in public or private through a text-based chat window
- Search and update the results in the virtual world (a capability currently restricted to internal users)

Users can talk with other concurrent visitors through a chat window (Figure 4, lower left), choosing private or public chat sessions. The user interface automatically saves transcripts of these dialogues; we drew the excerpts included in this article from these files. The session window displays status information including a list of virtual places available, the number of visitors in a particular virtual world, and short descriptions of visitors in the current virtual world.

The VR model constitutes the heart of StarWalker. Our emphasis on the contextual role of

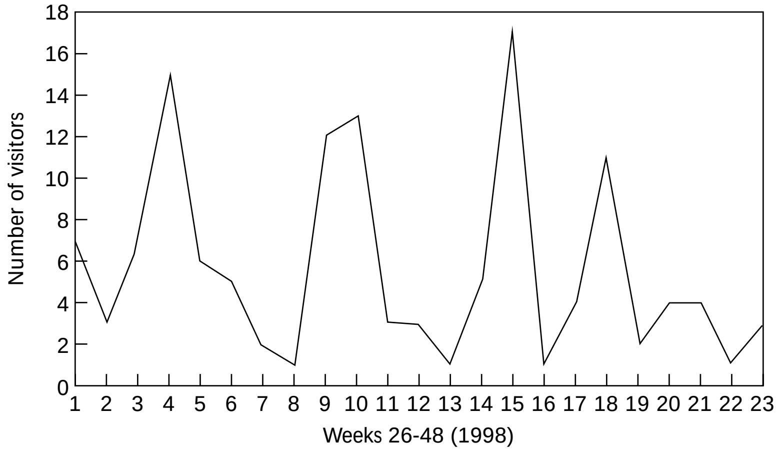

Figure 5. Number of visitors accessing StarWalker each week, June through December 1998.

this model distinguishes StarWalker from most other virtual environments. In Figure 4, each sphere in the spatial layout represents a paper from the online version of three ACM CHI conference proceedings (1995, 1996, and 1997). Links between spheres indicate that the corresponding papers have salient semantic connections derived through the modeling procedure described by Chaomei Chen.[13]

Users can move around within this VRML-based virtual world in several ways: typically, they walk, fly, or move directly to a predefined viewpoint. StarWalker currently has seven viewpoints from various perspectives: an overview, various landscape views in different directions, and a galaxy view. In a galaxy view, we could group several virtual worlds together for users to access. If several people visit the same virtual world, they can see the avatars of others moving through the virtual space. A user's spatial and semantic navigation patterns thus offer a meaningful and interpretable basis for social navigation. Not only can visitors make sense of events in a populated virtual environment, but they can also extract contextual and situational cues from others' behavior. This helps them adopt appropriate search and navigation strategies.

The content window displays the content of a paper downloaded when a user clicks on a sphere in the VR scene. StarWalker thus not only provides a forum for social interaction, it also functions as a gateway to literature resources on a particular subject domain.

StarWalker use

We have observed the use of the StarWalker virtual environment in several different settings, including invited private conversations as well as casual public chat. Figure 5 plots the number of requests per week for connecting to StarWalker during the second half of 1998. Four major peaks appeared during this period of time, each with more than 10 connections made to StarWalker. Six valleys occurred during this period with only one or two visits to StarWalker. More thorough data collection and analysis may reveal the predominant patterns of its use.

The following examples include excerpts from private chat sessions, public chat sessions, and a preliminary study of collaborative information retrieval in StarWalker.

Virtual meeting

We invited an expert on distributed virtual-environment systems to join the session from a remote site for a virtual visit. The meeting lasted about an hour and a half in private mode. Topics included the design of VRML models, multiuser virtual-environment platforms, and relevant research interests. The following excerpt focuses on multiuser virtual environments:

> A: DIVE is probably the best shared VR world but it has only recently been made available on WinNT
> B: How do I access DIVE
> A: via www.sics.se/dive
> B: what browser will I need?
> A: It has its own browser, scene and behavior language
> A: it's based on multi-cast rather than client-server
> A: it can use VRML 1.0
> B: is it relative stable?
> A: Yes it's at version 3.x now
> B: maybe I should check it out sometime.
> A: Yes it's worth it just to see a more advanced system than current Web-based shared worlds

StarWalker provides unique opportunities: Not only can we invite an expert into the virtual world to discuss many issues directly related to the design of the virtual world itself, but we can also experience the potential of connecting the trails of avatars in the virtual world to semantic interpretations.

In addition to discussing VR-related topics, we also took a photograph—a screenshot—of the expert's avatar in StarWalker. Figure 6 shows the snapshot taken in the virtual meeting—the expert was standing at the foot of a very tall red pole. Vertical red poles in the scene are landmarks that highlight the search results returned by the underlying LSI engine on articles about conversation analysis and related topics.

Our virtual meeting with the expert in distributed virtual environments required only the client browser plug-in in addition to a standard-configuration, networked PC. The plug-in is downloadable free of charge from Blaxxun Interactive's Web site (http://www.blaxxun.com). Thus, at its simplest, our environment provides a cost-effective communication channel for researchers to contact experts and elicit user feedback.

Social navigation

In art museums, visitors often gather in front of paintings, photographs, or other exhibits. They often use this clustering behavior of like-minded people as a reference for judging the value of a place or a point of interest. The virtual counterpart of this phenomenon—using similar heuristics to judge electronic information—is called social navigation.[16] Dieberger[4] discussed how a textual MOO called Juggler uses read-wear techniques to support social navigation. Read-wear provides a collaborative recommendation mechanism in a computer system by showing the current user how frequently other users have read a particular article in the past.

We designed StarWalker primarily to foster social interaction in a semantically rich context. It conveys a strong and novel metaphor to users, and the interaction model transcends the conventional models of buildings, cities, and geographic landscapes of other virtual worlds. Instead, StarWalker offers a world of meaning, knowledge, and cognition. Some visitors clearly recognize our design rationale. One said,

> I have not seen any multiuser environment this abstract before... I always thought the communities look too much like the "real" world ...

Another said,

> I saw it on the list and never been here. Star Walker sounded interesting, but instead of going to outer space I came into inner space ;-)

The strong semantic-oriented metaphor of StarWalker has also influenced how people perceive this virtual place and how they interact with others. One visitor asked if this place were only for trekkies (*Star Trek* fans) when he met another visitor with the nickname Star Trekys.

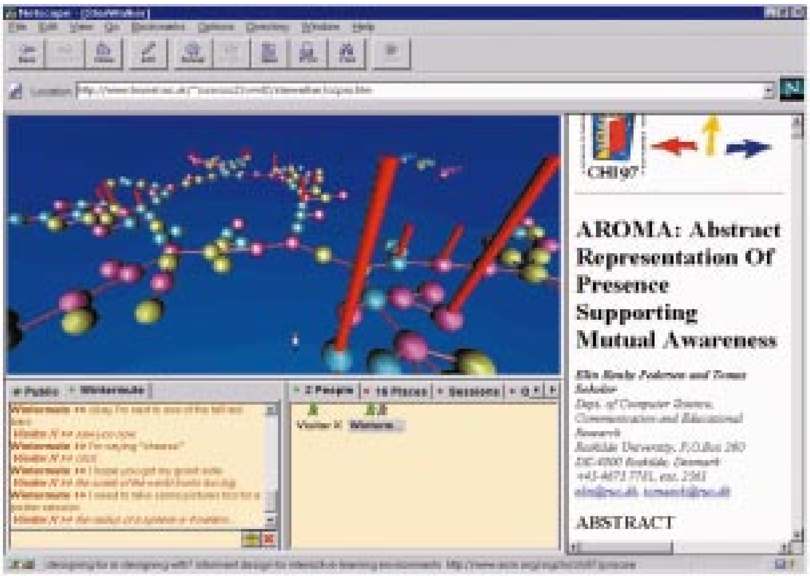

Figure 6. A snapshot taken during a virtual meeting in StarWalker.

Because the spatial model and the semantic structure are unified in StarWalker, the positioning, movement, and groupings of avatars in the virtual world offer meaningful indicators to peer visitors. For example, it becomes possible for users to infer the relevance of articles in particular areas based on the access patterns of colleagues or other users with known expertise.

Figure 7 shows a screenshot of StarWalker in real use. In this scene, four avatars have gathered into the areas dominated by a number of red vertical poles. This emergent pattern demonstrates the potential of social navigation—such trails and gatherings naturally foster social interaction.

The virtual world's spatial structure appears to have evoked several episodes that led to questions specifically concerning its semantic structure as well as its geometrical or topological properties. In further research, we are interested in the extent to which StarWalker's strong semantic cues stimu-

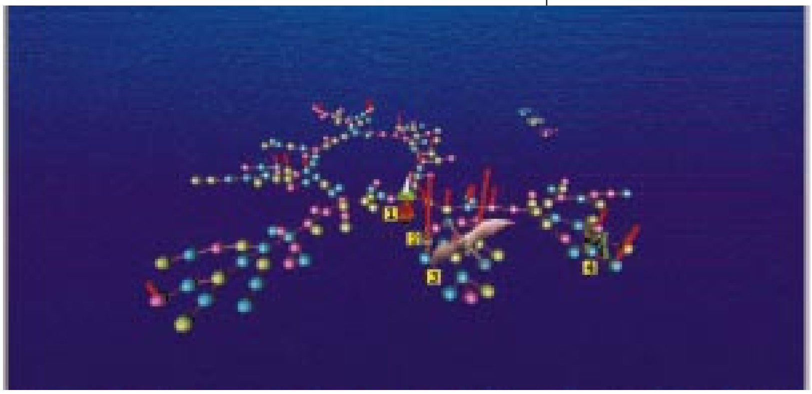

Figure 7. The trails of avatars in StarWalker provide a good starting point for social navigation.

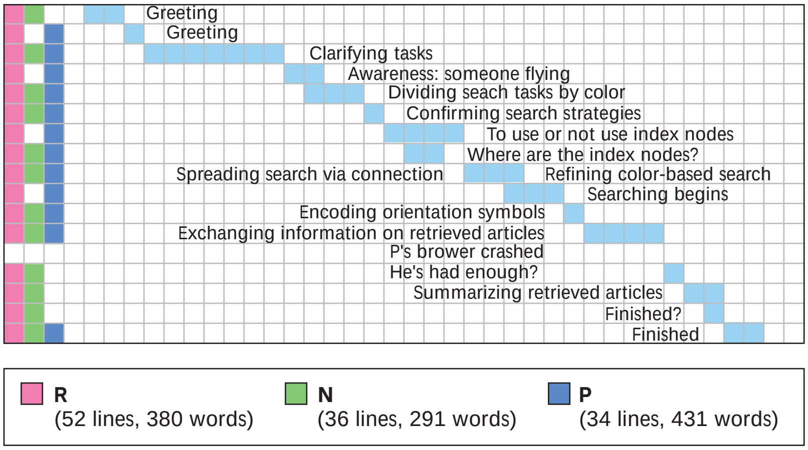

Figure 8. The flow of topics in collaborative information retrieval in StarWalker.

late social interaction. We would like to explore the interplay between the spatial structure of a virtual world and the discourse structure within the virtual environment. Integrating the trails of avatars with the concept of read-wear in the semantic space seems useful for future research and development of StarWalker. To investigate how such persistent environments should reorganize themselves in response to the dynamics of social navigation, we may need longitudinal studies on the emergence of a new community.

Collaborative information retrieval

To better understand how visitors would use StarWalker for collaboration, we invited three colleagues to use the StarWalker world based on the CHI conference proceedings for a collaborative information retrieval exercise. They accessed the StarWalker virtual environment from three computers running Microsoft Windows NT, Netscape Communicator 4.05, and the latest version of the Blaxxun Interactive Online Community client browser plug-in, Contact 4.0.

Prior to the experiment, we gave the participants a brief introduction to StarWalker. We then asked them to imagine that they were going to write a paper on a human-computer interaction topic for a conference or a journal, and that they would collaboratively search for some publications on potential topics using StarWalker.

The experiment lasted about an hour. A preliminary analysis of the chat log shows that one participant took 52 turns of talk in the whole session, one 36 turns, and one 34 turns. The participant with the least number of turns had the highest word count, 431. The other two word counts were 380 and 291. The participant with the highest number of turns in the session often acted as a moderator.

In conventional information retrieval, a search typically consists of several stages—namely, query formation, verifying search results, query refinement, storing new search results, and so on. In collaborative information retrieval, we expected a similarly structured process. We also expected the unique organization of the virtual world to influence the participants' search strategies.

Figure 8 shows a topic flow diagram of this collaborative information retrieval process. The color-coded boxes to the left of the diagram indicate which participants were involved in a particular communicative segment, which we call an episode. A sequence of blue boxes denotes the length of each episode. Each blue box corresponds to up to five turns of talk in the chat window. Next, we'll discuss a few episodes particularly relevant to the design of StarWalker.

Participants spent by far the longest time clarifying the task and deciding what topics to search for. They first discussed topics. They appeared to have agreed on searching for articles concerning avatars in virtual environments. The following excerpt shows that the three participants realized that they might have a different understanding of what they should search for.

R: paper on wayfinding strategies?...
N: yeah suppose so
P: okay
R: ok, so what, we're each going to find papers relating to VR and avatars and then discuss?
P: hey what happened to wayfinding or am I getting the wrong stick?
R: hey?
N: I thought they were the same thing ?
P: things that help you and how you navigate, that's what I thought

In the VR model, articles are color-coded by the year of publication. Participants in our experiment initially chose to divide the search space by these colors.

N: shall we each take a year then?
R: ok..green blue or pink?
N: I'll have blue
R: ok, me green...hehehe

In the following episode, participants noticed that the initial search-by-color strategy overlooked the connectivity in the semantic space.

P: ...... I'll try the right hand branches of the space
R: right hand?...what about years or did we abolish that?
N: well I'm searching all the blue balls
P: hmm what should we do? systematically check every (color of choice) blob, or split up and search
R: if we do it by color, then we miss out the connections?
R: just noticed that
N: ok start from the color and then if there are any connections search them too
R: ok

Although participants realized that searching by color would miss the connections, they decided to follow their initial search strategy anyway. In fact, a modified search-by-color strategy was proposed and accepted. One of the participants did mention a potentially more effective search strategy, "split up and search." This strategy would have taken advantage of the topology of the semantic structure because each branch represents a cluster of papers on similar topics. This suggests that how users select their search strategies demonstrates the level of their understanding of the spatial-semantic structure.

We also noticed that none of the participants in the collaborative information retrieval session asked the typical question, "What is this place?" This suggests that in-depth training sessions and background readings on the design rationale might help users develop a better understanding of the implications of the new metaphor on their course of action.

Discussion

We designed the StarWalker virtual environment to foster social interaction and stimulate focused conversations and to let users explore a semantic space collaboratively. A fundamental question is whether and how users might interact and collaborate with each other in such virtual environments.

Our preliminary analysis of the collaborative information retrieval dialogues suggests that StarWalker provides a meaningful context for peer learning through an engaging group task. For example, on several occasions participants found that they had to clarify the meaning of seemingly straightforward search topics, such as the use of avatars in VR and wayfinding strategies. The stand-by virtual world of the underlying semantic space appeared to contribute to the brainstorming process—users evidently explored some articles in the virtual world in parallel as they discussed search topics.

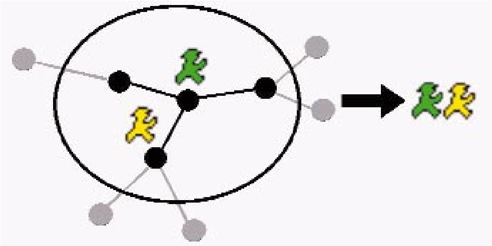

Figure 9. Maintaining awareness through the shortest paths in the semantic space.

The search-by-color strategy turned out not to be particularly compatible with the structure of the virtual world. Participants subsequently discussed the alternative of dividing the space into smaller areas. Focusing on the search strategy led the participants to realize that the search-by-color strategy overlooked connectivity patterns. However, they did not abandon the initial search-by-color strategy and work out a connectivity-based search strategy, which would be ideal for the given task.

These issues suggest that the conventional chat-room facility does not adequately support a highly task-oriented discussion. Brainstorming and decision-making support mechanisms, such as gIBIS (graphical issue-based information systems), would improve the effectiveness of brainstorming. Incorporating such facilities into StarWalker would be a natural step towards supporting collaborative work as opposed to traditional chat-room social interaction.

Another design issue is how we should support informal communication along with a smooth transition from loosely coupled to tightly coupled collaboration. Currently, although the virtual world's spatial configuration reflects its semantic structure, an avatar's position doesn't necessarily indicate the user's attention. To retrieve an article, the user does not have to place his or her avatar close to the sphere—a click on the ball loads the content into the content window.

We could use the semantic distance between articles activated by two users to provide a new framework for awareness support. Because the shortest semantic distance between any two articles is known in StarWalker, each user can define a semantic aura so that he or she will be notified when someone moves into the scope of the semantic aura. When two users are close enough to each other according to the semantic distance, they will maintain a tighter mode of mutual awareness (see Figure 9). This is a generalization of awareness maintained on the basis of geometric proximity.

Studies on classic MOOs such as LambdaMOO suggest that experienced users typically spend a considerable amount of time building their homes, and a large amount of engaging social interaction takes place at users' homes rather than in a public place. This raises a further design question for us: What would happen if users could create their own objects in the StarWalker virtual semantic world?

Conclusions

The design and use of the StarWalker virtual environment illustrate our novel approach to knowledge-based media space modeling. We focused here on how a semantically enriched spatial model mediates interactive behavior of users in a domain-specific context.

We found our experience with StarWalker encouraging—many visitors to the virtual environment are enthusiastic about the design principle, and some engaging conversations have taken place. The enriched spatial model may play a significant role in attracting people to visit, to interact, and to adapt the environment into their distributed working environments.

We recognize the significance of giving users the ability to adapt the virtual space according to their own interests. Users should be able to personalize a subspace and restructure it to meet their own needs. The persistent nature of StarWalker lets people access such resources as an integral part of their distributed collaborative environment. In this way, StarWalker merges and integrates a media space with a collaborative virtual environment, where users can explore and share their interpretation of knowledge structures effectively.

Of course, a lot of work remains. Pioneering media spaces and collaborative virtual environments have resulted in invaluable experience and lessons that we could draw upon. Information visualization, which has played a crucial part in our approach, is evolving rapidly, emerging as a new discipline. We need to tackle many challenging new issues as we start to deal with social, ecological dynamics as well as cognitive and individual differences in a wider context. MM

Acknowledgments

This work was in part funded by the Engineering and Physical Sciences Research Council (EPSRC) in the UK under the Multimedia and Networking Application (MNA) program (Grant Number GR/L61088). We thank the participants in the StarWalker sessions.

References

1. C. Greenhalgh and S. Benford, "Massive: A Virtual Reality System for Teleconferencing," *ACM Trans. Computer Human Interaction*, Vol. 2, No. 3, 1995, pp. 239-261.
2. K. Fairchild, S. Poltrock, and G. Furnas, "SemNet: Three-Dimensional Graphic Representations of Large Knowledge Bases," *Cognitive Science and Its Applications for Human-Computer Interaction*, R. Guidon, ed., Lawrence Erlbaum Associates, Hillsdale, N.J., 1988, pp. 201-233.
3. P. Curtis and D.A. Nichols, "MUDs Grow Up: Social Virtual Reality in the Real World," tech. report, Xerox PARC, Palo Alto, Calif., 1993; ftp://parcftp.xerox.com/pub/MOO/papers/MUDsGrowUp.ps.
4. A. Dieberger, "Supporting Social Navigation on the World-Wide Web," *Int'l J. Human-Computer Studies*, Vol. 46, No. 6, June 1997, pp. 805-825.
5. J. Bowers, J. Pycock, and J. O'Brien, "Talk and Embodiment in Collaborative Virtual Environments," *Proc. Computer-Human Interaction Conf,* ACM Press, New York, 1997, pp. 58-65.
6. E.D. Mynatt et al., "Design for Network Communities," *Proc. CHI*, ACM Press, New York, 1997, pp. 210-217.
7. P. Jeffrey and G. Mark, "Constructing Social Spaces in Virtual Environments: A Study of Navigation and Interaction," *Proc. Workshop on Personalised and Social Navigation in Information Space*, 1998, pp. 24-38; http://www.sics.se/humle/projects/persona/web/workshop/proceedings.html.
8. J. Muramastsu and M.S. Ackerman, "Computing, Social Activity, and Entertainment: A Field Study of a Game MUD," *CSCW Journal*, Vol. 7-1/2, 1993, pp. 87-122.
9. D.J. Schiano and S. White, "The First Noble Truth of CyberSpace: People Are People (Even When They MOO)," *Proc. CHI Conf.*, ACM Press, New York, 1998, pp. 352-359.

10. L. Toomey, L. Adams, and E, Churchill, "Meetings in a Virtual Space: Creating a Digitial Document," *Proc. 31st Ann. Hawaii Int'l Conf. System Sciences*, IEEE CS Press, Los Alamitos, Calif., 1998, pp. 236-244.
11. S. Harrison and P. Dourish, "Re-place-ing Space: The Roles of Place and Space in Collaborative Systems," *Proc. Computer-Supported Cooperative Work Conf.*, ACM Press, New York, 1996, pp. 67-76.
12. S. Bly, "Special Section on Collaboratories," *Interactions*, Vol. 5, No. 3, May/June 1998, p. 31.
13. C. Chen, "Generalized Similarity Analysis and Pathfinder Network Scaling," *Interacting with Computers*, Vol. 10, No. 2, Apr. 1998, pp. 107-128.
14. S. Deerwester et al., "Indexing by Latent Semantic Analysis," *J. Am. Soc. Information Science*, Vol. 41, No. 6, Sept. 1990, pp. 391-407.
15. R.W. Schvaneveldt, F.T. Durso, and D.W. Dearholt, "Network Structures in Proximity Data," *The Psychology of Learning and Motivation*, G. Bower, ed., Vol. 24, Academic Press, San Diego, Calif., 1989, pp. 249-284.
16. P. Dourish and M. Chalmers, "Running Out of Space: Models of Information Navigation," *Ancillary Proc. HCI,* Cambridge Univ. Press, Cambridge, UK, 1994; ftp://parcftp.xerox.com/pub/europarc/jpd/hci94-navigation.ps.

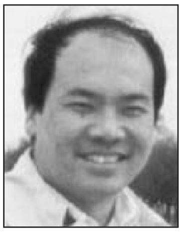

Chaomei Chen is a lecturer in information systems at Brunel University. His research interests include information visualization, hypermedia, digital libraries, and human-computer interaction. He is the author of *Information Visualisation and Virtual Environments* (Springer-Verlag, London, 1999) and has published over 40 articles in refereed journals and for international conferences. Chen received his BSc in mathematics from Nankai University, China, his MSc in computation from the University of Oxford, and his PhD in computer science from the University of Liverpool. He is a member of the ACM and the British Computer Society HCI specialist group.

Linda Thomas is a professor of education and a pro vice-chancellor at Brunel University. Her specialty is the application of phenomenographic methodology to investigations of the teaching and learning of economics. She is also interested in the application of new technologies to learning and teaching in higher education. Thomas received her BSc in economics and MEd in education from the University of Wales and her PhD from the University of London.

Janet Cole is a research assistant at Brunel University. Her research interests include group dynamics and discourse structures in a collaborative virtual environment. Cole received her BSc in computer science from Brunel University. She is a member of the BCS and the ACM.

Chiladda Chennawasin is a PhD candidate at Brunel University. Her research interests include information retrieval and information visualization in virtual environments. She received her BS in computer science from the University of Vermont and her MSc in information systems from Brunel University. She is a member of the IEEE Computer Society.

Readers may contact Chen at the Department of Information Systems and Computing, Brunel University, Uxbridge UB8 3PH, United Kingdom; chaomei.chen@brunel.ac.uk.