Chen, Chaomei, Katy Börner. 2005. From Spatial Promiximity to Semantic Coherence:
A Quantitative Approach to the Study of Group Dynamics in Collaborative Virtual Environments.
PRESENCE:Teleoperators and Virtual Environments,Special Issue on Collaborative Information Visualization Environments
14(1): 81-103. MIT Press.

Cover Page:

From Spatial Proximity to Semantic Coherence: A Quantitative Approach to the Study of Group Dynamics in Collaborative Virtual Environments

Chaomei Chen <Corresponding Author>

College of Information Science and Technology

Drexel University

3141 Chestnut Street

Philadelphia, PA 19104-2875, USA

Email: chaomei.chen@cis.drexel.edu

Phone : +1 215 895 6627

Fax : +1 215 895 2494

Katy Börner

School of Library and Information Science

Indiana University

10th Street & Jordan Avenue

Bloomington, IN 47405, USA

Email: katy@indiana.edu

From Spatial Proximity to Semantic Coherence: A Quantitative Approach to the Study of Group Dynamics in Collaborative Virtual Environments

Chaomei Chen

College of Information Science and Technology

Drexel University

3141 Chestnut Street

Philadelphia, PA 19104-2875, USA

Katy Börner

School of Library and Information Science

Indiana University

10th Street & Jordan Avenue

Bloomington, IN 47405, USA

Abstract

This article describes a quantitative approach to the study of group dynamics in Collaborative Information Visualization Environments (CIVEs). This approach characterizes group dynamics in terms of two concepts introduced in this article - *spatial proximity* and *semantic coherence*. The concepts are crucial to the understanding of profound interrelationships between spatial, semantic, and social navigation. Furthermore, this article describes three visualization techniques - semantic indentation chat sequence displays, activity maps, and clock-face maps - that permit the identification of important features of group interaction that are related to semantic coherence and spatial proximity. The approach is illustrated by applying it to the analysis of an empirical study in which four groups of subjects performed collaborative search tasks through three-dimensional visualizations of knowledge domains. The major contribution of the work is the conceptualization and quantification of group coherence as a generic methodology for the study of a range of collaborative virtual environments such as collaborative learning, distance learning, social networks, collaborative information visualization, or digital libraries. Further research challenges for the study of group behavior in collaborative information visualization environments are identified.

1. Introduction

Collaborative Information Visualization Environments (CIVEs) differ from Collaborative Virtual Environments (CVEs) because of the integral role of explicit spatial-semantic models such as abstract information visualization models. Research in information visualization has generated a broad range of graphical representations and three-dimensional structures of information that does not have inherited geometry. Notable examples of information visualization include Cone Trees (Robertson, Mackinlay, & Card, 1991), Hyperbolic Views (Lamping, Rao, & Pirolli, 1995), and Treemaps (Johnson & Shneiderman, 1991). A substantial amount of work has been done in the area of graph visualization due to the ubiquitous nature of the widely seen node-and-link representations (Chen, 1999a; Herman, Melançon, & Marshall, 2000). In this article, we are particularly concerned with CIVEs that draw upon graph visualization as part of the underlying virtual environments. Indeed, we further limit the scope of our study to three-dimensional landscape views of a visualized graph. The landscape metaphor not only closely mimics the physical world in which we live, but it also supports a number of well-developed ways that one may navigate. As users move back and forth in a three-dimensional landscape model of an abstract information structure, what do their movements tell us? In a virtual world, what does spatial proximity mean to collaborators? Will it alter how they communicate with each other? What if their activity space itself is represented as a virtual world? These are some of the typical questions that one may ask when trying to understand the usage of a CIVE. We aim to develop a quantitative approach that can help us to measure, monitor, analyze, and understand the dynamics of collaborative activities and other concurrent events in a CIVE.

In this article, we first introduce the notions of spatial, semantic, and social navigation in a virtual environment. Then we define the concepts of spatial and semantic coherence of a collaborating group as the basis of our quantitative approach. We explain the measurement of spatial and semantic coherence and how these measurements can be analyzed through the use of several visualization techniques. In particular, these visualization techniques are used in an empirical study in which four groups of users perform search tasks within a three-dimensional landscape model. Finally, we discuss the implications of the approach.

2. Related Work

A fundamental question in the study of a collaborative virtual environment is what factors influence how users navigate in a virtual world and in turn, when, where, and why they interact with others. A number of further questions that also need to be addressed may be derived from this core inquiry. For instance, how do users find what they need? How can a collaborative environment facilitate their tasks? What types of information are valuable to users? What strategies can they use to improve their task performance? Subsequently, we review research on different types of navigation as well as empirical studies of individual and group behavior in collaborative virtual environments.

2.1 Spatial, Semantic, and Social Navigation

Dourish and Chalmers (1994) suggested three types of navigation: spatial, semantic, and social. Spatial navigation relies on spatial attributes of a virtual or physical environment as a major source of guidance. For example, we walk along corridors in a building rather than walk through the wall. Spatial navigation has been studied by scientists in geography, urban anthropology and urban development as well as cognitive psychology. Similarly, when users navigate in a three-dimensional virtual environment, spatial attributes may have direct impact on how the users would perform, behave, and act. Indeed, this understanding has served as the premises of the design principles of many virtual environments (Benford, Snowdon, Colebourne, O'Brien, & Rodden, 1997; Benford et al., 1995; Darken & Sibert, 1996). To clarify the role of spatial navigation, way-finding in virtual environments has been studied within the field of virtual reality (Darken & Sibert, 1996). Findings indicate that collaboration in a virtual environment does not directly require users to maintain a certain degree of spatial proximity. Instead, researchers have shown that spatial proximity can be used to initialize transitions between different modes, for example, from a loosely coupled mode to a tightly coupled one. Spatial proximity, therefore, provides additional mechanisms for interacting in a virtual world.

Semantic navigation is guided by the meaning intrinsically attached to a virtual or physical environment. In an abstract space such as the Web, following hyperlinks is semantic navigation; we act based on the assumption that the other end of a hyperlink may contain what we need if we are properly informed by the semantic implication of the anchor of the hyperlink. Semantic navigation has been intensively studied in the field of hypertext, notably the issues of being “lost” in hyperspace and cognitive overload (Conklin, 1987). Various techniques have been proposed to reduce such cognitive overload, including using graphical representations of the underlying hyperspace and visualizing the history of the navigation.

Social navigation differs from both spatial and semantic navigation by drawing various navigational cues from usage signs (e.g., read-wear and edit-wear (Hill, Hollan, Wroblewski, & McCandless, 1992), footprints (Wexelblat, 1999; Wexelblat & Maes, 1997), dog-ears (Dieberger, 1997)), the gathering and movement of fellow tourists at the same attraction site, or the behavior of fellow users of the same information space (Crossley, Davies, McGrath, & Rejman-Greene, 1999). Interpersonal relationships in two- and three-dimensional virtual worlds have also been an important issue (Crossley et al., 1999; Erickson, 1993; Fry, 2003; Greenhalgh & Benford, 1995; Jeffrey & Mark, 1998). The concept of constrained navigation environments has been studied by Hanson, Wernert, and Hughes (1997). Constrained navigation appropriately restricts the user's degree of freedom when there is a mismatch between the goal of navigation and the user's search knowledge of the exploration domain. Social dynamics in virtual worlds has been studied from the perspectives of linguistics and small group behavior (Tromp et al., 1998). A recent study by Börner and Penumarthy (2003) visualizes the growth of virtual worlds along with an analysis of social diffusion in the virtual worlds. Furthermore, they define a measure of group dynamics in order to characterize groups as focused, unfocused, expanding, shrinking, or any time sequence of these characteristics. The virtual world in their study did not incorporate semantic aspects; therefore, a study of semantic navigation was not possible.

2.2 Empirical Studies of CVEs

CIVEs offer unique opportunities for users to foster collaborative activities within an environment that has intrinsic semantic values. Such contextual semantics provide valuable cues for users to determine how they may navigate and how they might interact with each other. There are relatively few empirical studies in the broader context of CVEs and even fewer empirical studies of CIVEs (see (Greenhalgh, 1997; Tromp, Steed, & Wilson, 2003). One reason for the lack of empirical studies is due to the lack of a unifying conceptual framework; the problem is also reflected in the fact that the basis of a generic quantitative approach is essentially missing. Such quantitative approaches are particularly important if one aims to make sense of the dynamics of diverse activities in a collaborative virtual environment that has unique spatial and semantic components.

A collaborative virtual environment should provide adequate cues for users to understand the underlying design rationale. Is it a formal place for visitors to discuss serious issues, or is it meant to be a virtual place for relaxed and informal gathering? How closely could two avatars stand next to each other without being considered rude? For example, a virtual world of a national park would replicate various attractions in the real world. Visitors to such virtual worlds would be able to glean specific clues from their environment and adapt appropriate social behavior. In contrast, visitors to a virtual world that consists of an open area with few contextual features would find it harder to anticipate how others might behave. In this article, we are mainly concerned with virtual environments that go beyond the typical chatroom-like virtual environment. Indeed, we are interested in virtual environments whose purpose is to help people find and manage information and expertise collaboratively. For example, a CVE of a digital library could help its visitors find electronic books of interest.

An early attempt to explore collaborative search in a collaborative information visualization environment was conducted in StarWalker (Chen, 1999a; Chen, Thomas, Cole, & Chennawasin, 1999). StarWalker is characterized by an information visualization model embedded in a multi-user virtual environment. The StarWalker study was motivated by the expectation that a visualization model in a virtual world could trigger more focused social interaction in relation to the content of the visualization model; the need for an integrative approach to spatial, semantic, and social navigation was also recognized. Although avatars were used in the StarWalker study, the usage data were limited to chat logs and a sequence of screenshots. Since the focus was on whether the visualization model indeed triggered various transitions in social episodes of collaborative search, the gathering of crowds around the visualization structure was verified by visual inspections of screenshots rather than by quantitative measures. In this article we aim to develop a quantitative methodology that can help analysts make sense of what users do in virtual environments in terms of where they go, what they chat about with special reference to the underlying virtual environment as a situated setting, and how they interact with each other.

Research in computer-supported cooperative work (CSCW) has inspired a number of concepts we will define in this study, including transitions between tightly and loosely coupled modes, and the role of spatial proximity as a non-verbal communication protocol. CSCW researchers have special interests in collaborative modes that can be characterized by the tightness of coupling between collaborators; a group can be tightly coupled or loosely coupled (Harrison & Dourish, 1996; Kraut, Egido, & Galegher, 1988; Olson, Card, Landauer, & Olson, 1993). The tightness of coupling is not necessarily based on spatial metrics since metrics may not be well-defined in some collaborative environments. For example, a group can be defined as tightly coupled if its members maintain and use wide-bandwidth communication channels all the time; consequently, group members may exchange information synchronously. In contrast, a loosely coupled group might use only email to communicate amongst members. In this case, all communications are asynchronous, thus they are less efficient than groups that use synchronous channels. Obviously, there is a wide range of options to define the tightness quantitatively if geographical or geometric metrics are available. If the tightness is defined as the distances between group members, then a group with all its members on the same floor of a building would be tightly coupled, whereas a group with its members across several continents would be loosely coupled. In many environments, spatial metaphors have intrinsic semantics that are subject to users' interpretation. Virtual environments with embedded abstract visualization models fall into this category.

3. Metrics for Group Dynamics

In this section, we introduce a set of metrics for the analysis of different aspects of collaborative navigation. In particular, we define a set of quantitative measures of the *spatial proximity* and the *semantic coherence* of a group.

3.1 Spatial Proximity

The spatial proximity of a group is defined as the ratio between the mean distance between group members over the entire session and the total distance traveled by the group as a whole. This definition compensates for groups with members who traveled a large amount of virtual miles. We are not interested in the extreme scenario in which no one would ever make a move during an entire session.

A generic definition of the spatial proximity δ is given as follows. Given a group of N members, if the total distance traveled by each member is d_{mi}, and the total amount of distance traveled by a group is $\sum_{i=1}^{N} d_{m_i}$, then the average of the total distance traveled is $\frac{\sum_{i=1}^{N} d_{m_i}}{N}$. The distance between group members i and j at time t is denoted by $d_t(m_i, m_j)$. The group diameter at time t is the mean of member-to-member distances $\frac{\sum_{i,j}^{N} d_t(m_i, m_j)}{N}$. The average of group diameter across the entire time series gives the overall group diameter; if T denotes the length of the time series, then the overall group diameter is $\frac{\sum_{i,j}^{N} d_t(m_i, m_j)}{N \bullet T}$. The spatial proximity of the group is therefore given by the following formula:

$$\delta = 1 - \frac{\dfrac{\sum_{i,j}^{N} d_t(m_i, m_j)}{N \bullet T}}{\dfrac{\sum_{i=1}^{N} d_{m_i}}{N}} = 1 - \frac{\sum_{i,j}^{N} d_t(m_i, m_j)}{T \bullet \sum_{i=1}^{N} d_{m_i}}$$

Note that this definition is applicable to single-person groups as well as for groups of two or more members. All groups that participated in the study reported in section 5 had two members.

3.2 Semantic Coherence

The semantic coherence of a group measures the spatial proximity between the group as a whole and the visualization models. Given a group of N members to search for O relevant target documents, or other types of targets, calculate the distance at time t between the group and the set of search targets; this is simply the shortest distance between any group member and any target object at time t $\min_{i \in N, j \in O} d_t(m_i, o_j)$. This shortest distance measures the proximity between the group and the most relevant objects as far as the search tasks are concerned. Finally, average the shortest distance time series over the entire session. In this article, we refer to the average shortest distance as *the group-to-target distance*. The semantic coherence is derived from the ratio between the group-to-target distance and the average of group travel distance.

$$\sigma = 1 - \frac{\frac{\sum_{t}^{T} \min_{i \in N, j \in O} d_t(m_i, o_j)}{T}}{\frac{\sum_{i=1}^{N} d_{m_i}}{N}} = 1 - \frac{N \bullet \sum_{t}^{T} \min_{i \in N, j \in O} d_t(m_i, o_j)}{T \bullet \sum_{i=1}^{N} d_{m_i}}$$

Given the same amount of total group travel distance, a group is semantically tight if it has a relatively small group-to-target distance; the opposite is obviously a semantically loose group. Like spatial proximity δ, the semantic coherence σ is also applicable to single-user groups as well as groups of two or more users.

3.3 Group Coherence

The overall coherence of a group combines the spatial proximity δ and semantic coherence σ with an emphasis on the semantic coherence. More precisely, the impact of spatial proximity is reduced to the square root of its original value.

$$\Sigma = \frac{\sigma\sqrt{\delta}}{\sqrt{\sigma^2 + \delta}}$$

Given the same semantic coherence σ, a group with higher spatial proximity will also have a higher group coherence measure. Similarly, given the same spatial proximity δ, the higher group coherence goes with the semantically more coherent group.

These definitions are also consistent with our general observations. The majority of existing studies of CVEs emphasize the necessity of spatial proximity in maintaining effective collaborations. However, CIVEs may lead to a different scenario in which the superiority of spatial proximity may be overridden by semantic coherence. Being spatially farther apart in a virtual world does not necessarily indicate a loosely coupled group. When two friends are connected by their cellular phones, they could be miles and miles away geographically but maintain an instant communication bond. Similarly, if a group as a whole can manage to stick to the vicinity of the most relevant part of the virtual world, then they don't have to maintain their spatial proximity. In fact, one may even argue that if the semantic navigation is strong, it shows a better quality of group if members are loosely coupled spatially, but tightly coupled semantically. At least in the setting of our study, the need for spatial proximity is no longer the predominant navigation concern.

In summary, an ideal group would be able to maintain a relatively small semantic distance to the visualization models while its members enjoy the freedom of traveling over large spatial distance without worrying too much about spatial proximity to their own members; and obviously, they need to navigate and access information in an efficient manner.

3.4 Group Coherence Space

The two factors that influence the coherence of a group, the total distance traveled (α) and its overall diameter (β), can be used to characterize various coherence properties of groups in an abstract space, called the 'group coherence space', as shown in Figure 1. Each point g in the space represents a group with its dynamics characterized by a set of five values, namely $g = (\alpha,$

β, δ, σ, Σ). The horizontal dimension represents the distance traveled by a given group. The length of the distance increases from the left to the right. The vertical dimension represents the diameter of a group; its value increases from the top to the bottom. Given a point in the coherence space, its spatial proximity δ, semantic coherence σ, and overall coherence Σ can be calculated. The three values determine the color of the group on the plane. Thus we can see the distribution of group dynamics over a wide variety of scenarios. We can find not only the positions of the groups in our empirical study, but also patterns of how the coherence measures can be influenced by the two factors.

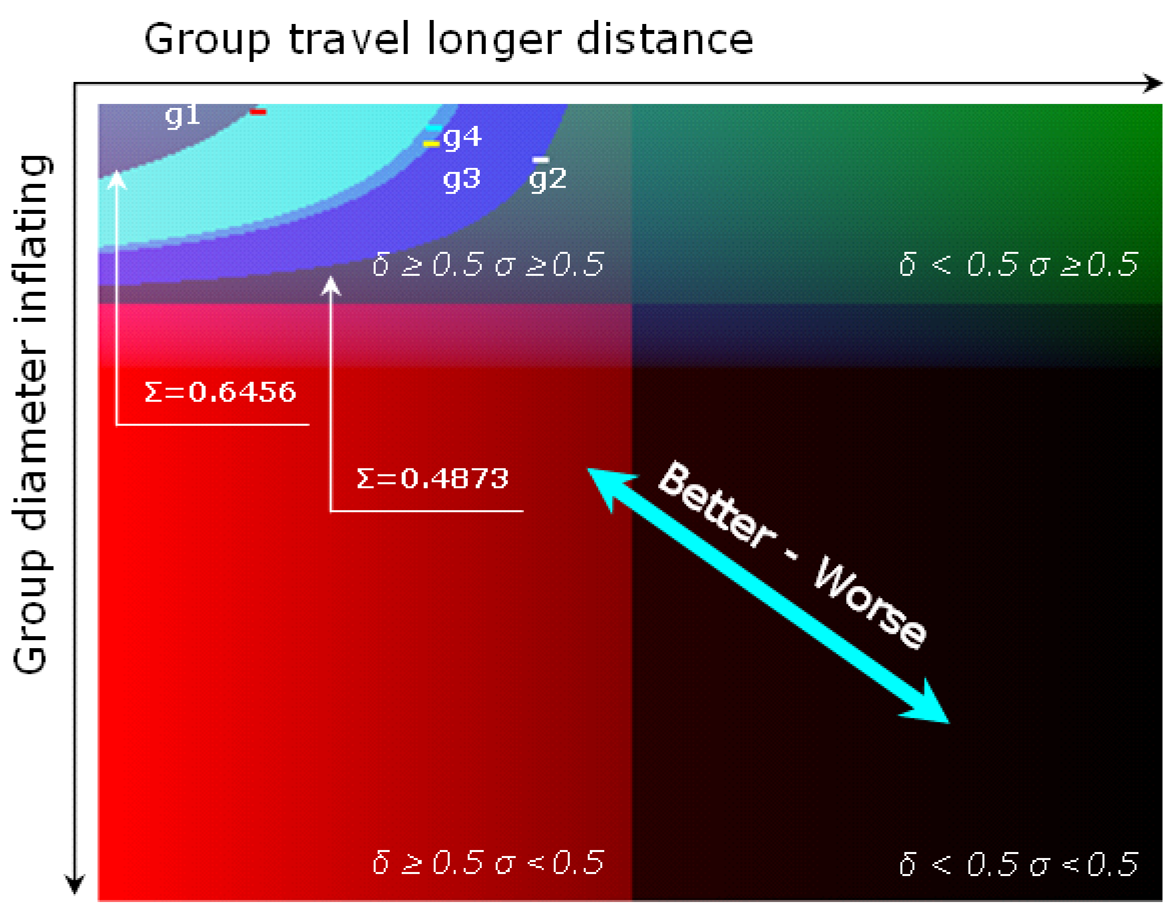

Figure 1. Group Coherence Space, colored by the different measures, showing the positions of four groups in the subsequently explained empirical study. The contour belts in the upper right

quadrant, which indicate regions of similar strengths of spatial-semantic coupling, are discussed in section 5.3.

Four quandrants can be distinguished based on the strength of the spatial and semantic coupling of groups. The upper left quandrant contains groups that have tight coupling both spatially and semantically; the upper right quandrant contains groups with strong semantic coupling but weak spatial coupling. Groups in lower left quandrant have strong spatial coherence, but are weak in semantic coupling. Groups in lower right quandrant are weak in both spatial and semantic coupling. In general, the closer a group is to the upper left corner, the stronger its coherence is; one may expect to find a highly focused and experienced group in this area. In contrast, the closer a group is to the lower right corner, the weaker it is coupled. We imagine a diffused social gathering in a virtual world with no central attraction objects and no semantic attachments would be positioned in this area.

4 Visualizations of Group Dynamics

As part of our quantitative methodology, a number of new visualization techniques were developed to facilitate the analysis of the spatio-temporal user interaction data typically generated in CIVEs. They comprise a semantic-indentation chat display, activity maps, clock-face maps, and a visual display of a group coherence space. These techniques enable analysts to study the interrelationship between spatial proximity and semantic coherence of a group in a consistent and unified conceptual framework. The emphasis of the role of visualization aims to identify connections between transitions in the discourse and patterns of group members' movements. The semantic-indentation chat sequence display and clock-map visualizations were developed at Drexel University, whereas activity maps were developed at Indiana University.

4.1 Semantic-Indentation Chat Sequence Display

The semantic-indentation display was developed at Drexel University, for this study, in an attempt to intuitively show how far away the group was from an information visualization model when chat utterances were made. It displays chat utterances in the original sequence, but the amount of the space indented from the left margin is proportional to the shortest semantic distance at the time of the current utterance. If the position of an utterance was shifted to the right, then it implies that the subject was moving away from the relevant target area. Our assumption is that because users need to inspect individual objects in the visualization models, they are more likely to perform their tasks by moving closer to areas of search than searching from a constant overview. This assumption is in part based on our everyday experience in map reading, when we want to take a closer look at where we search, and in part based on earlier studies of search behaviors with semantic maps (Chen, Cribbin, Kuljis, & Macredie, 2002). Sample visualizations are shown and discussed in section 5.4.

4.2 Activity Maps

Activity maps developed at Indiana University can be used to analyze and visualize spatio-temporal interaction and diffusion patterns in 3D virtual worlds (Börner, 2002; Börner, Hazlewood, & Lin, 2002; Börner & Penumarthy, 2003). In activity maps, user interaction data (e.g., movement, chat, web access, teleport activity) is overlaid on top of a map of the virtual world in which those actions were recorded. Activity maps can be generated for specific areas, time durations, users, and user groups. They can be used to display single, multiple, or all user activities. In addition, they are used to support social navigation, to evaluate and optimize three-dimensional virtual worlds, and to study their evolving communities (Börner, Lee, Penumarthy, &

Jones, Forthcoming; Börner, Penumarthy, DeVarco, & Kerney, Forthcoming). Sample visualizations are shown and discussed in section 5.4.

4.3 Clock-Face Maps

Clock-face maps, developed at Drexel University, are designed for the examination of various connections between chat sequences and the spatial configuration of the virtual worlds. A clock-face map displays one utterance at a time in the sequence of the original chat; the appearance of the utterance is also cross-referenced to the spatial position in the virtual world. Utterances made at different times may be linked by straight lines on the clock-face map if their timestamps were close enough to each other; therefore, it is possible to trace the movement of a subject by following such straight-line trails. These interactive maps enable us to examine all details associated with a chat session. Sample visualizations are shown and discussed in section 5.4.

5. Empirical Study

An empirical study was conducted to demonstrate and validate the developed measures and visualizations. Rather than testing a null hypothesis we wanted to collect evidence to support the notion that abstract semantic models influence spatial, semantic, and social navigation patterns. Correspondingly, the empirical study is designed to determine whether and to what extent subjects are attracted by the visualization models. In other words, how does the presence of abstract semantic structures in a collaborative virtual environment influence the navigation paths of the subjects? What variations of navigation paths across different groups exist and to what extent do they help to understand group dynamics in such settings?

5.1 Collaborative Information Visualization Environment

Diverse information visualizations pose a need for collaborative examination. The examples selected here are knowledge domain visualizations (KDVs) that are generated based on large amounts of publication data. KDVs help to visualize the semantic space of researchers, publications, major research areas, experts, institutions, grants in a research area of interest, the import and export of research between subdomains, the dynamics (speed of growth, diversification) of scientific fields, and scientific and social networks.

The accelerated speed of knowledge creation has led to increasing specialization of experts. No single person has all the expertise needed to examine a knowledge domain visualization. Collaborative virtual environments enable experts around the globe to examine information visualizations collaboratively.

The collaborative knowledge domain visualization used in our study was generated based on the citation structure reflected in a set of 2,485 journal articles published between 1983 and 2003 on the topic of *information visualization*. The bibliographic records were retrieved from the Web of Science of the Institute for Scientific Information by searching for articles that cited the work of Edward Tufte. Each record contains information on the title, authors, year and source of publication as well as a list of papers that it referenced, or cited.

Information visualization research efforts have grown considerably over the last decade. In order to investigate the evolution of the citation networks, we evenly divided the entire 20-year time interval of the dataset into two sub-periods. Two visualization models will be produced: one for the period between 1983 and 1992, and a second for the period between 1993 and 2003.

Knowledge domain visualizations can be generated using diverse approaches, see (Börner, Chen, & Boyack, 2003; Chen, 2003) for an extensive review. In general, the process entails determining the semantic similarity of all documents and then laying out documents in a way such that similar documents are closer in space and dissimilar ones are further apart. In addition, color coding and additional symbols can be used to denote specific properties of documents such as year of publication, topics covered, or number of received citations. Subsequently, we demonstrate the generation of knowledge domain visualizations using the method detailed in (Chen, 2003; Chen & Paul, 2001; Small, 1999).

Mapping the Semantic Space of Documents

Chen's approach uses co-citation frequencies as the measure of the strength of the connection between two documents. Two documents are assumed to be more similar to each other if they are frequently cited together by a third article (Chen, 1999b; White & McCain, 1998).

The node-and-link visualization of graphs was chosen because of the following considerations: 1) node-link visualizations have been widely studied and they come naturally with the citation network data; 2) graph-theoretical visualizations are compatible with the three navigation types of interest: spatial, semantic, and social. The concepts of spatial proximity and semantic coherence can be clearly defined with such visualizations; 3) users probably have used similar representations so they can better concentrate on tasks.

One way to increase the readability of a complex co-citation network of documents is pruning. A pruning process typically preserves salient links in a network but discards other links. The

Pathfinder network scaling technique was used because of its satisfactory performance in a number of earlier studies of co-citation networks (Chen, 2003). The resultant network represents documents as nodes and strong co-citation interrelations among them as edges. Pathfinder networks tend to capture local structures more accurately than conventional methods such as multidimensional scaling (MDS) (Jolliffe, 1986; Thurstone, 1931).

Augmenting the Semantic Space of Documents

Besides mapping documents and their semantic relationships, it is beneficial to augment the visualization with information on the topics covered by documents or the citation history of individual documents.

Topics covered by an article can be identified using principal component analysis (PCA). PCA is a commonly used method for dimensionality reduction (Robertson, Card, & Mackinlay, 1993). It can characterize the original dataset with fewer dimensions, or factors. The factor loading coefficient can be regarded as a quantitative measure of the extent to which a document is defined by a given dimension. A document with a large factor loading coefficient tends to be one of the most salient representatives of an underlying dimension.

A large volume of citations to a document is generally regarded as an impact indicator of the document in its own field. The number of citations will be indicated by a pole, the height of which is proportional to the total amount of citations to the corresponding document accumulated over the entire period.

Figure 2 shows a screenshot of one generated visualization model. Spheres denote documents, and cylinders connecting spheres denote the strongest co-citation links as preserved by

Pathfinder network scaling. Article nodes are color coded using the factor loading coefficients of the largest three PCA factors. For example, spheres in red or colors close to red would be documents from the largest specialty, or the mainstream of the field; spheres in green or nearly green would identify the second largest specialty; spheres in blue or similar colors would identify the third. Any colors in between are due to the combination of the three basic colors; for example, a white or gray colored sphere would be a document that is almost equally recognized by all three major specialties. In other words, those specialties are not mutually exclusive.

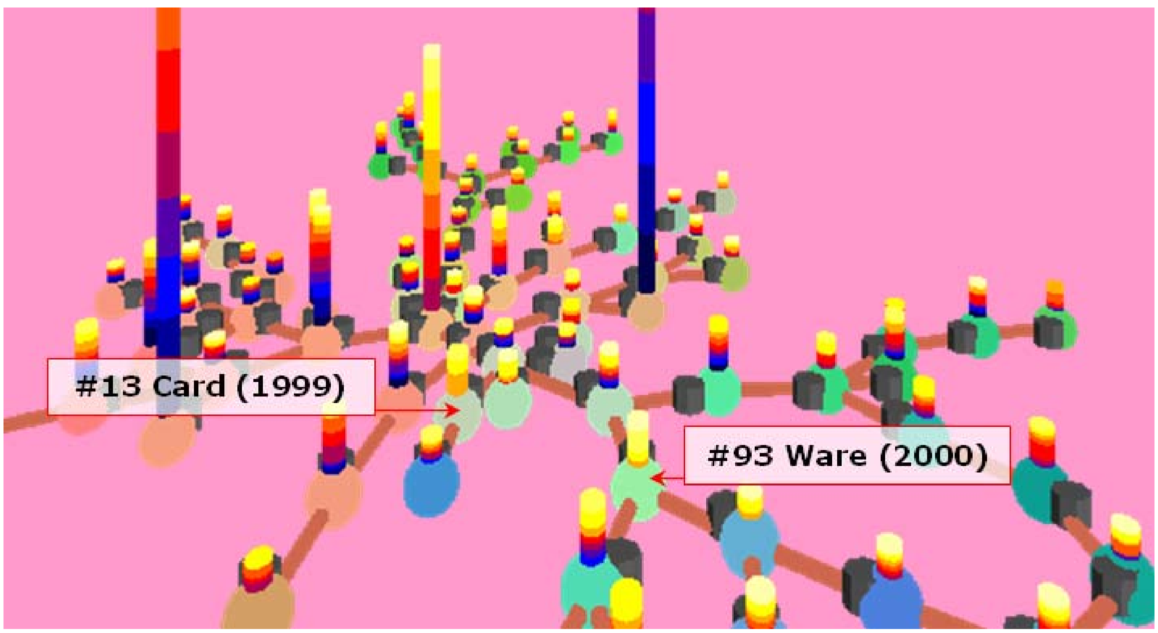

Figure 2. A visualization model of citation patterns.

The poles stemming out of the spheres represent the citation history of individual documents. The poles are color coded to indicate in which years citations have been acquired. The earlier years are given in darker colors and later years lighter colors. Along with the total height of a citation pole, the color spectrum of a citation pole is designed to give the viewer a glimpse of the citation history. For example, an all-time popular document would be shown as the tallest citation pole with an almost evenly distributed color band, whereas a rising-star document would have a considerably larger portion of its citation colors in light colors as it must have attracted many citations in a relatively short and recent period of time.

Although a detailed analysis of the structure of the field of information visualization is beyond the scope of this article, its major specialties in two periods between 1983 and 2003 are outlined as follows. The first period (1983-1992) was dominated by the highly cited Tufte books, Bertin's *Semiology of Graphics*, and Cleveland's several articles on statistical graphics. The citation visualization also featured documents on graphical data analysis and Exploratory Data Analysis (EDA). In essence, the second period (1993-2003) heralded the emergence of information visualization as an independent field. Tufte's books still attracted by far the highest citations. New documents with high citation rates in the second period include Robertson, Card, and Mackinlay's classic article on information visualization and Cleveland's *Visualizing Data*. The single most noticeable trend between the two citation landscapes would be the emerging documents on information visualization and the receding of statistical graphics from its prominent position in the 1980s.

Creating CIVE's Using Active Worlds

The information visualization models discussed so far were generated in Virtual Reality Modeling Language (VRML). They were then subsequently converted to the RenderWare format (RWX), and imported into a 100-by-100-meter collaborative virtual world powered by Active Worlds technology (http://www.activeworlds.com/).

The Active Worlds user interface and its various controls are shown in Figure 3. The interface is divided into four major areas. The left area is a list of the virtual worlds available for access. The upper middle area provides access to a three-dimensional virtual world, here filled with information visualization models. The lower middle area is the chat area, including an input window and a window showing the history of the current chat session. The right area is a Web page which is here used to show details of documents.

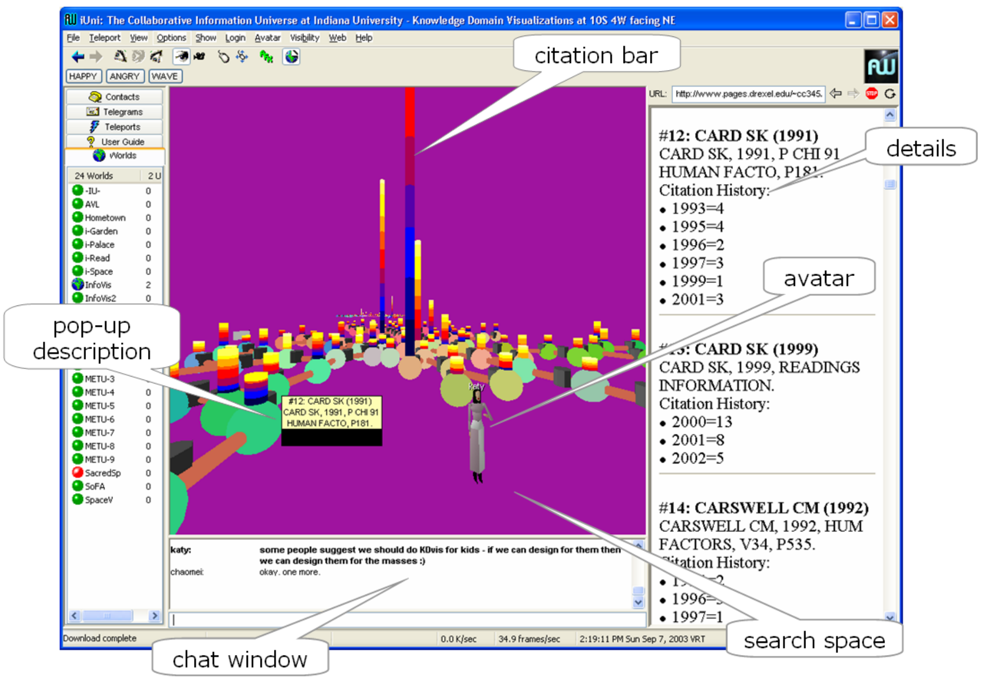

Figure 3. The user interface of the collaborative information visualization environment.

Users have the option to explore the virtual world via flying or walking in first- or third-person views; users can see their own avatars in the third-person view. Although spheres and cylinders in the original VRML version were clickable, the information was lost during the conversion to the RWX format for Active World. To amend this problem, a small black anchor was placed next to each sphere so that users may move their mouse cursor over the anchor to view a pop-up description, or click on it for the display of further details in the adjacent text document window.

5.2 Procedure

In order to evaluate the group dynamics measures and visualize them with tools introduced in sections 3 and 4, we conducted team-based information searches in two thematic CIVEs on research in information visualization in 1983-1992 and 1993–2003, respectively. Each subject first completed a pre-test questionnaire on simple demographics, including age, sex, experience using both three-dimensional virtual worlds and computers in general, and knowledge of the subject of information visualization. After a 5-minute practice task, each pair of subjects had 40 minutes to collaboratively examine the two visualization models in order to answer 10 questions. Feedback on search strategies and encountered problems was collected from a post-test questionnaire.

Participants

A total of eight subjects participated in the initial study, including five males and three females of an average age of 26.5 years (range of 22-40 years). All subjects were right handed. Six were native English speakers, one Chinese, and one Gujarati. The highest level of education included high school (1), college (4), graduate (2), and postgraduate (1). Major areas of study were information science and informatics (4), computer science (1), instructional systems technology (1), English (1), and history (1). The self-reported number of hours spent on a computer per week was 38.25 (range 16-80) and number of hours spent online per week was 29.12 (range 5-80).

Domain knowledge was self-rated on a 5-point scale (1-naïve, 5-expert) on ‘computers’ (3.75) and ‘information visualization’ (2.87). Two subjects took an information visualization course at Indiana University; one subject is an experienced user of Active World. Subjects had various experiences in related virtual environments: 7 used a text based chat environment (e.g., Yahoo, IM, IRL, MSN, Pow Wow, or Online Games), 3 used a three-dimensional single user virtual environment (e.g., CAVE, 3D Games), and 4 used three-dimensional multi-user virtual environment (e.g., AW, Online Games). When asked to identify their main problems with finding relevant information and/or documents, the subjects answered: identifying/finding correct query/index/phrase.

Equipment

The experiment was conducted in a PC computer lab at Indiana University. The lab is equipped with Pentium 4, 2GHz, 256MB RAM, Dell PCs running Windows 2000, and have 19-inch monitors that are connected via a 100MB per second connection to the host server of the virtual worlds. Subjects could only communicate to each other by chat in the virtual world. Each pair of subjects did not know each other before their session. They neither had a chance to talk about the tasks beforehand, nor had they performed group tasks in other CIVEs.

Tasks

Ten questions were given to the subjects. Subjects were instructed to search through the Data Set 1 (1983-1992) model for the first five questions, and the Data Set 2 (1993-2003) model for the second five questions. Subjects were instructed to answer the retrieval questions by filling in the identification numbers of corresponding documents.

<u>Q 1. – Q 5. Go to Data Set 1 (1983-1992)</u>

Q 1. Find Tufte, E. R. (1983) and Tufte, E. R. (1990). #: ___________ #: ________

Q 2. Find at least two entries of Bertin, J. #: ____________ #: ____________

Q 3. Find an entry on Exploratory Data Analysis. #: ________________

Q 4. Find Chrnoff, H. (1973). # ____________

Q 5. Open-end questions: Can you identify the common themes of documents colored in red, green, and blue?

Q 6. – Q 10. Go to Data Set 2 (1993-2003)

Q 6. Find Tufte's three books #: _______ #: _______ #: ______

Q 7. Find two most highly cited and recently published papers, one was published in 1999 and the other was published in 2000. #: ________ #: ________

Q 8. Open-end questions: Can you identify the common themes of documents colored in red, green, and blue?

Q 9. Open-end question: What was new in the second period?

Q 10. Open-end question: Can you identify any 1st-period topics that subsequently disappeared in the 2nd period?

The ten questions varied in terms of the level of difficulty and the amount of information to be gathered. The simplest questions required the subjects to identify Edward Tufte's three books. Because all his books have the highest citations in both periods they can be easily spotted in the virtual world. A more difficult question required the subjects to find at least two documents from another prominent figure, Jacques Bertin, but no other clues were given. A tougher question required the subjects to find two recently published and highly cited documents. The clue to this question was in the color patterns of their citation poles: their citation poles should be covered by bands of light colors because of the recent publication years, and the height of these poles should be the tallest among poles with similar color patterns. To answer this question, the subjects would have to compare the visual-spatial patterns carefully. An even tougher question was for the subjects to identify the nature of document groupings according to their sphere colors. To answer this question, the subjects would have to compare and examine documents of similar colors and identify patterns. The most difficult question was about the major differences between the two time periods. This question would entail the most extensive search and comparison between the two models at a global level.

5.3 Group Dynamics Analysis

In addition to the data collected from the pre- and post-test questionnaires as well as the answers to the ten questions, time stamped user positions, object clicks, and chat logs were also collected. The positions and sizes of all objects in the virtual worlds were also available. Based on those data sets, the spatial-semantic impact of a CIVE on group dynamics was examined.

Table 1 summarizes the performance scores in terms of task accuracy and movement measures of all four groups. Definitions are repeated in the table for convenience.

Table 1. Performance scores and group dynamics measures.

Group	Task Accuracy (Qs:1-4, 6-7) Excluding open-end	Total Distance Traveled by Individual	Total Distance Traveled by Individual	Average Total Distance Traveled	Group Diameter $\frac{\sum_{i,j}^{N} d_t(m_i, m_j)}{N \bullet T}$	Group-to-Target Distance $\frac{\sum_{t}^{T} \min_{i \in N, j \in O} d_t(m_i, o_j)}{T}$

	questions	$\max d_{m_i}$	$\min d_{m_i}$	$\frac{\sum_{i=1}^{N} d_{m_i}}{N}$		
1	0.36	68868.19	6579.76	37723.98	1746.78	5261.93
2	0.46	18072.38	11189.53	14630.96	3225.29	6081.48
3	0.47	29645.03	5555.89	17600.46	842.61	5397.32
4	0.52	22851.16	20929.50	21890.33	2821.77	6417.49
Mean	0.45	34859.19	11063.67	22961.43	2159.11	5789.56

Task Accuracy

Factual search questions were scored in terms of the percentage of relevant answers found in the total number of relevant target documents. The first part of the 10 questions mainly required factual search. Most groups found correct answers, although some groups achieved higher scores than others. The second half of the questions required teams to not only to find and compare documents, but also to identify what they have in common and derive a synthesized common theme. The difficulty was in part due to the intensity of the amount of information to be assessed and synthesized within a relatively short period of time. Another reason is related to the minimum amount of domain knowledge required to answer the open-end questions. Although the necessary information was available in the virtual worlds, it could still be very challenging, especially when all the subjects learned about the field was from just one course. Therefore, we were surprised when Group 4 came up with a very accurate and comprehensive answer, because the group neither had a particularly active chat session, nor an extensive travel record in the virtual worlds. Despite the fact that one member of the group took an information visualization course in the past, we suspect there may be other reasons to be discovered.

As seen in Table 1, Group 4 has the highest score; Group 3 the second; Group 2 the third; and Group 1 the lowest. Open-end questions were assessed and aggregated, but not scored.

Travel Distance

Group 1 has by far the largest group travel distance, the third largest group diameter, and the smallest group-to-target distance. In contrast, Group 2 has the least amount of combined travel distance, the largest group diameter, and the second largest group-to-target distance. According to these measures, Group 2 was loosely coupled semantically and spatially.

Group Diameter

Figure 4 shows the group diameter time series of Group 1 and Group 3 for easy comparison. Although Group 1 has the highest group travel distance, it has the second smallest overall group diameter (1746.78 meters), whereas Group 3 has the smallest among the four groups (842.61 meters).

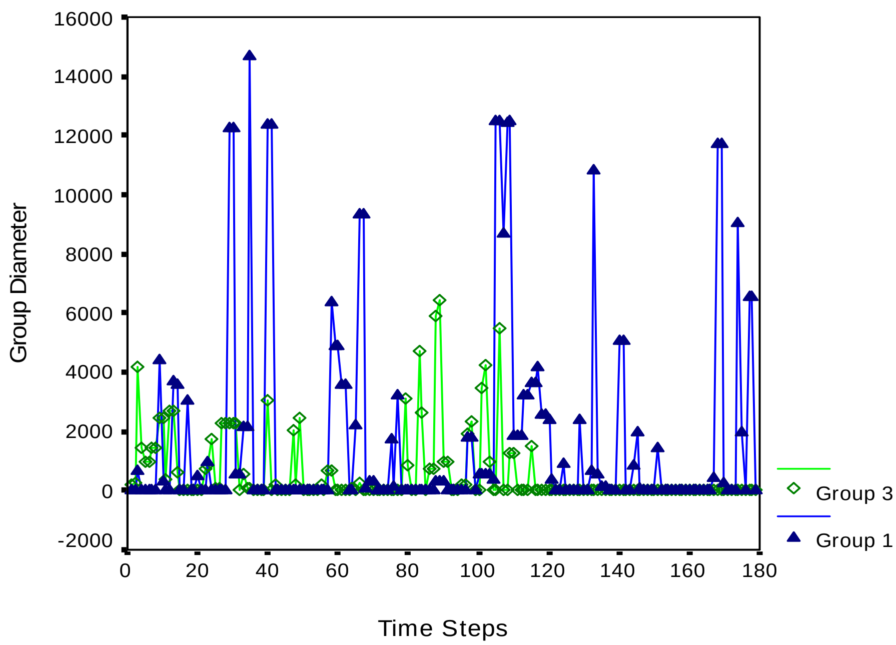

Figure 4. The group diameter time series of Group 1 and Group 3

Subjects reported their search strategies in the post-test questionnaire. For example, some subjects minimized their travel distance in the virtual world to avoid getting lost because of lack of experience in this type of navigation and lack of confidence in handling various navigation controls. Other subjects covered a large amount of travel distance in order to explore the virtual world.

Group-to-Target Distance

The group-to-target distance is a principal component of the semantic distance between a group and a visualization model. In the definition of semantic distance, the group-to-target distance is normalized by the average amount of distance traveled by the members of a group in their session. The smaller the group-to-target distance is, the stronger the semantic tightness of the group. Group 1 has the smallest group-to-target distance, and Group 2 has the second smallest measure. We incorporated the group-to-target distance sequence into a display of the sequence of chat utterances (Figure 5 and Figure 6). Substantial changes of such distances may be the sign of significant transition positions in chat sequences.

Group Coherence

Table 2 lists the spatial proximity, semantic coherence, and group coherence for all four groups. Remember that the coherence coefficients range between 0 and 1. The value of 1 indicates a tightly coupled group, whereas the value of 0 indicates a loosely coupled group. If two groups have the same level of spatial coherence, the one with a stronger semantic coherence will have a higher overall coherence.

Table 2. Three measures of group dynamics

Group	Spatial Proximity	Semantic Coherence	Group Coherence	Performance Score (0-1)

	(δ)	(σ)	(Σ)	1 - highest
1	0.954	0.861	0.646	0.36
2	0.780	0.584	0.487	0.46
3	0.952	0.693	0.565	0.57
4	0.871	0.707	0.563	0.82
Mean	0.889	0.711	0.565	0.55

Group 1 has the strongest group coherence (0.646), whereas Group 2 has the weakest (0.487). Group 3 and Group 4 have similar values. Group 1 has the highest spatial proximity, whereas Group 2 is the lowest. The difference between Group 3 and Group 4 in their spatial proximity was reduced in the overall coherence measure because they have very similar semantic coherence.

Group Coherence Space

The group coherence space shown in Figure 1 gives us a generic framework with which to identify the position of a particular group in the context of other groups performing similar tasks. The four groups in our study are all located in the upper left region because their spatial and semantic coherence measures are all greater than the dividing threshold of 0.5. The positions of the four groups are marked by the contour belts. Points on the same belt represent groups – hypothetical ones as well as actual ones - that have the same strengths of spatial-semantic coupling but might differ in their group travel distances and group diameters. For example, the particular group dynamics of Group 1 is just one among many possibilities – a different point on the contour line represents the same group coherence level as Group 1, which may come from the same group performing the same experiment, or even a different group. The diagram shows that it is possible for a group to demonstrate the same degree of coherence if it maintains a shorter traveling distance even if its group diameter is larger. Group 1 has the best coherence measures in this study followed by Group 4, 3, and 2.

5.4 Group Dynamics Visualizations

This section demonstrates the application of the social visualization techniques introduced in section 4 to the data acquired in the empirical study.

Semantic-Indentation Chat Sequence Display

The semantic distance of a group contains important clues for detecting the boundaries of individual episodes in a chat sequence. A smaller semantic distance does not necessarily mean that users must be engaging in a tighter coupling in terms of their group dynamics. More valuable implications of semantic proximity are related to the boundaries of episodes rather than the level of detail. For example, in Figure 5, the first significant increase of the group-to-target distance shown was linked to the utterance: “Should we concentrate on answering the rest of Data Set 2 questions for now?” The change in the semantic distance signifies the subjects were about to start another episode of collaborative search. Because they became more familiar with the structure of the search space, they could comfortably perform similar tasks further away from the visualization model.

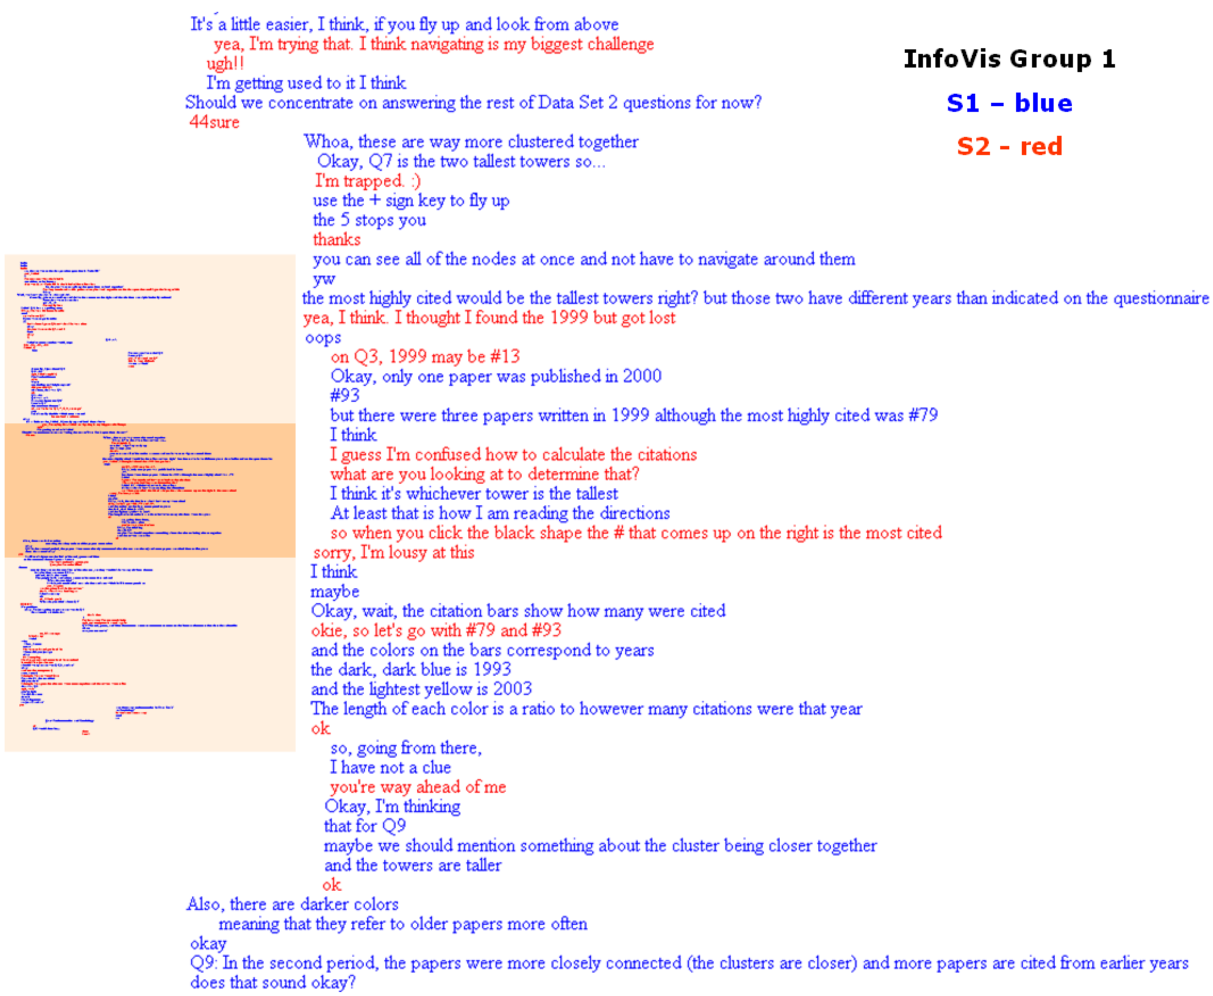

Figure 5. The chat sequence of Group 1 is indented in proportion to their group-to-target distance. Sudden changes in the group-to-target distance may be associated with potentially interesting transitions of the discourse. An overview map of the complete chat log is shown on the left. Highlighted is the enlarged text shown on the right.

Figure 6 shows another example of a semantically indented chat sequence from Group 3. The first major semantic distance increase was linked to a similar utterance of a suggestion: “[let’s look for] more Tufte.” The group subsequently switched their search to a different area.

Figure 6. The chat sequence of Group 3 is indented in proportional to their group-to-target distance. The significant jump corresponded to the search for the second half of the questions, marked by Subject 5: "more Tufte."

Activity Maps

In order to provide a visual summary of the history of a group exploring and interacting with the virtual worlds, so-called activity maps were generated that overlay user interaction data (positions, chat and click activity) onto a map of the virtual world. Figure 7 shows the activity maps of Groups 2, 3 and 4 run in this study. Larger-sized maps of Group 1 are shown in Figure 8.

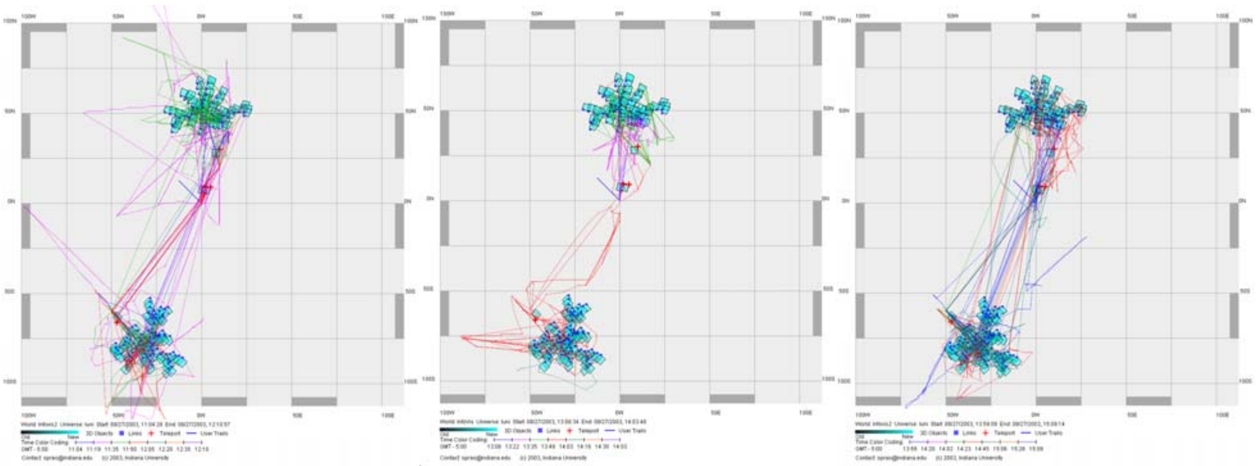

Figure 7. Group activity maps: Group 2 (left), Group 3 (middle), and Group 4 (right).

The two clusters of objects correspond to the two document visualization models for the two periods of time. Data set 1 is in the top middle, data set 2 in the lower left corner. Blue dots denote Web links to additional information about the documents. Red plus signs represent teleports. The colored lines represent time-coded user trails (See Figure 8). Straight lines indicate that users teleported using the four orientation and teleport signs placed in the world for convenient navigation between the two models. Different groups demonstrated not only distinct global patterns of navigation, but also diverse communication patterns.

Figure 8 shows detailed maps on the activity of Group 1. The map on the left shows locations in which users chatted (red triangles for user 1, green triangle for user 2) and the locations of anchor clicks (dark red dots), present when subjects retrieved detailed information on documents.

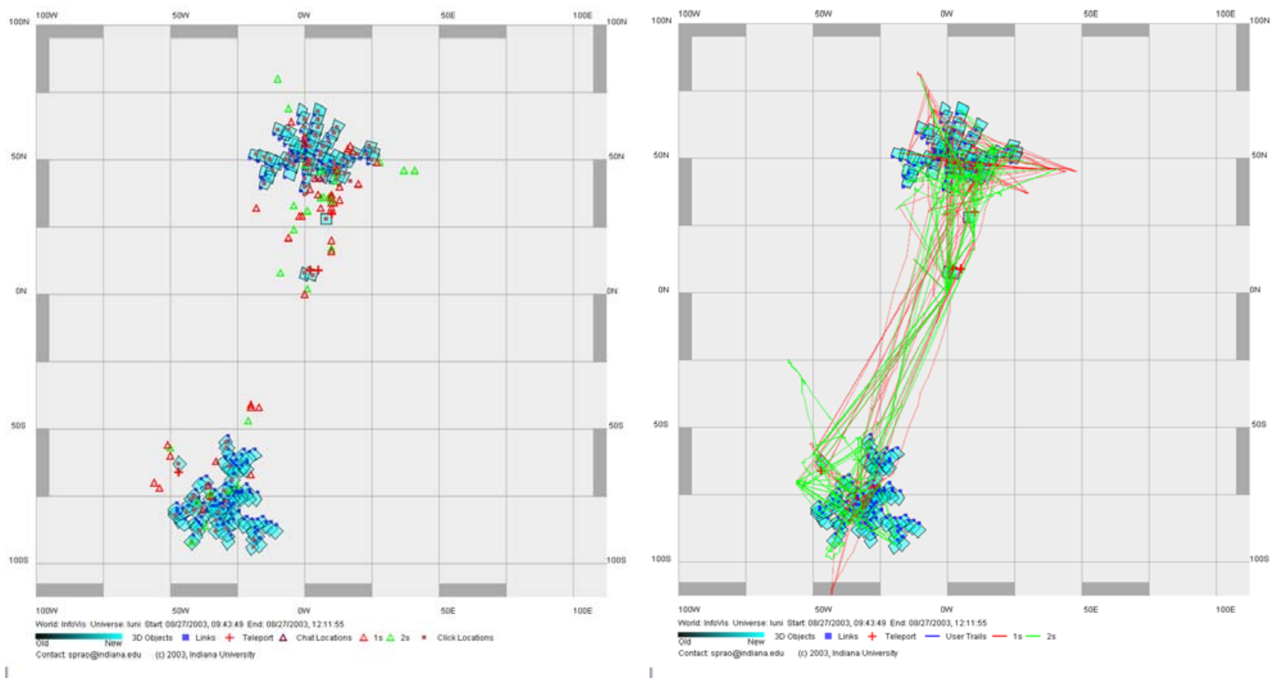

Figure 8. An activity map for Group 1 (left) and a trail map for Group 1 (right). User activity locations and click locations are overlaid on the spatial configuration of the virtual world.

Figure 8, right shows a larger-size trail map of Group 1. The trails of user one is given in red, the trail of the other user is shown in green. The predominant trails of this group connect the two clusters of documents.

Clock Face Maps

Clock-face maps visualize the temporal positions of an utterance in a given session. Clock faces show the time swept from the beginning of the session. Utterances within the same temporal neighborhood have similar patterns on their clock faces. Analysts can use clock-face maps to identify where and when users were particularly talkative with reference to spatial, semantic, and social aspects of their navigation. As shown in Figure 9, Group 3 worked on the upper cluster in the first part of the session, and then in the lower cluster towards the end of the session. The lines connecting utterances in the visualization highlight the utterances made within a short interval; in this study, linked utterances were made less than 25 seconds apart from each other.

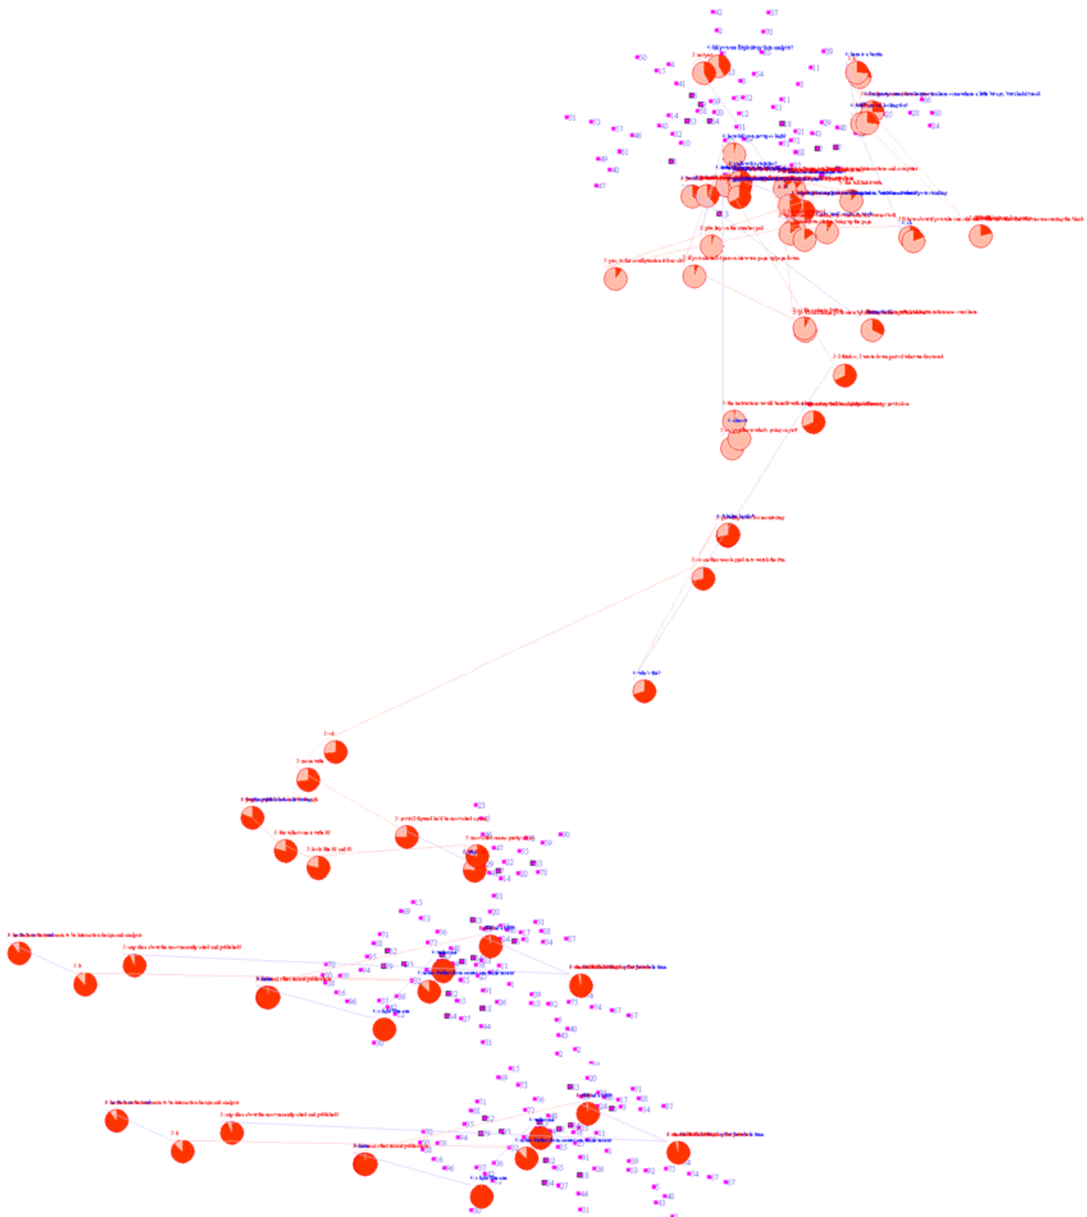

Figure 9. The clock-face visualization of the timing and the locations of individual utterances.

Figure 10 shows a more detailed analysis of the collaborative search session of Group 3. Significant episodes of the chat sequence were identified by extracting sub-sequences of utterances that were within each other's temporal vicinity. The shading of each cluster was added by hand. An episode typically began with a question or a suggestion, especially those associated with substantial changes of semantic distance. In addition, the temporal clusters of utterances may also identify an episode. We successfully identified some interesting episodes using the 25-second silence margin. For example, a short tutorial episode started with the question "How did you get up so high?", which was followed by a number of chat exchanges as one subject was learning from the other subject how to navigate in Active Worlds. The concentration of clock faces in Figure 10, just below the cluster of small rectangle objects, indicates that the subjects spent a considerable amount of time in the vicinity of the visualization model.

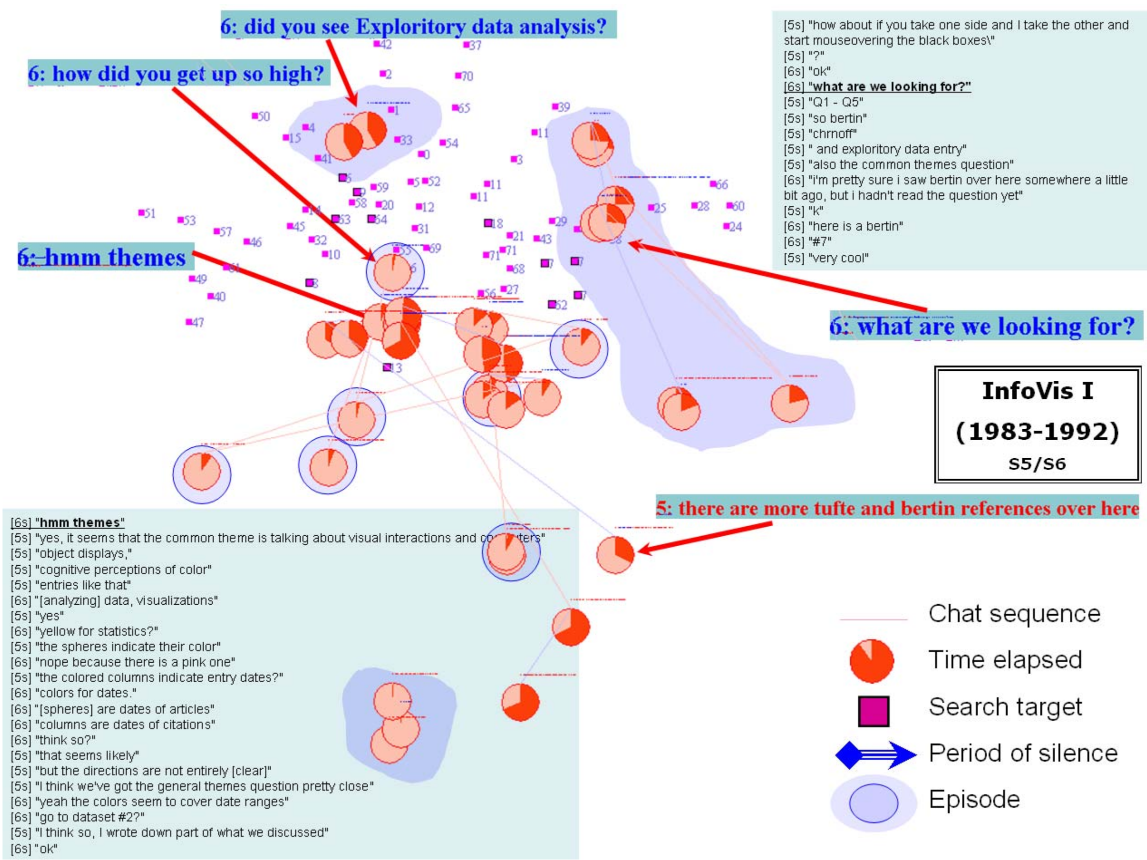

Figure 10. Episodes of collaborative search by Group 3.

We also explored the chat sequences using a longer silence margin and actually found a different type of episode. These episodes consisted of a substantial period of silence during which no chat communications were made; the 25-second silence margin would separate the utterances before and after the silence into two different episodes. However, we considered them as utterances from the single episode if we could infer from the utterances that subjects were still engaged the same task. Figure 11 is an example. During the silent interval, the subjects apparently went to search and resumed the chat only after they had found something. The "bridge" between the opening line of such an episode and the main body of the episode could be as long as 25 seconds as shown in the middle of Figure 11. The episode was opened with a question: "[Do you have] any idea about the most recently cited and published [documents]?" There was a 25-second silence in the chat window before the other subject responded: "No. I've been looking for that the whole time." Both subjects subsequently moved from their original positions to the cluster of objects where the target documents reside.

During the silent interval, the first subject did not even ask the other one to clarify the status of their collaboration mode. We are interested in how representative this case is. We might consider this phenomenon as a series of transitions of collaboration mode, from a tight mode, to a loose mode, and then back to a tight mode.

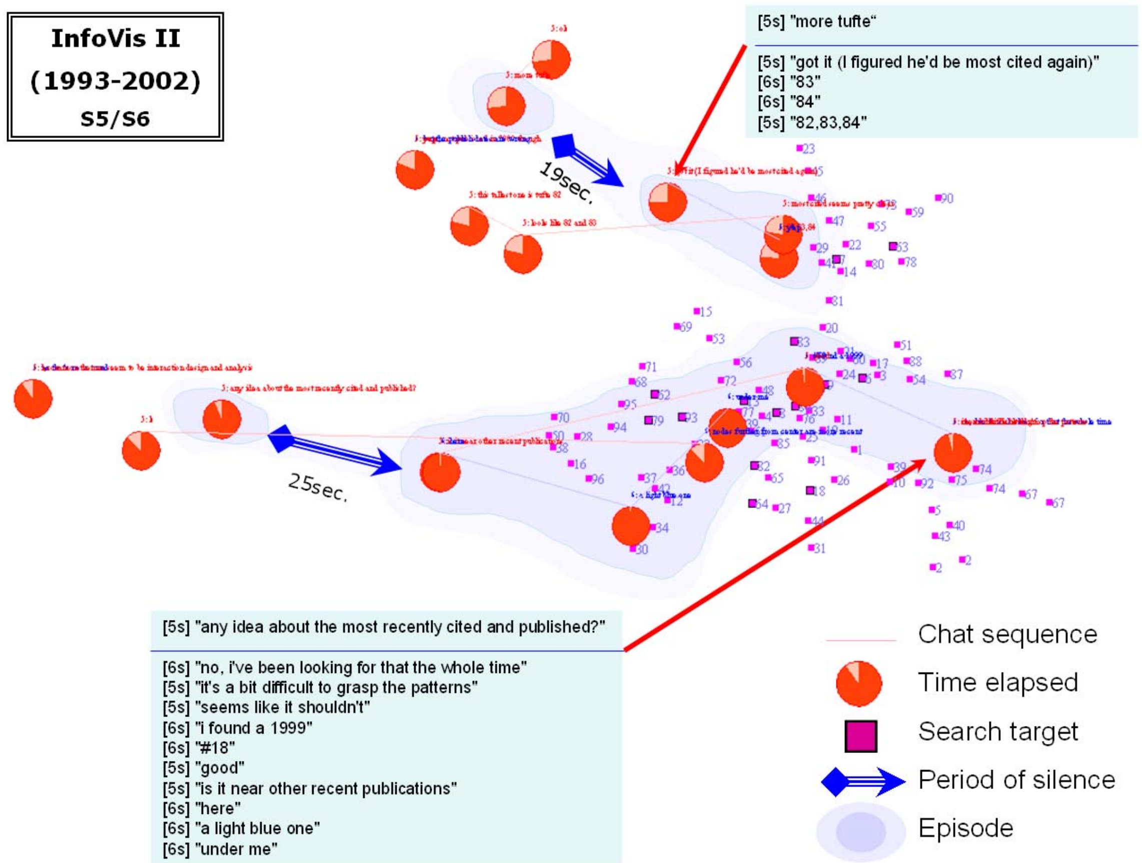

Figure 11. Episodes containing extended intervals of silence while group members being engaged in individual search.

5.5 Discussion

The empirical study generated a number of interesting results. The new group dynamics measures and visualization techniques introduced in this article provided the opportunity to cross-reference the spatial configuration of the virtual world with the chat, movement, and click activities of users.

The application of the group dynamics measures resulted in a number of observations. For example, users maintained their spatial proximity in order to avoid getting lost in the virtual world. When they did so, our quantitative approach would reveal a spatially tight group. In contrast, users who are familiar with virtual worlds may not be afraid of getting lost. As our analysis shows, instead of making moves back and forth within the virtual world, advanced users made many fewer moves than other users and took the full advantage of having instant access to the entire or almost entire set of documents via an overview.

In addition, the visualization techniques discussed in section 4 were applied to visualize and contrast group dynamics across teams and to study the impact of information structures intrinsic to a CIVE on collaborative behavior. Semantically indented chat sequences turn out to be a very encouraging technique for analysts to identify boundaries between episodes based on visual inspection. We expect this technique will be particularly useful when analyzing longer chat sequences. At the theoretical level, the fact that semantic distance provides a good identifier of transitions between episodes is particularly encouraging, suggesting that spatial, semantic, and social navigation can be understood and quantitatively monitored as an integral part of group dynamics in a CIVE.

The clock-face maps also provide an interesting way to make sense of group dynamics in terms of the interrelationships between spatial, semantic, and social dimensions. For example, the time-stamp on a clock face can help us to track where and when users moved from one place to another and what they said. A commonality between semantically indented chat sequences and clock-face maps is that both are designed to reveal underlying connections between users' navigational movements and the contents of their discourse in a situated manner. Therefore, we recommend that this quantitative methodology should be considered in the study of group dynamics in a wider range of CIVEs, especially where the interrelationship between spatial and semantic attributes is concerned. In addition to the node-and-link type of information visualizations, further studies should investigate alternative designs such as relief maps, self-organizing maps, and other innovative visualizations.

The analysis of the post-test questionnaires provided additional insight into the information foraging and collaboration strategies subjects used. Some used a divide-and-conquer strategy whereby group members searched in their part of the virtual world and then merged what they found. Others said that they did not use any strategy. Most subjects found the search interesting: "This is a very interesting way to look for good papers to read," said one subject. Another one liked it because it was like playing a treasure hunt game. Subjects also found that they could fly up and have a bird's eye view of the clusters. Comments from the post-test questionnaires are valuable for the development of a deeper understanding of situated activities. Qualitative and quantitative approaches together provide the most comprehensive and convincing accounts of users' experiences.

The initial results are particularly encouraging as we move towards the establishment of a generic framework for the study of the spatial-semantic impact of CIVEs on group dynamics. Our ultimate goal is to understand not only how people navigate in (virtual) worlds, but also how we may draw insightful feedback from quantitative measures of group dynamics as they interact with CIVEs so that we can specifically tailor the design of CIVEs and make it a better place for future visitors.

There are diverse avenues to improve collaborative information visualizations. For example, all the citation details of documents in the experiment were available from a single HTML document. The chat log implies that some subjects might have taken advantage of this design; instead of conducting a visual search in the spatial interface, one could answer at least the first few questions via text search in the HTML page. In a new design, each bibliographic record will be isolated from the rest of records to ensure that visual search is used.

The sample size of eight individual subjects in four pairs should be increased in subsequent experiments. Currently, we cannot draw any connections between demographic characteristics of subjects at this sample size, although we have paid special attention to factors that may distinguish one subject's performance and behavior from the other. A larger sample size, repeated sessions, and interviews are among the possible ways to consolidate the methodology. There are some factors that one may control more explicitly in an experimental study, including the size of a group, the length of collaboration, and the level of prior experience.

The intensity of person-person communication during a chat session, or more broadly, in any communicative sessions, would be an additional aspect of group coherence; possible measures include the number of words per minute, the number of lines per minute, and the gaps between consecutive utterances. This particular possibility and other options are worth consideration in refining the concepts of spatial and semantic coherence.

Some interesting feedback also drew our attention to the overall design of the virtual environment, namely, the general support for information search and more specific supports for carrying out in-depth analysis of patterns and trends. In the former, examples include making the clickable anchor objects in the scene larger and/or higher so that one can easily select them and click for further details. In addition to being able to click objects in a visualization model and examine further details outside the spatial interface, subjects also expressed the desire to be able to click the HTML back into the visualization model. In other words, both spatial and textual interfaces need to be integrated more tightly.

The design of the CIVE itself can be improved so that more specific hypotheses can be tested. For example, in the empirical study described in this article, one of the assumptions is that users searching for documents are expected to move closer to potentially *profitable* search areas, concurrent with findings in earlier studies (Chen et al., 2002). Research in information foraging has generated many results regarding relevant issues (Pirolli & Card., 1995). However, the current visualization models merely reward users with rather limited advantages when they move closer to document spheres in the three-dimensional structure – due to the small size of the black anchor of a document, the closer we are close to the object, the easier it is to click on it. Although clicking on black anchors does bring further details to the user, the direct advantage for an experienced CVE user becomes marginal; this in part explains why experienced users did not move very close to their targets. In follow-up experiments, one can increase the degree of reward for users who move close enough to their targets by, for example, showing increasingly more important information as the distance decreases. Once a clear and distinct rewarding mechanism is in place, users would be more strongly motivated to make moves to the vicinity of their targets. It is reasonable to expect that such mechanisms will considerably improve the accuracy of spatial-semantic coupling measures.

A tougher challenge is to facilitate users to perform tasks involving in-depth analysis beyond an individual document. More challenging and potentially more rewarding facilities would be a search mechanism that integrates visual-spatial properties with semantic properties of the underlying information space. For example, users may formulate queries such as "give me a list of all documents on data layout shown in red" or "show me where the documents on data layout in red from the first period are located in the second model." In the current experimental setting, these facilities are not available to users, which in part explained the difficulties experienced by the subjects in searching for answers to the last few questions. This experiment has provided valuable input concerning how to further investigate and improve the match between user tasks and facilities available in the virtual environment. This is also a necessary step to incorporate information visualization into collaborative virtual environments in practical applications such as distance learning, digital libraries, and online scientific communities.

6. Conclusion

The major contributions of this study are as follows: 1) the establishment of the first integrative methodology for quantitatively analyzing group coherence with special reference to spatial, semantic, and social navigation; 2) the development of a set of visualization and analysis tools that can be made widely available; and 3) empirical findings from performing potentially realistic tasks, namely, searching for information in domain visualizations.

The generic measures of group dynamics and visualization techniques introduced and exemplified in this article are applicable to a wide range of CIVEs and other virtual or real world environments with explicit spatial and semantic properties and opportunities for collaboration. The ability to detect significant transition points in a discourse is also an important contribution to the toolkit of analysts. They are an important step towards the establishment of analytical and descriptive methods for understanding group dynamics. These methods will considerably strengthen the connection between information visualization and CVEs in terms of cross-domain collaboration and applications.

We expect that the findings and experiences of this study are constructive and informative and that, as an interesting exemplar, they will stimulate further studies in this area of interdisciplinary nature, and strengthen the connections between individual fields involved - information visualization, collaborative virtual environments, and knowledge domain visualization.

Acknowledgements

The authors wish to thank Shashikant Penumarthy of Indiana University for his help in the collection of user activity data and the generation of the activity maps. Thanks also to the eight participants in the experiment and for their valuable feedback. This research is supported by Indiana University's High Performance Network Applications Program and an academic equipment grant by Sun Microsystems.

References

Benford, S., Snowdon, D., Colebourne, A., O'Brien, J., & Rodden, T. (1997). *Informing the design of collaborative virtual environments.* Paper presented at the the international ACM SIGGROUP conference on Supporting group work: the integration challenge (GROUP '97).

Benford, S., Snowdon, D., Greenhalgh, C., Ingram, R., Knox, I., & Brown, C. (1995). VR-VIBE: A virtual environment for co-operative information retrieval. *Computer Graphics Forum, 14*(3), C.349-C.360.

Börner, K. (2002). Twin Worlds: Augmenting, evaluating, and studying three-dimensional digital cities and their evolving communities. In M. Tanabe & P. v. d. Besselaar & T. Ishida (Eds.), *Digital Cities II: Computational and Sociological Approaches* (pp. 256-269): Springer Verlag.

Börner, K., Chen, C., & Boyack, K. (2003). Visualizing Knowledge Domains. In B. Cronin (Ed.), *Annual Review of Information Science & Technology* (Vol. 37, pp. 179-255). Medford, NJ: Information Today, Inc./American Society for Information Science and Technology.

Börner, K., Hazlewood, R., & Lin, S.-M. (2002). *Visualizing the spatial and temporal distribution of user Interaction data collected in three-dimensional virtual worlds.* Paper presented at the Proceedings of the 6th International Conference on Information Visualization, London, England.

Börner, K., Lee, G. J., Penumarthy, S., & Jones, R. J. (Forthcoming, January 18-22, 2004). *Visualizing the VLearn3D 2002 conference in space and time.* Paper presented at the SPIE Conference on Visualization and Data Analysis, San Jose, CA, USA.

Börner, K., & Penumarthy, S. (2003). Social Diffusion Patterns in Three-Dimensional Virtual Worlds. *Information Visualization Journal, 2*(3), 182-198.

Börner, K., Penumarthy, S., DeVarco, B. J., & Kerney, C. (Forthcoming, September 19-21, 2003). *Visualizing social patterns in virtual environments on a local and global scale.* Paper presented at the Digital Cities 3: Local information and communication infrastructures: Experiences and challenges, Amsterdam.

Chen, C. (1999a). *Information Visualisation and Virtual Environments*. London: Springer.

Chen, C. (1999b). Visualising semantic spaces and author co-citation networks in digital libraries. *Information Processing and Management, 35*(2), 401-420.

Chen, C. (2003). *Mapping Scientific Frontiers: The Quest for Knowledge Visualization*. London: Springer.

Chen, C., Cribbin, T., Kuljis, J., & Macredie, R. (2002). Footprints of information foragers: Behaviour semantics of visual exploration. *International Journal of Human-Computer Studies, 57*(2), 139-163.

Chen, C., & Paul, R. J. (2001). Visualizing a knowledge domain's intellectual structure. *Computer, 34*(3), 65-71.

Chen, C., Thomas, L., Cole, J., & Chennawasin, C. (1999). Representing the semantics of virtual spaces. *IEEE Multimedia, 6*(2), 54-63.

Conklin, J. (1987). Hypertext: An introduction and survey. *Computer, 20*(9), 17-41.

Crossley, M., Davies, J., McGrath, A., & Rejman-Greene, M. (1999). The knowledge garden. *BT Technology Journal, 17*(1).

Darken, R. P., & Sibert, J. L. (1996). *Wayfinding strategies and behaviors in large virtual worlds.* Paper presented at the CHI '96, Vancouver, BC.

Dieberger, A. (1997). Supporting social navigation on the World-Wide Web. *International Journal of Human-Computer Studies, 46*, 805-825.

Dourish, P., & Chalmers, M. (1994). *Running out of space: Models of information navigation.* Paper presented at the HCI '94.

Erickson, T. (1993). *From interface to interplace: The spatial environment as a medium for interaction.* Paper presented at the the European Conference on Spatial Information Theory.

Fry, B. (2003). *anemone*. Benjamin Fry. Retrieved Sept 8, 2003, 2003, from the World Wide Web: http://acg.media.mit.edu/people/fry/anemone/index.html

Greenhalgh, C. (1997, Setp 7-11, 1997). *Analysing movement and world transitions in virtual reality tele-conferencing.* Paper presented at the The 5th European Conference on Computer Supported Cooperative Work (ECSCW'97), Lancaster, UK.

Greenhalgh, C., & Benford, S. (1995). MASSIVE: A virtual reality system for tele-conferencing. *ACM Transactions on Computer Human Interaction, 2*(3), 239-261.

Hanson, A., Wernert, E., & Hughes, S. (1997). *Constrained navigation environments.* Paper presented at the Scientific Visualization: Dagstuhl '97.

Harrison, S., & Dourish, P. (1996). *Re-place-ing space: The roles of place and space in collaborative systems.* Paper presented at the CSCW '96.

Herman, I., Melançon, G., & Marshall, M. S. (2000). Graph visualization and navigation in information visualization: a survey. *IEEE Transactions on Visualization and Computer Graphics*(1), 24-44.

Hill, W. C., Hollan, J. D., Wroblewski, D., & McCandless, T. (1992, May 3-7, 1992). *Edit wear and read wear.* Paper presented at the CHI'92 Conference on Human Factors in Computing Systems, Monterey, CA.

Jeffrey, P., & Mark, G. (1998, 16-17 March). *Constructing social spaces in virtual environments: A study of navigation and interaction.* Paper presented at the Workshop on Personalised and Social Navigation in Information Space, Stockholm.

Johnson, B., & Shneiderman, B. (1991, October 1991). *Tree-maps: A space filling approach to the visualization of hierarchical information structures.* Paper presented at the IEEE Visualization 91.

Jolliffe, I. T. (1986). *Principal Component Analysis*. New York: SpringerVerlag.

Kraut, R., Egido, C., & Galegher, J. (1988, 1-12 September 1988). *Patterns of contact and communications in scientific research collaboration.* Paper presented at the CSCW '88.

Lamping, J., Rao, R., & Pirolli, P. (1995). *A focus+context technique based on hyperbolic geometry for visualizing large hierarchies.* Paper presented at the ACM CHI'95 Conference on Human Factors in Computing Systems.

Olson, J., Card, S., Landauer, T., & Olson, G. (1993). Computer-supported cooperative work: Research issues for the 90s. *Behaviour and Information Technology, 12*(2), 115-129.

Pirolli, P., & Card., S. K. (1995). *Information foraging in information access environments.* Paper presented at the CHI '95, Denver, CO.

Robertson, G., Card, S., & Mackinlay, J. (1993). Information visualization using 3D interactive animation. *Communications of the ACM, 36*(4), 57-71.

Robertson, G. G., Mackinlay, J. D., & Card, S. K. (1991, April 28 - May 2, 1991). *Cone trees: Animated 3D visualizations of hierarchical information.* Paper presented at the CHI '91, New Orleans, LA.

Small, H. (1999). Visualizing science by citation mapping. *Journal of the American Society for Information Science, 50*(9), 799-813.

Thurstone, L. L. (1931). Multiple factor analysis. *Psychological Review, 38*, 406-427.

Tromp, J., Steed, A., Frecon, E., Bullock, A., Sadagic, A., & Slater, M. (1998). Small group behavior experiments in the Coven project. *IEEE Computer Graphics and Applications, 18*(6), 53-63.

Tromp, J. G., Steed, A., & Wilson, J. R. (2003). Systematic usability evaluation and design issues for collaborative virtual environments. *Presence: Teleoperators and Virtual Environments, 12*(3), 241 - 267.

Wexelblat, A. (1999). *History-based tools for navigation.* Paper presented at the Proceedings of the 32nd Annual Hawaii International Conference on Systems Sciences, Hawaii.

Wexelblat, A., & Maes, P. (1997). *Visualizing histories for Web browsing.* Paper presented at the RIAO'97: Computer-Assisted Information Retrieval on the Internet, Montreal.

White, H. D., & McCain, K. W. (1998). Visualizing a discipline: An author co-citation analysis of information science, 1972-1995. *Journal of the American Society for Information Science, 49*(4), 327-356.

Figure Captions

(Preferred Size: 2 column wide, except Figure 4)

Figure 12. Group Coherence Space, colored by the different measures, showing the positions of four groups in the subsequently explained empirical study. The contour belts in the upper right quadrant, which indicate regions of similar strengths of spatial-semantic coupling, are discussed in section 5.3.

Figure 13. A visualization model of citation patterns.

Figure 14. The user interface of the collaborative information visualization environment.

Figure 15. The group diameter time series of Group 1 and Group 3

Figure 16. The chat sequence of Group 1 is indented in proportion to their group-to-target distance. Sudden changes in the group-to-target distance may be associated with potentially interesting

transitions of the discourse. An overview map of the complete chat log is shown on the left. Highlighted is the enlarged text shown on the right.

Figure 17. The chat sequence of Group 3 is indented in proportional to their group-to-target distance. The significant jump corresponded to the search for the second half of the questions, marked by Subject 5: "more Tufte."

Figure 18. Group activity maps: Group 2 (left), Group 3 (middle), and Group 4 (right).

Figure 19. An activity map for Group 1 (left) and a trail map for Group 1 (right). User activity locations and click locations are overlaid on the spatial configuration of the virtual world.

Figure 20. The clock-face visualization of the timing and the locations of individual utterances.

Figure 21. Episodes of collaborative search by Group 3.

Figure 22. Episodes containing extended intervals of silence while group members being engaged in individual search.

Table Captions

Table 3. Performance scores and group dynamics measures.

Table 4. Three measures of group dynamics