



Social diffusion patterns in three-dimensional virtual worlds

Katy Börner¹
Shashikant Penumarthy²

¹Indiana University, SLIS, Bloomington, IN, U.S.A.

²Indiana University, Computer Science Department, Bloomington, IN, U.S.A.

Correspondence:

Katy Börner, Indiana University, SLIS, 10th Street & Jordan Avenue, Main Library 019, Bloomington, IN 47405, U.S.A.
Tel: +1- 812 855 3256; Fax: +1 812 855 6166
E-mail: katy@indiana.edu WWW: <http://ella.slis.indiana.edu/~katy>

Abstract

This paper presents a visualization tool set that can be used to visualize the evolution of three-dimensional (3D) virtual environments, the distribution of their virtual inhabitants over time and space, the formation and diffusion of groups, the influence of group leaders, and the environmental and social influences on chat and diffusion patterns for small (1–100 participants) but also rather large user groups (more than 100 participants). The techniques are applied to analyze and visualize data recorded during events in virtual worlds, as well as simulated data, but are also applicable to real-world data. Resulting visualizations can and have been used to ease social navigation in 3D virtual worlds, help evaluate and optimize the design of virtual worlds, and provide a means to study the communities evolving in virtual worlds. The visualizations are particularly valuable for analyzing events that are spread out in time and/or space or events that involve a very large number of participants. The paper reviews and builds upon research in information visualization, scientific visualization, geography, architecture, and social science. It discusses intended user groups and their tasks and how the proposed techniques support those tasks. Three dimensional virtual world technologies are briefly described before the visualization tool set is explained in detail together with sample applications. The paper concludes with a discussion of results and an outlook for future work.

Information Visualization (2003) 2, 182–198. doi:10.1057/palgrave.ivs.9500050

Keywords: information visualization; social visualization; social diffusion; group behaviour; spatial navigation; data mining; virtual worlds; collaborative environments; geographic information systems

Introduction

How do cities evolve over time? How are they used by different social groups? How do people, inventions, and diseases spread over time and space? These are questions that are frequently dealt with by a myriad of professionals, including cartographers, architects, and social scientists.

Today, information about the design/deletion/modification of most manmade objects (e.g., buildings, places, etc.) is available in the form of time-stamped cartographic maps, digital architectural blueprints, floor plans, or product records. Census figures, scientific publication data, etc., provide a wealth of spatially explicit information about us and our creations. Various data sets, such as email, news, telephone traffic, mail, etc., can be analyzed to identify social groups based on the frequency and kinds of interactions among participants. Sharing and correlating samples, data, and techniques is beneficial, for example, to identify and combat the spread of diseases such as SARS.

However, the correlation and analysis of large-scale (spatially and temporal explicit) data sets and the interpretation of results is highly complex. We claim that information visualization techniques can be

Received: 6 August 2003
Revised: 3 September 2003
Accepted: 12 September 2003



applied to ease the understanding and communication of analysis results. In particular, we are going to introduce techniques to visualize the evolution of three-dimensional (3D) virtual environments, the distribution of their virtual inhabitants over time and space, the free and guided diffusion of social groups, and the environmental and social influences on chat and diffusion patterns for small (1–100 participants) but also rather large user groups (more than 100 participants). Object creation/deletion/modification data as well as user position/interaction/chat statistics obtained from virtual world events as well as simulated data sets will be used to demonstrate the proposed techniques.

The use of virtual world data is of increasing interest to researchers in accordance with the continuing rise of virtual online spaces for education, leisure, and general social interaction. A number of virtual worlds and events have reached dimensions that are very much comparable with real world counterparts in terms of size or number of participants. Virtual tours, learning environments, conferences, art shows, etc., can easily attract several hundred participants that quickly distribute in time and space.¹ To give an example, the Avatars! Conference in 1998 had over 4000 attendees and has continued to grow over the last 5 years.² Similar to participants in real-world events, virtual participants typically have a limited view in terms of space and time. That is, they can see only nearby participants and be at one point in time.

Today, 3D virtual environments and events are far less effective than their real-world counterparts.³ Part of the problem is that existing 3D browser systems (including AW) provide very limited support for spatial and social navigation; overview maps are rare. Saving a chat log (oftentimes restricted to a fixed number of nearby avatars) and taking screenshots are the most common methods of logging usage information. Hosting an online demo for a large number of participants or finding something or somebody in a large world can easily become a nightmare. Offline and online analysis of events with many users or events that are spread out in space and time is also difficult. Frequently, spatial and temporal patterns and their interactions have to be analyzed and understood to arrive at valid conclusions. The features and functionalities of online spaces that do not have analogs in physical spaces, for example, larger hearing range, chat instead of voice, etc. make it necessary to develop new means for describing and evaluating online spaces.

This paper summarizes and extends work on visualization tools that can and have been used to ease social navigation in 3D virtual worlds, to help evaluate and optimize the design of virtual worlds, and to provide a means to study the communities evolving in virtual worlds.^{4–6} Note that all subsequently presented visualizations are designed to be highly readable and understandable. They are simplified to a point that they appear to be obvious solutions for, example; visualizing the growth of a digital learning environment. This ease is

intended, as the major user groups (teachers, kids, etc.) are not able or do not have the time to study instructions or go through extensive teach-in phases to understand highly complex visualizations.

The remainder of the paper is organized as follows: The next section introduces the main user groups and their tasks that are served by the analysis methods and visualizations presented in this paper. Later existing work in information visualization, scientific visualization, geography, architecture, and social science are discussed. The section after that describes the virtual world technology we are using as well as the data sets used to exemplify the proposed techniques. Further, the main part of this paper, describes and exemplifies a visualization tool set that can be used to analyze the growth of virtual worlds, world activity, chatting patterns, and the spatial-temporal distribution and diffusion of users are discussed. The paper concludes with a discussion of challenges and opportunities for the presented work.

User groups and their information needs

The visualizations presented in this paper (see the section on three dimensional virtual worlds and their inhabitants) have three main user groups: Users in need of social navigation support, designers concerned with evaluating and optimizing their worlds, and researchers studying virtual worlds. The information needs of these three groups, as discussed in detail below, were extracted from two sources: interviews with a large number of users and owners of virtual worlds conducted over the last 3 years; and results of an online survey that investigated design principles for educational virtual worlds.

Users of virtual worlds – like users of other digital resources (e.g., web pages) – are confronted with a space that typically shows no common usage signs, such as footsteps, dog ears on books, bite marks on toys, etc. Given that the virtual worlds are collaborative spaces allowing multiple users to interact (see next section), users may be drawn toward crowds of people and benefit from querying others. However, seeing closeby people are typically the only clues that aid social navigation.

Almost all of the virtual worlds that we have studied, analyzed and visualized thus far serve educational purposes. Hence, our users represent children ages 9–13 and their teachers and mentors. Educational objectives might comprise improving learner communication strategies, improving English language and 3D design skill, developing a sense of student community, and understanding complex processes.

Within 3D worlds, children are interested in finding and interacting with their peers, teachers, and mentors, and therefore need to navigate effectively. Teachers who use 3D virtual environments for a variety of teaching methodologies, including inquiry based learning, learning tied to interaction and sharing, learning by explaining, constructivist learning,⁷ etc., need to glean the effectiveness of the environment and their mentoring. Therefore, teachers are interested to find out if, and how,

children completed a certain learning experience, as well as the details of their interactions with each other and the environment. During an online lecture, for example, they may need to know if a student got lost or otherwise needs help.

Designers of virtual worlds may include artists, computer scientists or social science researchers who have strong technical and 3D modeling skills. They aim to build worlds that efficiently serve the information and interaction needs of a certain community. In order to evaluate and optimize designs they need information on which places, teleports, or web links are accessed, how often, and by which users. Based on this information they might decide to make rarely accessed resources more accessible (or to delete them) and to accommodate high access for frequently used resources. Information on where people engage in which activity (e.g., chatting⁸ or solitary work) helps to design environments that support the types of planned activities.

Researchers, such as information scientists, social scientists, or educators, might be interested in the commonalities and differences between real world and virtual communities.⁹ They may wish to know how social groups form, interact, and vanish, or what educational purposes can be well served by using virtual worlds (as opposed to classroom activities). Answers to those questions require not only close collaboration between the owners and users of virtual worlds, but also tools with which to analyze world growth or user interaction patterns for specific events and over time.

Section on three-dimensional virtual worlds and their inhabitants will present analysis techniques and visualizations that aim to serve all three user groups. Prior to this, we will discuss related research and provide details on the 3D virtual environment we are using.

Related work

At least five rather distinct research areas are concerned with the development of theories, techniques, and systems that support the study of (virtual) spaces, their inhabitants and social navigation patterns, as well as the design, evaluation, and optimization of these spaces.

The most closely related research area is social visualization. Here, research-to-date has concentrated on the visualization of rather small user groups. However, the visualization of 10 users is qualitatively quite different from the visualization of, say, 1000 users. The latter requires that data are aggregated in a form that shows major patterns and trends. A second, related research area is scientific visualization, exploration of which is concerned with the depiction of massive amounts of spatially referenced data. It has developed diverse techniques that can be advantageously applied to visualize, for example, the diffusion of user crowds over time and space. Most scientific visualization approaches batch process rather large data sets as opposed to life data streams. In addition, they are typically not very interactive. Last but not least, work in other scientific fields,

for example, geography, architecture, and social sciences, can be used to analyze movement and interaction patterns of entities. Approaches and techniques from all five areas are discussed subsequently.

Social visualizations are a special type of information visualizations that focus on analysis of social behavior. For example, lifeline visualizations reveal migrations, transitions and trajectories of users or user groups.^{10,11} Other research aims at the visualization of very large-scale conversations, such as those that take place on the Usenet,¹²⁻¹⁵ or visualize Web activity or user trails.¹⁶⁻¹⁸ The representation of people in text-based or graphical VVs by avatars is yet another topic. Avatars are varied and range from purely textual descriptions over 2D smiley faces, cartoon characters or photographic images, to abstract or highly realistic 3D models.¹

Further, there is interest in visualizing and supporting social interactions in text-based or 2D graphical systems.^{19,20} *Chat Circles*¹³ is a 2D graphical interface for synchronous conversation. It visualizes the non-textual components of online chatting, such as pauses and turn-taking behavior, that can be key to fully understanding the nature of discussion²¹ and that are lost in regular chat log files. The conversational archive of *Chat Circles* can be visualized as a 2D, interactive conversational landscape in which each vertical line shows the activity of one participant and the horizontal lines are postings. *People-Garden*²² uses a particularly apt flower metaphor to create individual data portraits of chat participants and a garden metaphor for combining these portraits to present the conversation activity of a group of participants. Work by Marc Smith *et al*²³ analyzed gestures and movement of users in VChat, a graphical chat system. They compared the average distance and orientation of users in relation to users targeted in their chat and randomly selected users. They concluded that people were standing closer to their chat target, but kept some distance from targeted users to maintain personal territories.

Another line of research focuses on mapping MUDs and 3D virtual worlds. Martin Dodge's *Atlas of Cyberspaces* section on MUDs & Virtual Worlds provides an excellent overview.²⁴ Elaborated maps by Andrew Smith show the urban density and the teleport systems of his 3-D world.²⁵ The *AlphaWorld Mapper* (<http://mapper.activeworlds.com/aw/>) by Greg Roelofs and Pieter van der Meulen provides access to a complete, zoomable 2D map of a virtual world that is roughly the size of California (429,025 km²).

Maps help users to orient themselves in an environment,²⁶ they also equip their users with survey knowledge that may be hard to acquire purely by navigation of the environment.^{27,28} Maps have been used to support users' navigation in virtual environments (e.g. by allowing them to drag an icon of themselves to a desired new position)²⁹ as well as in graphical multi-user domains (MUD) consisting of spaces and landmarks.³⁰

The existing systems teach many valuable lessons about how to visualize online spaces, their users, navigation and conversation activity, as well as (social)

relationships. However, the analysis of computer-mediated communication is dominated by textual analysis without spatial reference, using linguistic discourse analysis and ethnographic methods.³¹ Studies that examine graphical chat rooms (e.g., Suler's Palace study) use field observations, email interviews, studies of mailing lists, and participant observation. Naper³² was among the first to analyze chat text logged in a 3D virtual worlds and calls for analysis of graphical elements. Most systems also analyze a static set of data. Notable exceptions are *Chat Circles*,¹³ Erickson *et al's* *Babble* system,²⁰ *Footprints*,¹⁷ and *Anemone*.³³

We are not aware of any analysis or visualizations of spatio-temporal diffusion patterns of single users or social groups. However, user positions and temporal aspects of user interactions are key variables in the identification, analysis and understanding of user–environment and user–user interaction data. Very few social visualizations approach scale. However, the visualization of more than, say, 100 users requires qualitatively different visualization techniques from the visualization of smaller groups (e.g. 1–20 users).

Scientific visualization research is typically concerned with the visualization of rather large data sets such as those collected in fluid flows or wind tunnel studies. Frequently, vector field techniques are applied to analyze and visualize pressure, heat, or contamination distributions. Vector field techniques aim to visually convey large amounts of 3D directional information. They can also be applied to represent complex diffusion patterns, e.g., of objects or users in virtual worlds.

Two properties of vector fields are of special interest: (1) the strength and direction of flow at a given location and (2) information on where the flow comes from and where it will go (interesting, e.g., for spread of pollution or diseases). 'Local' flow can be visualized with respect to a fixed point using glyphs that show the direction and magnitude. 'Global' flow is shown by visualizing the trajectory of a (mass less) particle transported by the flow via particle traces, streamlines, stream ribbons, or stream tubes. Some systems allow user to inject 'dyes' of various colors into the flow field to probe its shape and behavior.

Diverse methods exist to compute the magnitude and direction of the vector field at each point. ([http://](http://www.cg.tuwien.ac.at/research/vis/dynsys/frolic/)

www.cg.tuwien.ac.at/research/vis/dynsys/frolic/). Isoclines are curves within the vector field, where either the x or the y derivative is zero. They can be seen as boundary demarcations between different flow areas.

Given the large size of most data sets, only simple interactions with the visualization (such as zooming, filtering, etc.) are typical. To our knowledge there exist no techniques that could be applied to track the evolution of global flow features over time.

Geographers, in particular cartographers, have developed a number of techniques to visualize spatio-temporal diffusion patterns. Pioneering work by Tobler represents and visualizes diffusion potentials and gradients as vector fields and as continuous spatial gravity models.^{34,35} To compute the diffusion of features over time, a movement table is compiled. It resembles a square matrix indicating movement from every point to every other point. Then, using a continuous version of the gravity model, a map can be computed, the result being a set of partial differential equations solved by a finite difference iteration to obtain a potential field³⁶ (see Figure 1, left). The spatial pressure field is then used to obtain information on the direction and the intensity of the flow at a certain point in geographic space and time (see Figure 1, right). Figure 1 shows the pressure to move in US based on a continuous spatial gravity model. Clearly, New York exhibits the highest diffusion pressure while Florida has the highest (inward) absorption, that is it acts as a sink.

Conventional cartographic flow maps have also been used to represent diffusion patterns. Handcrafted flow maps by TeleGeography (<http://www.telegeography.com/>) represent the volume of international telephone traffic flow between European nations via arrows overlaid over cartographic maps where the thickness of the arrow corresponds to volume of the traffic.

Architects know that the layout of a building/place has fundamental effects on how and where people move, sit, browse, buy, or get lost.³⁷ They study the performance of a building/place by studying the way people move though it—either by literally following them or by analyzing time lapsed crowd movement recordings. For example, Whyte³⁸ determined the influence of steps, fountains, green spaces, sitting places, building arrangement, etc. on the crowd flow and social interaction in




Figure 1 Migration gradients and potentials generated by Tobler. Reprinted with permission.

New York plazas.³⁹ Space index techniques and ‘spatial integration’ techniques have been applied to forecast the accessibility of individual pedestrian route segments based on their position in the overall movement network.³⁷ Fascinating work by *Space Syntax Limited* (<http://www.spacesyntax.com/>) aims to quantify the degree to which different factors influence pedestrian flow³⁹ and how function follows form in general.⁴⁰ Results augment strategic design as well as the selection of design alternatives that best serve the needs of a certain population.

Due to the limited spatial coverage (data is collected for specific places, buildings, transport systems, etc.), the quality of the data (manual counting of people), and the kinds of data available (movement data, but typically no social interaction data, etc.), the resolution of the spatial analysis of user data as well as the correlation of different data types, for example, movement and chatting, is limited.

Social science research has developed a diverse set of social network analysis methods.⁴¹ Social networks represented by matrices or graphs are analyzed in multiple ways, for example, in terms of their attributes, relations and channels. Their density, cliques, small world properties, and power laws are examined, as are the centrality and roles of certain nodes. Most social networks are compiled from surveys, diaries, or ethnographic observation. Only recently, digital data such as email and news logs are being obtained and analyzed in a more automated way.

Social network analysis is grounded in the systematic analysis of empirical data and guided by formal theory. Frequently, multi-agent modeling approaches are applied to analyze and understand the evolution of a social system.

An apt example is the research of Whyte,⁴² who spent 16 years studying the streets of New York with cameras and clipboards. His findings inform the best urban design around the world. He noted that urban congestion – the hustle and bustle of pedestrians on the sidewalks – is one of the most attractive things about cities.

Pictorial images are used extensively by social scientists to understand network data and to communicate that understanding to others. Freeman reviews the long history of image use in the field from hand-drawn images, the development of systematic procedures for locating points, to the usage of computers to produce interactive drawings of networks.⁴³

Space and time questions, for example questions such as, ‘How does the behavior x differ for different places?’ or, ‘How does behavior x within a place change over time?’ are typically separated. However, this ignores a potentially richer understanding arising from space–time interactions.

Three-dimensional virtual worlds and their inhabitants

Today, a large variety of commercial online browser systems is available for the design of customized 3D

virtual worlds. Among the most frequently used systems are Blaxxun’s online community client-server architecture (<http://www.blaxxun.com/community>), Active Worlds (AW) technology by Activeworlds, Inc. (<http://www.activeworlds.com/>), and the new Adobe Atmosphere browser (<http://www.adobe.com/products/atmosphere/>). Damer¹ and Börner⁴⁴ provide reviews of different browser systems and existing virtual communities.

Available 3D browsers vary in many details, such as their programming language, storage methods for objects, and how 3D objects are rendered for 2D display. However, all the systems mentioned above facilitate the creation of multi-modal, multi-user, navigable, and collaborative virtual worlds in 3D that are interconnected with standard Web pages and are accessible from standard computer platforms via the Internet, 24 hours a day and 7 days a week. Mouse, keyboard, and screen are used as the main human computer interface, as are commonly used interactions such as point and click, icons, and menus. However, the browsers provide access to a 3D virtual world that users can enter and explore together.

Browser systems differ in the ease of building within the worlds, their scalability and extensibility, and the size of the user communities they attract. Active Worlds (AW) stands out due to its exceptionally fast browser download, easy installation, and surprisingly low system requirements. AW is based on *Render Ware*, an interactive 3D Graphics API. It differs from VRML-based systems in the ease with which participants can build within the world. A user simply selects an existing object, makes a copy of it, and changes its properties so that it points to a different object, follows a different animation sequence, glows or not, links to web page, acts as a teleport, etc. A large object library is available. The real-time object download is based on proximity; this enables the creation and efficient exploration of very large worlds (Some worlds are the size of the UK!). AW has a large user community dating back to mid-1990s. As one of the most popular VW systems, it hosts over 6000 different worlds in its main entertainment-oriented universe, and more than 160 worlds in *EduVerse* (<http://www.activeworlds.com/edu/>), a special universe with an educational focus.

Figure 2, left shows the AW interface. It contains three main windows: a 3D graphics window populated by avatars (middle), a Web browser window (right), and a chat window (bottom). At the top are a menu bar and a toolbar for avatar actions. Users can collaboratively navigate in 3D, move their mouse pointer over an object to bring up its description, click on 3D objects to display the corresponding Web page in the right Web frame of the AW browser, or teleport to a different area. The browser maintains a history of visited places and Web pages so that the user can easily return to previous locations and pages.

Quest Atlantis (<http://atlantis.crlt.indiana.edu/>) is both a learning and teaching project that uses a 3D multi-user environment to immerse children, ages 9–12, in educational




Figure 2 The AW browser interface showing a snapshot of the culture world (left) and a map of the culture world (right).

tasks. In spring 2003, the Quest Atlantis Universe contained about 20 virtual worlds in which children travel to virtual places to perform educational activities (known as Quests), talk with other users and mentors, and build virtual personae. Quest Atlantis activities can be integrated into many settings, including classrooms, after-school programs, public libraries, and museums. Given that the main user group of this universe is children, access to this universe is restricted.

Culture and Otakhub are two of the more active, mature worlds. Culture is a 400x400 m world and Otakhub measures 500x500 m (All AW worlds are square shaped). Most subsequently discussed data sets were recorded in Culture and Otakhub. About 30 users visit both worlds daily. The movement within these worlds is restricted to walking, running, and flying to a certain height upwards. Users cannot go through objects nor go underground.

The 3D graphics window in Figure 2, left shows a snapshot of the Culture entrance area populated by diverse avatars. Figure 2, left, right shows a clickable overview map of the world to be explained in the next section.

User activity can be recorded using so-called log bots (see discussion in the sub section on chat analysis). Simple statistics on the number of unique chat utterances posted, object clicks (e.g., to access web pages, select objects or use teleports), as well as user positions for both worlds for a two month period—April through May, 2003—are shown in Figure 3. The very large number of different user positions was divided by 100 to improve readability.

Obviously, the number of users in a world positively correlates with the number of different chat utterances, clicks, and user positions (or moves) recorded.

Visualization tool set

Subsequently, we describe the set of tools designed to support the users and tasks discussed in the section on user

groups and the information needs. The tools analyze and visualize spatially- and temporally-referenced user interaction such as navigation, object manipulation, Web access, or chatting. They aim to assist users in making sense of the world, its information resources, and collaboration possibilities; to aid designers with the organization and layout of world content and the selection of interaction possibilities; and to study virtual communities. Parts of the toolkit were presented previously.⁴⁻⁶

We start by introducing the Worldmapper tool and visualizations of world growth, continue with tools to analyze chat activity and the spatio-temporal diffusion of (guided) single participants and social groups.

World layout and growth patterns

Real world cities are approached from their periphery, on roads with signs displaying the names of the city, nearby cities, major interstates, etc.; virtual worlds are entered at their center. Hence, visiting a new 3D virtual world is comparable with being blindfolded, driven (here teleported), and dropped off in the middle of an unknown city. The resulting 'lost in space' effect is amplified by the fact that users are typically unaware of the size and layout of a world, existing interaction possibilities, or major places to go. Guided tours by experienced users for first time visitors, teleports in the entrance area to major places in this world, and explanations on the world's accompanying web page are common ways to help users become familiar with a world.

This subsection introduces techniques to automatically generate 2D maps of 3D virtual worlds that show the size and layout of a world, existing interaction possibilities, and major places to go. In accordance with the user requirements discussed in the section on user groups and the information needs, the maps aim to support navigation as well as global awareness of the world's layout and interaction possibilities as well as its usage by other participants.




Figure 3 Statistics of Culture (left) and Otakhub (right.)

The layout of a world can be visualized by mapping the position, size and rotation of 3D objects in 2D. In AW, information on objects that make up a virtual world—such as time of building, builder, position, orientation, object name, its description (displayed when user moves mouse pointer over this object), and any actions such as teleporting or display of web links—are stored in the object file (also called propdump file). The dimensions of each object—used to determine and avoid encroachment during object design and manipulation—are specified in the so-called registry file (A sample AW registry file is available online at <http://www.activeworlds.com/help/registry.html>). In order to render a map, the object file is parsed for the complete set of objects, their positions, and rotations. The exact dimensions for each object are derived from the registry file; teleports and web links are identified and stored separately. Subsequently, object positions and sizes, as well as positions of teleports and web links, are scaled according to the size of the requested map.

The visualization itself consists of a reference grid (indicating the size of the world), coordinate labels (ground zero is in the center of the world at 0, 0), objects, marks for teleports and web links, and a legend. All objects are rendered in transparent green to preserve the visibility of layered objects. To show the evolution of a world, lighter colors are used for younger objects and darker colors for older (Given the availability of detailed logs about the modification or deletion of objects, previously deleted or modified objects can be mapped as well). Web links and teleports are indicated by green squares and purple plus signs respectively.

The map is clickable; i.e., a mouse click on a certain area on the map instantaneously teleports the user into the corresponding area in the 3D VW, easing navigation.

A sample map of the Culture world based on the list of objects that made up the world on Feb 20th, 2003, is shown in Figure 4. Obviously, objects in the center part of the world and the lower left part of the world were created first. There are four identical copies of the flower-like building structure in the upper end of the map. A large number of web links—typically close to objects—but very few teleports exist in this world.

The map also shows chat locations, user trails, and object click locations (to be explained in the subsequent subsections). The reference grid provides an indication of the world size; see Figure 7 for a comparison of Culture (400 × 400 world) and Otakhub (500 × 500 world).

The Active WorldMapper interface can be used to generate world maps automatically based on propdump and registry files. It is available online at <http://iuniverse-slis.indiana.edu/map/>.

Interestingly, different virtual worlds show diverse growth patterns in terms of the number of objects. Using AW, each newly created world has exactly one object. This object (and any subsequently created object) can be copied and its properties can be changed to create a wealth of exciting landscapes and architectures. Alternatively, objects from other worlds or those

created outside in 3D modeling tools can be loaded into a world.

The graph in Figure 5 shows the increase in number of objects for four different Quest Atlantis worlds: Culture, Ecology, Ocean, and Otakhub.

A gradual increase in the number of objects corresponds to in-world building activity. Otakhub appears to be the only world that was built from scratch. The other three worlds were started by loading a considerable number of objects (about 4,500 in the case of Ecology), an activity represented by steep increases in the number of objects in Figure 5.

Chat analysis

Chat log files collected from event participants have been analyzed to answer questions such as: How many users (citizens and visitors) participated in the discussion surrounding the demo as logged in the chat files? How much do users chat and who chatted the most? How many utterances are devoted to greeting, explanation, commands, questions, or other topics? How long is the average utterance length (number of words in an utterance) for different users? How often do users whisper?

Chat log bar graph visualizations⁶ provide an overview about how many citizens and visitors participated in an event, for example; a sightseeing tour, as well as the number, lengths, and type of their unique chat utterances.

A sample chat log bar graph visualization for an event on 4/14/2003 in the Culture world is shown in Figure 6. The (modified) virtual names of all users are given on the *x*-axis. Names in quotes would indicate users with visitor status. However, only citizens are allowed to enter Quest Atlantis and hence no names appear in quotes. The *y*-axis represents the number of chat utterances generated by particular users. Each utterance is represented by a rectangle of uniform height. The width and color of the rectangle corresponds to the number of words in the utterance.

The visualization in Figure 6 depicts short messages in light yellow and long messages in red. Very few people posted longer utterances (with the exception of user 008). Some people, such as 008 and 016, very actively participated in the discussion. In consultation with the owners of this world they were identified as the lead teachers. However, most users contribute about 20 utterances on average. Eight users added only 1–2 utterances.

Using chat log bar graphs for events with more than 40 people requires the application of a threshold to avoid cluttering the visualization. Commonly, the event duration that is analyzed is reduced or only users who post a certain number of utterances are displayed.

The color of the rectangles has also been used to represent the type of an utterance.⁶ This requires the semi-automatic classification of chat utterance into categories, for example; Greetings, Explanations, Commands,




Figure 4 Map of culture world with interaction possibilities and user activity.

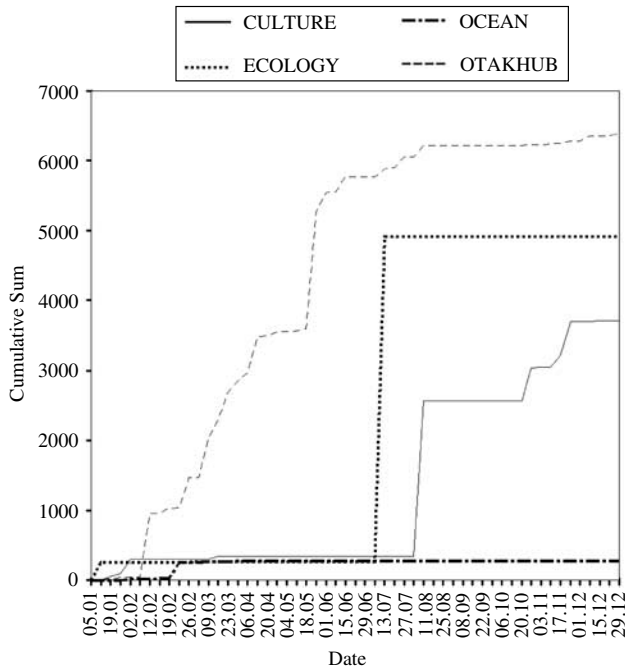


Figure 5 Increase in the number of objects over time for four Quest Atlantis worlds.

Questions, and Other. Greetings contain worlds such as ‘hello’, ‘Hi’, ‘bye’. Explanations describe certain places or objects or answer questions. Commands start with a verb and advise certain actions. Questions end with a question mark. Other utterances are typically short. The number of these kinds of utterances per event can be used to reveal the nature of the event. Linguistic techniques can also be applied to improve the chat analysis (see the discussion on future work in the section on Discussion of results and future work).

The presented chat log file analysis has a number of limitations. First, only nearby avatars are in the hearing

range and recorded. Missing time stamps for utterances makes the identification of duplicate utterances (harvested from different chat log files with overlapping hearing ranges) impossible. No spatial coordinates are available to correlate chat activity with spatial user positions. Users that may have followed the demo but did not contribute to the conversation, so-called ‘lurkers’, cannot be identified.

Subsequent work concentrated on implementing tools to record time stamped information on 3D position, navigation (e.g., teleport usage), and manipulation (e.g. web access) of the users’ avatar as well as chat activity using so-called log bots. Figure 6 (right) shows a sample bot log file. Two users, here named 001 and 002, are entering the Culture world. They move in space, exchange greetings and depart. Each line of the log file (except for the header) has seven entries: a time stamp (e.g., 417080507), the user’s name (e.g., 001), the users recorded activity (e.g., Arrival, Moving, Talking, Departure), $x-y-z$ -position (e.g., 0|0|0), and the chat utterance if the user talked.

The recorded data can be used to plot utterances of different users in the corresponding 3D space and time. Each chat utterance is represented by a triangle. The color of the triangle corresponds to the time at which it was uttered, but could also be used to identify the user who chatted. A sample visualization (using QA data) is shown in Figure 4. The center of the world shows the highest chat activity.

Spatio-temporal diffusion of participants

The log bot files (see Figure 6, right) can also be used to analyze and visualize the position of users at different points in time.

For small user groups (e.g. 1–20), the trails of participants can be plotted as polylines that interconnect discrete user positions and are overlaid on the map of the world in which user actions took place. Color coding of trails helps to capture the temporal sequence of user




Figure 6 Chat log bar graphs for Culture world (left); Bot log file (right).




Figure 7 User trail visualization for Culture (left) and Otakhub (right).

positions. Examples of user trail visualizations are given in Figures 4 and 7. Both visualizations come with a legend that provides information on what world from which universe was mapped; the start and end time of the user data log that is displayed; the color code used to indicate the age of 3D objects; symbols used to indicate links and teleports available in this world; and user actions such as chat locations, user trails, and object click (e.g., to access web pages, activate teleport, etc.) locations. The time-based color coding for the user trails is given as well. The maps displayed in Figure 7 do not show links, chatting, or click locations. They do show teleport locations and user trails as well as idle times. Idle times are represented by circles, the size of which correspond to the amount of time that elapsed before a user started moving again. By correlating user trails and available teleports, the importance and usage of different teleports can be quickly examined.

The visualization of larger user groups requires qualitatively different strategies because too many trails (see also Figure 7, right) easily clutter and make the visualization unreadable. Visualizing diffusion potentials or pressure gradients as suggested by Tobler (see the section on related work) can help to visualize general trends of user movements or the pressure to chat as the exemplified in Figure 8.

The visualization shown in Figure 8, left, was generated by dividing the world into cells of size $n \times m$. Next, a movement table was compiled, representing the number of user movement from every cell to every other cell. Subsequently, we computed the pressure to move for any cell. Cells with no movement have a value of 0. Cells which draw many users (e.g., ground zero—the entry point to the world) have negative values; cells which are left by many users (e.g., teleport areas) have positive values. The resulting movement table was then used to calculate the direction and the intensity of the flow at a certain point in geographic space. Here, for each cell, all user movement 2D vectors that started in that cell are used to calculate the average vector or ‘cell vector’. The starting point of the cell vector corresponds to the mean x and y values of the starting points of all vectors in this cell. The diffusion potentials for the Culture world are shown in Figure 8, right.

Figure 8, right shows a 3D map of the ‘pressure to chat’. The x and y axis correspond to the two horizontal dimensions of the world. The Z -axis reflects the number of chat events at this location recorded in the chat log file. This Z -axis value was scaled by a factor of 100 for readability. The resulting 3D data set was loaded into ArcMap. Then the 3D analyst extension was used to create a network of nonoverlapping triangles, also called TIN (Triangulated Irregular Network), from the height values taken as mass points. Natural breaks in the data were used to classify it into 32 classes of values. The number of breaks in the data at lower values were found to be densely concentrated. This result was expected since the data had a large number of small values

(number of chats less than 14) and a small number of large values (number of chats > 40). The elevation data set was then shaded based on height values and visualized in ArcScene. Obviously, most chat activity happens in the center of the world.

The visualizations in Figure 8 are generated by averaging the activity (chat or movement) data of all users. The results can be used to analyze and understand the general usage of a world. However, they do not reveal the activity patterns of simultaneously active social groups.

Spatio-temporal diffusion of social groups

“What attracts people most ... is other people”.⁴² Even during free exploration, participants frequently form groups and follow a leader. In particular, in large worlds, there is frequently more than one active social group. The members of real world or virtual events can be identified and clustered into social groups either by using additional knowledge about the events (e.g., classes taught by a teacher, conference sessions with one presenter and a larger audience) or by an automatic analysis of user trails. The former requires close collaboration with the owners and users of virtual worlds. The latter can be achieved by, for example, analyzing the spatial proximity of users to each other over time.

The spatial movement of a larger group can then be visualized by determining the X – Y position of the centroid of all group members for each time step. Interconnection of these centroids results in an aggregate movement trail for this particular group.

In addition, the spatial distribution or spatial homogeneity—measured by the average distance among participants—is of interest. We define the spatial homogeneity of a group to be 1.0 if all members of a group are in the same place at the same time. Obviously, a spatial homogeneity of one can only be achieved in virtual environments where multiple people can share a 3D position. The homogeneity decreases as the average distance among participants increases. The spatial homogeneity of a group is represented by the radius of a circle at the corresponding centroid point.

Mapping the continuous change of group centroid and group homogeneity resembles drawing with a brush of changing size. A sample visualization showing trails of four groups with 10 members each is given in Figure 9, left. Aggregate trails are shown in Figure 9, right. Each group is represented by a different color. While user trails for many single users easily become cluttered, aggregate user trails provide a means to visualize the trails of very large user groups.

A close examination of aggregate user trails from diverse virtual events reveals four types of group behaviors: (1) Groups can be *focused*, maintaining a high homogeneity level throughout their existence; (2) Groups can be *unfocused*—members of the group are always spread out in space; (3) Groups can be *spreading*—their members are in close proximity in the beginning,




Figure 8 Diffusion potentials (left) and chat pressure gradients (right) and for an event in the Culture world.

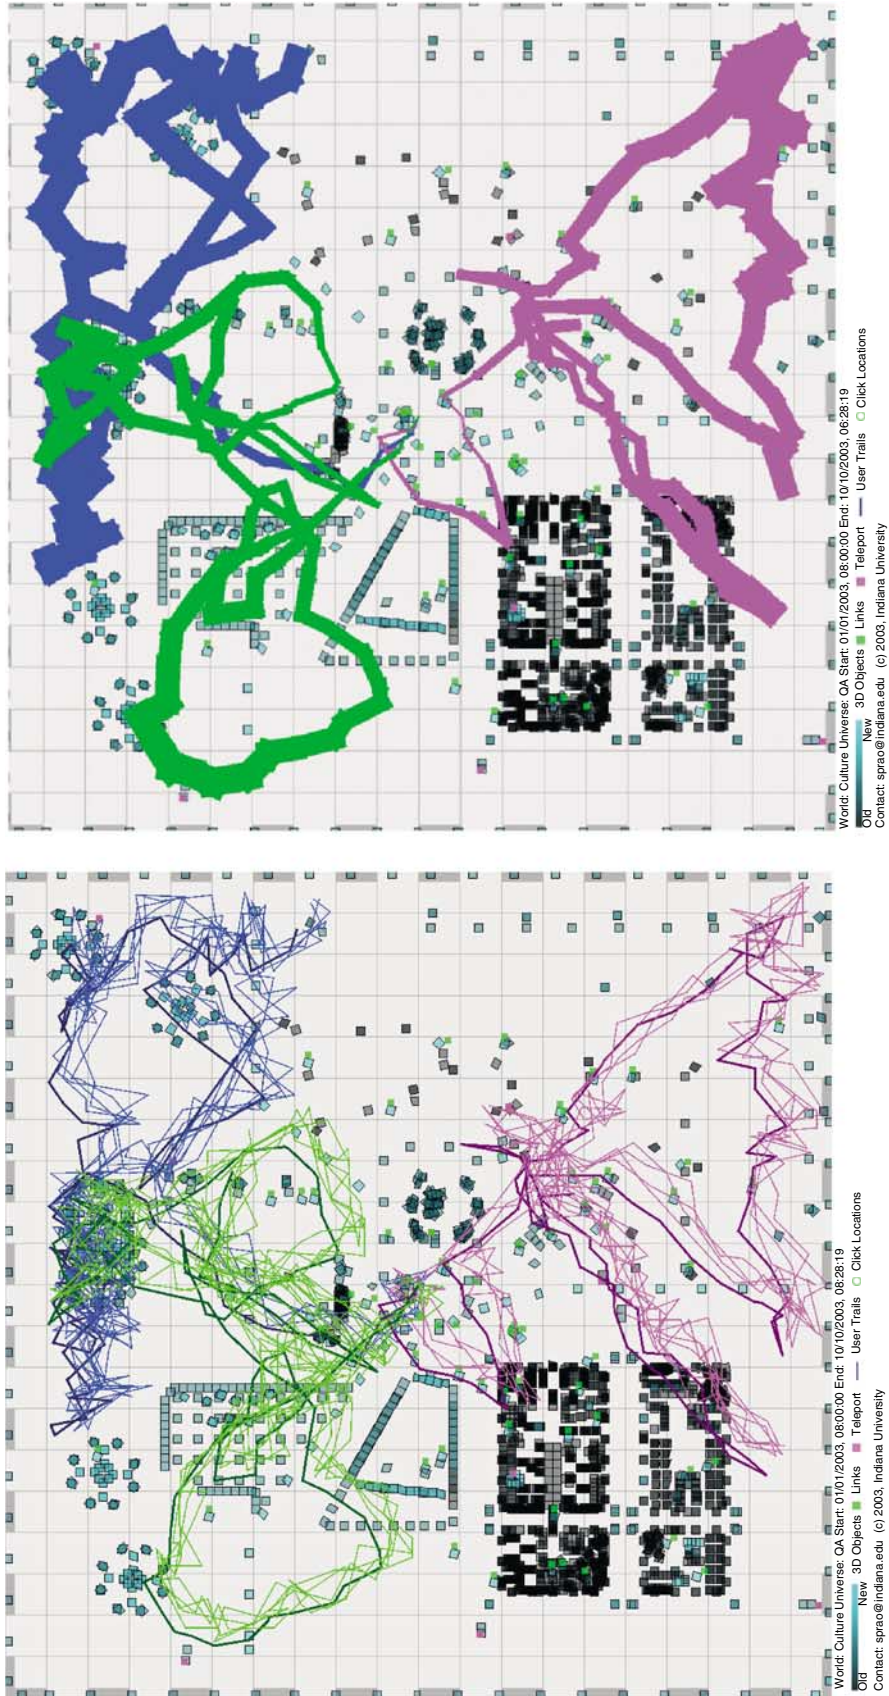


Figure 9 Visualizing user groups by single trails (left) and by aggregated trails (right).

but the average distance among them continuously increases over time; (4) Groups can *focus*—group members are far apart initially, but their average distance to each other decreases over time. Commonly, groups go through a sequence of a subset of those behaviors, e.g., phases of spreading and refocusing.

Figure 10 plots the homogeneity of major user group types over time. We assume that all users started in Ground Zero (the center entrance point of the world) at the same time. Hence, initially, the spatial distance among all participants is zero, or the homogeneity of the group is 1.0. For all other points in time, a homogeneity value of 1.0 for groups with more than one member is very atypical. Note that the spread of group participants is bound by the size of the world.

Other user group attributes can be identified by correlating the activity data of group members in different ways. The level of a group's chat activity, task focus, or semantic closeness (see discussion in next section) are just a few promising routes for future work.

Discussion of results and future work

The high spatial and temporal resolution of data collected in virtual worlds as well as the completeness of coverage of world and participant logs enables the correlation and analysis of data which was previously impossible to analyze and correlate. The presented visualizations go beyond the related work discussed in the section on Related work by providing a means to visually inspect different data sets (world layout, chat, movement, etc.) and their interrelations. They can be applied to augment the usage, design, evaluation, and study of virtual worlds. Correlating these data with event information provides a means to evaluate the success of a certain design, virtual tour, lecture unit, etc. The dynamically generated overview maps, user activity visualizations, and aggregated user group




Figure 10 Types of diffusion behavior of user groups.

visualizations can be used to support and evaluate world usage as well as to study virtual worlds and their evolving communities.

As our environment becomes increasingly smarter—sensing where we are and what we are doing via motion sensors, passive sensors (e.g., Radio Frequency Identification (RFID) tags), or active sensors (e.g., global position sensors (GPS))—the extensive collection of live user position and interaction data becomes possible. While the analysis and visualization of real world events may result in different diffusion patterns, initial results on using GPS data indicate that the techniques introduced in this paper are equally applicable to map, analyze, support, and study real world events.

Future work comprises three main areas. The first one is rather practical and motivated by the high interest of world owners and users to use the presented visualizations. It comprises the implementation of a toolset interface that lets users (mostly teachers and researchers) specify a time duration, world and/or spatial areas, and a set of users to be mapped, together with the type of map (map of the world and/or user trails and/or chat, single users or aggregated visualization, etc.) that is to be generated. In addition, our collaborators are interested in a real-time visualization of user interaction data that can be used to find currently active users, to steer an event, to change the pace of a lecture, etc.⁴⁵

Second, extending the work on the spatial-temporal diffusion of groups presented in the section on spatio-temporal diffusion of social groups, user studies will be conducted to examine the influence of spatial, semantic, and social factors on dynamic group behavior. While spatial maps can be used to depict the influence of a world layout and positions of other users on the diffusion of users, they may also help to visualize the influence of spatially referenced semantic information, information access points, on the pathways users take. In the section 1, four different user group types have been defined based on their spatial closeness. User groups types could be analogously defined based on semantic closeness. A more detailed understanding of spatial, semantic, and social influences on the usage of virtual worlds will contribute to the design of highly usable and effective learning environments, information access spaces, etc.⁴⁶

Third, linguistic techniques will be applied to improve the chat analysis and subsequent cross-indexing of linguistic/conversational data with other activity data (e.g., movement, web access data). Using dialogue act modeling,⁴⁷ not only chat unit types such as Statement, Question, Backchannel, Agreement, Disagreement, and Apology, but also the discourse coherence of a dialogue act sequence can be detected based on lexical, collocational, and prosodic cues. Chat types can be employed to color code the chat log bar graph visualization shown in Figure 6 or the map-like chat visualizations illustrated in Figure 4. The discourse coherence of a group as well as changes in discourse coherence can be used to characterize user groups not only based on their spatial diffusion



(see Figure 10), but also based on the interaction of group members via chat.

Given the increasing ease with which social activity can be monitored in both the real world and virtual space, care has to be taken to ensure that extensive user tracking and analysis benefits the users rather than invading their privacy. At any point in time, users need to be fully aware of what interaction data is being recorded, analyzed, and visualized. Otherwise, the system is likely to approach the status of a panopticon.

Acknowledgments

Over the years, a number of students have been involved in the design of the visualization tools discussed in this paper. Among them are Mitja Hmeljak, William R. Hazlewood, Alan Lin, Sy-Miaw Lin, Hakan Tuzun, Gyeongja Jun Lee, and Min Xiao. Elijah Wright conducted the user study on experiences with educational virtual world environments.

References

- Damer B. *Avatars! Exploring and Building Virtual Worlds on the Internet*. Peachpit Press: Berkeley, CA, 1998.
- Damer B, Gold S, Bruin Jd, Bruin D-Jd. *Conferences and trade shows in inhabited virtual worlds: a case study of avatars 98 & 99*. In: Heudin J-C (Ed). *Virtual Worlds*. Lecture Notes in Computer Science. Springer: Berlin; 2000; 1–11.
- Börner K. *iScape: a collaborative memory palace for digital library search results*. International Conference on Human-Computer Interaction, Vol. 1: Usability Evaluation and Interface Design (New Orleans, LA, 2001), Lawrence Erlbaum Associates: London; 1160–1164.
- Börner K. *Twin worlds: augmenting, evaluating, and studying three-dimensional digital cities and their evolving communities*. In: Tanabe M, van den Besselaar P, Ishida T (Eds). *Digital Cities II: Computational and Sociological Approaches*. Springer Verlag: Berlin; 2002; 256–269.
- Börner K, Hazlewood R, Lin S-M. *Visualizing the spatial and temporal distribution of user interaction data collected in three-dimensional virtual worlds*. Sixth International Conference on Information Visualization (London, England 2002), IEEE Press: New-York; 25–31.
- Börner K, Lin Y-C. *Visualizing chat log data collected in 3-D virtual worlds*. International Information Visualisation Conference (London, England 2001); 141–146.
- Savery JR, Duffy TM. *Problem based learning: An instructional model and its constructivist framework*. *Educational Technology* 1995; **35**: 31–38.
- Herring S, Börner K, Swan M. *Spatial reference in communication in a 3-D virtual world*. Colloquium on Language and the New Media British Association of Applied Linguistics (Cardiff, Wales 2002).
- Rheingold H. *The Virtual Community: Homesteading on the Electronic Frontier*. Harper Perennial: New York, 1993.
- Kullberg RL. *Dynamic timelines: visualizing historical information in three dimensions*. in: Media Laboratory. MIT: Boston, MA, 1995.
- Plaisant C, Milash B, Rose A, Widoff S, Shneiderman B. *Life Lines: Visualizing personal histories*. ACM CHI '96 Conference Proceedings 1996, ACM Press: New York; 221–227.
- Smith M. *Invisible crowds in cyberspace: measuring and mapping the social structure of USENET*. In: Smith M, Kollock P (Eds). *Communities in Cyberspace: Perspectives on New Forms of Social Organization*. Routledge Press: London, 1998.
- Donath JS, Karahalios K, Viegas F. *Visualizing conversation*. *Journal of Computer Mediated Communication* 1999; **4**(4).
- Paolillo J, Heald D. *Demographic participation in the discursive management of usenet*. 35th Hawaii International Conference on the System Sciences (Wailea, Maui, HI, 2002), IEEE Press: New York.
- Sack W. *Discourse diagrams: interface design for very large scale conversations*. Proceedings of the 34th Annual Hawaii International Conference on Systems Sciences, Persistent Conversations Track (Los Alamitos, CA, 2000), Institute of Electrical and Electronics Engineers (IEEE) Computer Society: Los Alamitos, CA.
- Skog T, Holmquist LE. *WebAware: Continuous visualization of Web site activity in a public space*. Proceedings of the CHI 2000 Conference on Human Factors in Computing Systems 2000; 351–352.
- Wexelblat A. *History-based tools for navigation*. Proceedings of the 32nd Annual Hawaii International Conference on Systems Sciences (Hawaii), Institute of Electrical and Electronics Engineers (IEEE) Computer Society: Los Alamitos, CA, 1999.
- Hochheiser H, Shneiderman B. *Using interactive visualizations of WWW log data to characterize access patterns and inform site design*. *Journal of the American Society for Information Science and Technology* 2001; **52**: 331–343.
- Donath JS. *Visual Who: Animating the affinities and activities of an electronic community*. ACM Multimedia 95 1995 (San Francisco, CA, 1995), ACM Press: New York.
- Erickson T, Smith DN, Kellogg WA, Laff M, Richards JT, Bradner E. *Socially translucent systems: social proxies, persistent conversation, and the design of "Babble"*. Proceedings of the CHI' 99 Conference on Human Factors in Computing Systems: The CHI is the limit 1999, ACM Press: New York; 72–79.
- Saville-Troike M. *The Ethnography of Communication*. Basil Blackwell: New York, NY, 1982.
- Xiong R, Donath J. *PeopleGarden: creating data portraits for users*. 12th Annual ACM Symposium on User Interface Software and Technology, Asheville: North Carolina, USA, 1999; 37–44.
- Smith MA, Farnham SD, Drucker SM. *The social life of small graphical chat spaces*. SIG CHI 2000. ACM Press: New York.
- Dodge M. *An Atlas of Cyberspaces. Multi-User Dimensions (MUDs) & Virtual Worlds*. <http://www.cybergeography.org/atlas/atlas.html> 2000.
- Smith A. *Mapping 30 days in ActiveWorlds*, <http://www.casa.ucl.ac.uk/30days/mapping.htm>. 2000.
- Darken R, Cevik H. *Map usage in virtual environments: orientation issues*. Proceedings of Virtual Reality (Houston, TX, 1999), IEEE Computer Society Press: Los Alamitos, CA; 133–140.
- Darken RP, Sibert JL. *Wayfinding strategies and behaviors in large virtual worlds*. SIGCHI '96, Vancouver: British Columbia, Canada, 1996.
- Lynch KJ. *The Image of the City*. MIT Press: Cambridge, MA, 1960.
- Bowman D, Wineman J, Hodges L, Allison D. *The educational value of an information-rich virtual environment*. Presence: Teleoperators and Virtual Environments 1999; **8**: 317–331.



- 30 Schafer WA, Bowman DA, Carroll JM. *Map-based navigation in a graphical MOO*. 2001, Computer Science, Virginia Tech.
- 31 Herring SC (ed). *Computer Mediated Communication. Linguistic, Social and Cross-cultural Perspectives*. John Benjamins Pub Co: Amsterdam/Philadelphia, 1996.
- 32 Naper I. *System features of an inhabited 3-D virtual environment supporting multimodality in communication*. 34th Annual Hawaii International Conference on Systems Sciences (Maui, Hawaii, 2000), Los Alamitos, CA: Institute of Electrical and Electronics Engineers (IEEE) Computer Society; 4040.
- 33 Fry B, *Organic information design, in Media Lab*. MIT; Master's Thesis, 2000.
- 34 Tobler W. *Smooth pycnophylactic interpolation for geographical regions*. American Statistical Association, 1979; **74**: 519–536.
- 35 Tobler W. *push pull migration laws*. Annals Association of American Geographers 1983; **73**: 1–17.
- 36 Tobler W. *A model of geographic movement*. Geographical Analysis 1981; **13**: 1–20.
- 37 Hillier B, Hanson J. *The Social Logic of Space*. Cambridge University Press: Cambridge, 1984.
- 38 Whyte WH, *The Social Life of Small Urban Spaces*. The Conservation Foundation: Washington, DC, 1980.
- 39 Stonor T, Campos MBdA, Smith A. *Towards a walkability index*. Walk 21: 3rd International Conference, Steps Towards Liveable Cities, 8-11.05 2002 (San Sebastian, Spain, 2002).
- 40 Stonor T. *From the passengers' perspective: function follows form*. Interchange Conference 2001 (Docklands), <http://www.spacesyntax.com/news/presentations.htm>.
- 41 Wasserman S, Faust K, *Social Network Analysis : Methods and Applications (Structural Analysis in the Social Sciences, Vol. 8)*. Cambridge University Press, Cambridge, 1994.
- 42 Whyte WH, *City: Rediscovering the Center*. Anchor Books: New York, 1971.
- 43 Freeman LC. *Visualizing social networks*. Visualizing Social Networks 2000; **1**.
- 44 Börner K, *Virtual World Technologies for Digital Cities of Tomorrow*. 2001, Invited Talk at the Panel Collaboration in Building Future Digital Cities. The Second Kyoto Meeting on Digital Cities, Kyoto, Japan, <http://ella.slis.indiana.edu/~katy/research/dc01-panel.pdf>.
- 45 Börner K, Penumarthy S, DeVarco BJ, Kerney C. *Visualizing social patterns in virtual environments on a local and global scale*. Accepted for Digital Cities 3: Local Information and Communication Infrastructures: Experiences and Challenges 2003 (Amsterdam), (in press).
- 46 Stolcke A, Ries K, Coccaro N, Shriberg E, Bates R, Jurafsky D, Taylor P, Martin R, Ess-Dykema CV, Meteer M. *Dialogue act modeling for automatic tagging and recognition of conversational speech*. Computational Linguistics 2000; **26**: 339–373.