From Latent Semantics to Spatial Hypertext — An Integrated Approach

Chaomei Chen
Department of Information Systems and Computing
Brunel University
Uxbridge UB8 3PH, UK
Tel: 44-1895-203080
E-mail: chaomei.chen@brunel.ac.uk

Mary Czerwinski
Microsoft User Interface Research
One Microsoft Way
9N/2290
Redmond, WA98052, USA
Tel: 1-425-703-4882
E-mail: marycz@microsoft.com

ABSTRACT
In this paper, we introduce an integrated approach to the development of spatial hypertext. This approach brings together several theories and techniques concerning semantic structures, and streamlines the transformation from implicit semantic structures to a semantic space rendered in virtual reality. Browsing and querying become natural, inherent, and compatible activities within the same semantic space. The overall design principle is based on the theory of cognitive maps. Techniques such as latent semantic indexing, Pathfinder network scaling, and virtual reality modelling are used in harmony. The value of this integrated approach is discussed based on initial results of a recent empirical study, which suggests that the spatial metaphor is intuitive and particularly useful when dealing with implicit information structures, or when a highly flexible and extensible virtual environment is required. Search strategies in association with the spatial hypertext and further work are also discussed.

KEYWORDS: Spatial hypertext, latent semantic indexing, virtual reality, digital libraries

INTRODUCTION

Generating flexible and extensible hypertext systems is a challenging task as one must capture and track implicit and emergent structures in an evolving environment [13, 22, 28]. There has been a rapidly growing interest in approaches such as open hypermedia systems (OHS), in which dynamic node-link binding strategies are often used in order to achieve desired flexibility and maintainability (e.g., [14]).

The notion of spatial hypertext relies on implicit structures that can be derived from how text is spatially organised by people. Marshall and Shipman [20] used the term *linkless structure* to describe how people use spatial layout to represent implicit structures in three different hypertext systems. They argued that the ability to find and use implicit structures is important to users in spatial hypertext. In this paper, spatial hypertext refers to hypertext systems that maintain a strong connection between spatial representations of information and underlying semantic structures.

In our previous work, we have developed the *Generalised Similarity Analysis* (GSA) framework, which integrates several structuring mechanisms for generating virtual hypertext link structures [5]. In this paper, we will describe how these structures can be transformed into spatial hypertext. We will explore the notion of virtual reality-enabled spatial hypertext and whether browsing and querying can both fit into the same semantic space naturally.

This paper is organised as follows. First, we will describe the context of the work and introduce techniques to be used in our approach, especially Latent Semantic Indexing (LSI) [10] and Pathfinder network scaling [27]. Second, we will introduce the theory of cognitive maps and its role in our subsequent virtual reality modelling. We will discuss how these theories and techniques can be integrated in order to resolve a number of challenging issues. For example, what information is needed for browsing and querying intuitively and seamlessly within the same spatial hypertext? We will also briefly describe some empirical findings concerning the spatial user interface. Finally, we will discuss the implications of this approach for the design of hypertext systems.

SPATIAL ORGANISATION OF INFORMATION

In this paper, we introduce an integrated approach that addresses a series of interrelated issues concerning the generation of spatial hypertext, including the nature of cognitive maps, extracting patterns in proximity data, and incorporating these patterns into a visual user interface. Although most of these issues have been explored separately from various perspectives, few studies have examined them collectively in a larger context. To our knowledge, these techniques have not been streamlined to form a comprehensive and far-reaching approach.

Information visualisation techniques, such as Fisheye Views

HyperText 98 Pittsburgh PA USA

[11] and Cone Trees [24], are designed to help users to access information at various levels based on visualised structures. However, most visualisation techniques are designed to deal with existing and explicit structures, such as hierarchical structures that already exist in a file system. The focus of our work is, instead, on implicit and emergent structures.

A spatial metaphor is often used in association with the notion of a context, which is typically implemented as nodes and links in hypertext systems. Spatial hypertext, such as VIKI [21], has been regarded as a unique design paradigm. Marshall and Shipman [20] studied how people used spatial layout to imply structures in three different spatial hypertext systems, including VIKI. They noted that a spatial metaphor allowed authors to convey volatile, implicit, and emergent patterns in a document space, and that users were able to interpret implicit structures based on spatial proximity and perceptual conventions.

Orendorf and Kacmar [23] described a spatial approach to organising digital libraries by taking advantage of an existing geographical layout. However, a straightforward layout may not always be available or appropriate for generic information visualisation (see also [19]). Structuring abstract digital documents in general presents a challenging issue, which has been the focus of our work.

In addition to the approaches described above, implicit structures can be extracted and derived from other sources of heuristics. Most of these approaches fall into two categories. Approaches in the first category simply take advantage of pre-defined association rules, such as SemNet [12], which used existing rules defined in a knowledge base. Approaches in the second category derive new association rules, such as BEAD [3], which was based on content analysis. Our approach belongs to the latter category. In fact, our GSA framework unifies a number of ways of inferring and visualising implicit structures [5].

Visualising complex information structures is much more difficult than representing regular hierarchical structures. Zizi and Beaudouin-Lafon described the design of a document retrieval system, called SHADOCS, in which interactive dynamic maps were incorporated into the user interface [32]. SHADOCS used a dynamic clustering technique to divide a large set of document descriptors into smaller clusters. Graphical overview maps were subsequently generated on the screen using a space-filling algorithm. Each region in a map corresponds to a cluster of descriptors. The size of a region is proportional to the relative importance of those descriptors in the underlying documents. However, some significant and challenging issues are yet to be fully addressed, for example, the scalability issue (e.g., [16]).

There are two types of approaches towards the issue of scalability, focusing on either the size of the network (in terms of the number of nodes) or the density of the network (in terms of the number of links). SHADOCS has convincingly demonstrated that a large network can be separated into a number of smaller networks by dynamic clustering algorithms. The scalability issue, in terms of the size of new networks that one must deal with, has been largely resolved [32]. However, a density-related scalability issue turns out to be more difficult. The total number of links in a network consisting of n nodes could be as many as n^2. A commonly used strategy is to set a threshold value and only consider links whose values are above the threshold. For example, SHADOCS used a straightforward threshold to control the number of links to be displayed on the screen map. Since the spatial relations have not been taken into account, the linkage in a pruned network may look rather arbitrary and incompatible with the spatial layout. After all, scalability implies the ability to maintain the original integrity, consistency, and semantics associated with the network representation of an implicit structure. In this paper, we intend to show that Pathfinder network scaling algorithms provide a useful means of dealing with this challenging problem in a more harmonious way.

Pathfinder algorithms were originally designed to help psychologists to extract structural patterns from proximity data estimated by human beings [27]. Pathfinder network scaling can be seen as a link reduction mechanism that preserves the most salient semantic relations. A key assumption is the triangular inequality condition and only those links that satisfy this condition will appear in the final network. In essence, the rationale is that, if the meaning of a semantic relation can be more accurately or reliably derived from other relations, then this particular relation becomes redundant and therefore it can be safely omitted. In the GSA framework, we extend this notion in order to handle proximity data estimated by statistical and mathematical models [5]. A distinct advantage is that the same spatial metaphor can be consistently used across a range of proximity data. This is a significant advantage for maintaining the integrity of the semantic structures generated by different theories and techniques.

It is clear now that we need some efficient statistical and mathematical models in order to generate a high-quality semantic structure. LSI work was originated in automated semantic indexing for information retrieval [10]. In our previous work [5], we followed the classic $tf \times idf$ vector space model [25, 26], whereas in this study LSI was used instead (see Figure 1). We will explain the reasons of this change shortly.

In a broader sense, there are some similarities between our approach and self-organised feature maps produced by artificial neural network techniques (e.g., [18]). The major difference between our approach and neural network-based approaches lies in the way that the network structure is derived and represented. Comparing the two approaches more closely is certainly an interesting area of further research.

The significance of our approach is also due to the integral role of user interface design and empirical evaluation. For example, how can we use virtual reality techniques to widen and enrich the ways that users can access complex structures? Chimera and Shneiderman [6] studied the use of three interfaces for browsing a large table of contents. Their empirical work has highlighted the significance of the ability to explore high-level information continuously and to access lower-level information easily on demand when dealing with a large amount of information. In this paper, we will describe the design of a virtual reality-based user interface in order to provide users a sense of continuity and direct manipulation as well as traditional link following features.

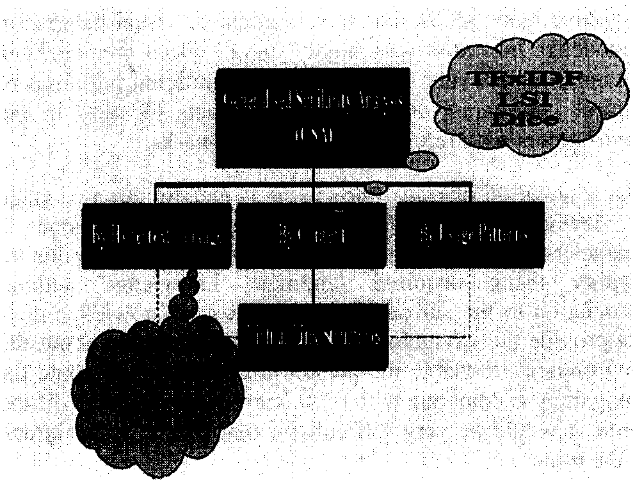

Figure 1. An integrated approach and its components.

In our previous work, virtual hypertext link structures were derived from interrelationships among documents. For example, content-based similarities were computed using the classic *tf × idf* vector space model in information retrieval [25, 26]. However, this model relies on an assumption that terms in document vectors are independent. It has been noted that the validity of this assumption may need to be re-examined in a wider context [e.g., 10]. Therefore, in this paper, we estimate document similarities based on LSI instead. The shift from the classic *tf × idf* model to LSI not only underlines the flexibility and extensibility of the GSA framework, but also highlights the need of comparative studies in related areas. We will explain the significance of using LSI and Pathfinder networks in the following sections.

LATENT SEMANTIC INDEXING AND PATHFINDER

LSI was designed to overcome the so-called *vocabulary mismatch* problem faced by information retrieval systems [10]. It has been shown that the relevance of documents to a given topic may not be reliably estimated based on individual words used in these documents. LSI is based on an assumption that there exists an underlying semantic structure in association with a collection of documents, but this structure is partially obscured by the vocabulary mismatch problem. The role of LSI, therefore, is to uncover the latent semantic structure by taking into account various evidence that some words may be in fact related by a unique concept.

In LSI, a semantic space is constructed based on a large *term×document* matrix. Each cell in the matrix represents the number of times a term occurring in a particular document. LSI relies on a mathematical technique called *singular value decomposition* (SVD) in order to approximate the original matrix with a smaller, truncated SVD matrix. A proper truncation usually results in a more concise and accurate model of the underlying semantic structure.

Perhaps the most compelling claim from the LSI work has been that it allows an information retrieval system to retrieve documents that share no words with the original query [10]. Another potentially appealing feature is that the underlying semantic space can be mapped into a geometric space for visualisation purposes. For example, one can project the semantic space into an Euclidean space for a 2D or 3D visualisation (Figure 2). However, large complex semantic spaces in practice may not always fit into low-dimension spaces comfortably.

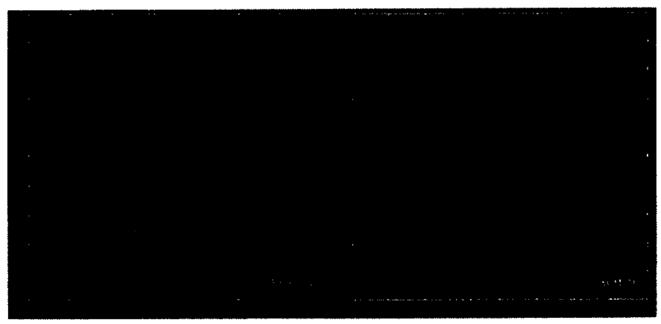

Figure 2. Scatter plots of CHI (left) and the CACM collection (right).

Pathfinder network scaling [27] shares some similarities with widely used methods such as Multidimensional Scaling (MDS) [17]. These methods are often used to map an abstract or high-dimensional space into a 2-dimension or 3-dimension Euclidean space.

Unlike MDS, Pathfinder network scaling takes a special step to safeguard the integrity in the abstract space. The direct outcome of this safeguard step is a network with a reduced number of links. More precisely, Pathfinder network scaling relies on the triangular inequality condition described earlier, which specifies a link selection rule based on weights associated with the links. Pathfinder network scaling removes links that do not satisfy the triangular inequality condition from the final network representation. The rationale is that the underlying structure is likely to be better captured and communicated in this way.

The spatial layout of a Pathfinder network is determined by a force-directed graph drawing algorithm [15]. Such graph drawing techniques are becoming increasingly popular in information visualisation due to their simplicity and intuitive appeal [5, 32]. In the next section, we will describe how LSI and Pathfinder network scaling fit into an integrated methodology.

INTEGRATED APPROACH

The focus of our work is on an integrated, iterative design framework that allows us to extract and represent latent semantic structures and transform these structures into spatial hypertext. An integrated approach brings together several fundamentally related theories and techniques in order to achieve a series of transformation and yet preserve the most salient semantic structures. The need of an integrated and highly automated approach becomes clear when one must capture implicit structures from a large collection of documents, such as a digital library or other sources of information. In this section, we will explain how these theories and techniques are connected together.

LSI was used to generate a *document* × *document* similarity matrix based on the title, author names and the abstract of each document. A standard set of common English words, known as *stopwords* in the information retrieval literature, were identified and excluded from subsequent indexing and analyses. Document vectors in LSI were generated using the logarithm of term-document occurrences as local weightings and the entropy as global weighting. This is a recommended choice [10].

We then generated Pathfinder networks by imposing the triangular inequality condition throughout the entire network in order to minimise the total number of links in the final networks. If the number of links in the resultant network is still too large, a Minimum Spanning Tree (MST) option is supported as an alternative in our software based on [31]. However, a Pathfinder network has a very desirable feature — the structural representation is unique in that a Pathfinder network is the set union of all the possible MSTs.

Finally, the result of force-directed graph drawing of the network was automatically transformed into virtual reality models in Virtual Reality Modeling Language (VRML).

We have applied this approach to the latest three ACM SIGCHI conference proceedings on Computer-Human Interaction (CHI) and the ACM *Hypertext Compendium*™ (HTC) [1]. The CHI collection includes 169 papers from CHI'95, CHI'96 and CHI'97. The HTC collection includes 128 papers and panels from conference Hypertext'87, Hypertext'89, ECHT'90 and other sources.

In addition to semantic structures derived separately from individual sources of documents, a combined semantic structure can be derived from a merged document collection, which may consist of several unique collections. An integrative view can be used for information filtering or managing personalised digital libraries. Researchers can simulate and see how queries, relevant documents and retrieved documents are located in the semantic space.

In this paper, the concept of a cognitive map is used in our user interface design to optimise the cognitive mapping between users' understanding of the environment and the abstract information space. The following section explains concepts relevant to this notion in more detail.

COGNITIVE MAP

The concept of a cognitive map plays an influential role in the study of navigation strategies, such as browsing in hyperspace and wayfinding in virtual environments [8]. A cognitive map could be seen as the internalised analogy in the human mind to the physical layout of the environment [29, 30]. The acquisition of navigational knowledge proceeds through several developmental stages from the initial identification of landmarks in the environment to a fully formed mental map [9].

Levels of Knowledge

Landmark knowledge is often the basis for building our cognitive maps [2, 9]. The development of visual navigation knowledge may start with highly salient visual *landmarks* in the environment such as unique and magnificent buildings or natural landscapes. People associate their location in the environment with reference to these landmarks.

The acquisition of *route* knowledge is usually the next stage in developing a cognitive map. Route knowledge is characterised by the ability to navigate from one point to another using acquired landmark knowledge without association to the surrounding areas. Route knowledge does not provide the navigator with enough information about the environment to enable the person to optimise their route for navigation. If someone with route knowledge wanders off the route, it would be very difficult for that person to backtrack to the route.

The cognitive map is not considered fully developed until *survey* knowledge has been acquired [30]. The physical layout of the environment must be mentally transformed by the user to form a cognitive map.

Dillon *et al.* [9] have noted that when users navigate through an abstract structure such as a deep menu tree, if they select wrong options at a deep level they tend to return to the top of the tree altogether rather than just take one step back. This strategy suggests the absence of survey knowledge about the structure of the environment and a strong reliance on landmarks to guide navigation. As hypertext designers, we are interested in exploring ways to help users overcome a reliance on landmarks so that they can discover optimal routes or paths during information navigation. Fortunately, some studies have suggested that there are ways to increase the likelihood that users will develop survey knowledge of an electronic space. For instance, intensive use of maps tends to increase survey knowledge in a relatively short period of time [8, 29]. Other studies have shown that strong visual cues indicating paths and regions can help users to understand the structure of a virtual space [7].

By and large, visual information navigation relies on the construction of a cognitive map and the extent to which users can easily connect the structure of their cognitive maps with the visual representations of an underlying information space. The concept of a cognitive map suggests that users need information about the structure of a complex, richly interconnected information space. However, if all the connectivity information is displayed, users would be unlikely to navigate effectively in spaghetti-like visual representations. How do designers of complex hypertext visualisations optimise their user interfaces for navigation and retrieval given this conundrum?

One problem faced by designers is that detail concerning an explicit, logical structure may not be readily available in a visualisation. An explicit organising structure may not always naturally exist for a given data set, or the existing structure may simply be inappropriate for the specific tasks at hand. What methods are available for hypertext designers to derive and expose an appropriate structure in the user interface? How can we connect such designs with the user's cognitive map for improved learning and navigation?

In the following section, we will address issues concerning now to single out important structural characteristics to make visual navigation easier, as well as how to filter out redundant information in order to increase the clarity and simplicity of the visual environment.

VIRTUAL INFORMATION SPACES

In this section, we introduce the design of visual representations of various semantic entities. We identify relationships between the user's cognitive map and visual representations of abstract entities that users may encounter as they navigate through the environment. In Table 1, we classify visual representations of objects in accordance with the three types of cognitive knowledge about the underlying environment, namely landmark, route and survey knowledge.

Table 1. Transforming a cognitive map into virtual reality.

Cognitive Map	VR Model	Concrete	Abstract
Landmarks	Reference Points	Document Size Creation Time	Search Results User Profiles
Route Knowledge	Node/Link	Guided Tour	Minimal Cost Paths
Survey Knowledge	Overviews	Geographical Map Underground Map	Semantic Networks Pathfinder Networks

Visual representations in information visualisation systems often fall into two categories. A natural, concrete representation relies on an existing explicit structuring model, for example, data organisation according to a geographical layout. An abstract representation usually does not have an inherited organisation model to convey latent, implicit structures in the data, such as semantic networks. CHI proceedings and other textual documents used in our study are associated with the latter category.

VIRTUAL REALITY MODELLING

Virtual reality modelling is an integral part of our approach. It transforms the blueprint provided by Pathfinder and force-directed graph drawing algorithms to virtual worlds in VRML so that users can visually explore the virtual structure (see Table 2). Several direct manipulation tasks are supported in such virtual worlds, such as walk, spin, slide and examine. When users click on a document sphere, the document, whether it is local or remote, will be downloaded to their client-side browsers.

Direct manipulation-based user interfaces can be easy to learn and use [6]. Virtual reality models provide new ways of interacting with the semantic space, such as walking back and forth through the space, which effectively overcome the traditional focus-versus-context problem [11]. VRML supports the notion of *level of detail* (LOD)— as the user approaches to an object in the virtual world, the virtual world increasingly reveals more information about the object.

Table 2. Visualisation model.

Digital Objects	Geometric Model	Attribute	Semantics
document	sphere	radius	size
document	sphere	colour	source of data
link	cylinder	radius	semantic similarity
link	cylinder	length	latent semantic distance
query	cylinder	height	matching similarity
query	cylinder	colour	keyword

By explicitly representing salient relationships between two documents in a virtual link structure, users are able to see the connectivity patterns in the entire semantic space. Virtual link structures of different natures, be they hyperlinks, content similarity, navigation patterns or bibliographic citations, can be combined and animated to help users to make sense of the complex semantic structure.

VIRTUAL STRUCTURES AND SPATIAL HYPERTEXT

We present the following examples in order to illustrate design and usability issues concerning the spatial hypertext.

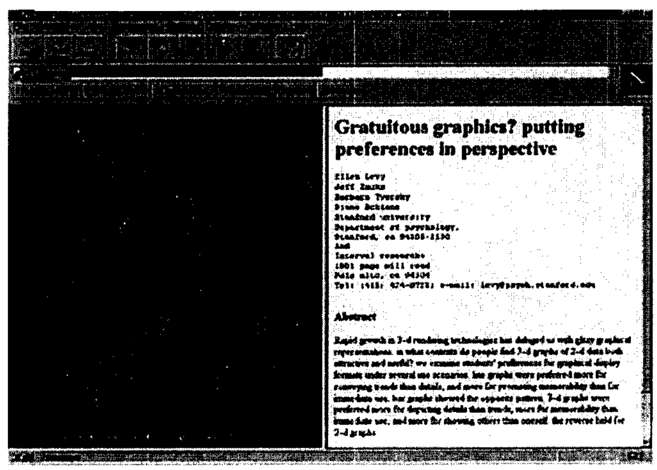

Figure 3. The virtual structure is used with a WWW browser.

Figure 3 shows the virtual space of the recent CHI proceedings (1995—1997). This virtual space is based on the latent semantics characterised by LSI and link structures determined by Pathfinder network scaling. When the user moves the mouse cursor over a document sphere in the structure, the title of the document will appear at the point of the cursor. If the user clicks on the sphere, the abstract of the document will appear in the right-hand side frame.

Landmarks

In our spatial hypertext, predominant landmarks are related to search relevance rankings. A cylinder will appear on a document if the document is sufficiently similar to the query. If the query has a number of distinct terms, the resultant cylinder will consist of cylinders for terms that reached sufficiently high rankings. These landmarks are coloured and labeled to enable users distinguish amongst them easily. Neighboring documents are often likely to contain more keywords, in our experience. The structuring techniques used to build the information visualisation tend to group documents on similar topics near to each other.

Once the user identifies the document with the highest cylinder landmark (indicating the most relevant neighborhood of documents to search through), then he/she can use this document as a starting point to explore the semantic space. For example, some documents nearby may not contain particular terms used in the query, but since they are grouped together by LSI they are likely to have something in common, and thus are worth exploring. The user may simply want to click on the bar's corresponding node and read the most relevant retrieved paper directly.

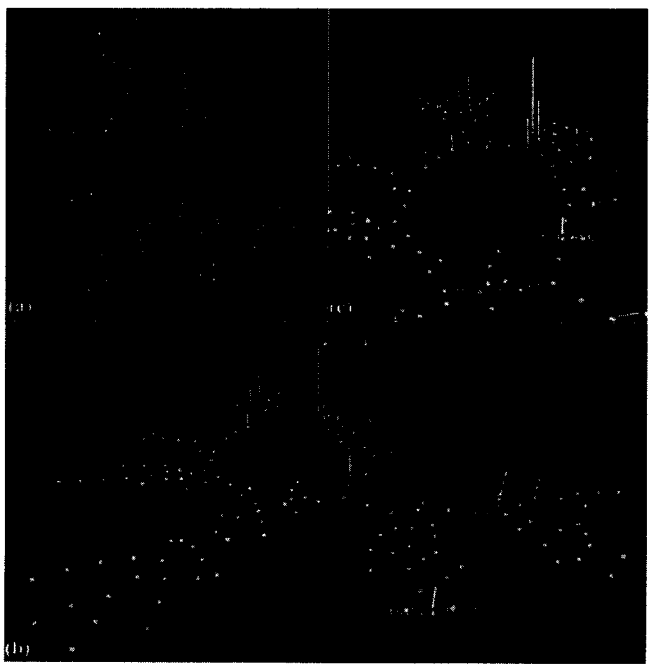

Figure 4. Search and browsing in the semantic space of CHI proceedings (a) Overview, (b) Landsape View, (c) Zoom in.

The virtual space in Figure 4 visualises the result of a search of keywords *digital library* and *spatial map* on the basis of the overall semantic structure of CHI proceedings. In the landscape view, for example, vertical bars highlighted papers that have good match to these words. The height of each bar is proportional to the strength of the match. For example, the best match for *spatial map* (similarity=0.724) is at the far end of the scene in Figure 4b with the highest vertical bar.

There are two general types of hypermedia networks—homogenous and heterogeneous. In a homogenous network, all the nodes are of the same type; for example, the network contains papers and nothing else. In a heterogeneous network, one may deal with different types of nodes; for example, the network not only contains papers, but also contains user profile of their information interests and sample queries (even though many studies have regarded queries as a special type of document). These nodes can be regarded as a special type of landmarks, or reference points. There is a similar notion known as *unfolding* in psychology [17], in which subjects and stimulus are embedded into the same space.

Figure 5 shows an aggregated virtual structure based on three independent collections, namely, CHI papers, HTC papers and papers from one of the authors. CHI papers were coloured in light blue (1995), light green (1996) and light red (1997). HTC papers were rendered as red spheres and papers by one of the authors were in dark blue. Each document is situated in a larger context. Users now can access the three data sets from the single virtual structure, while the original data sets remain intact.

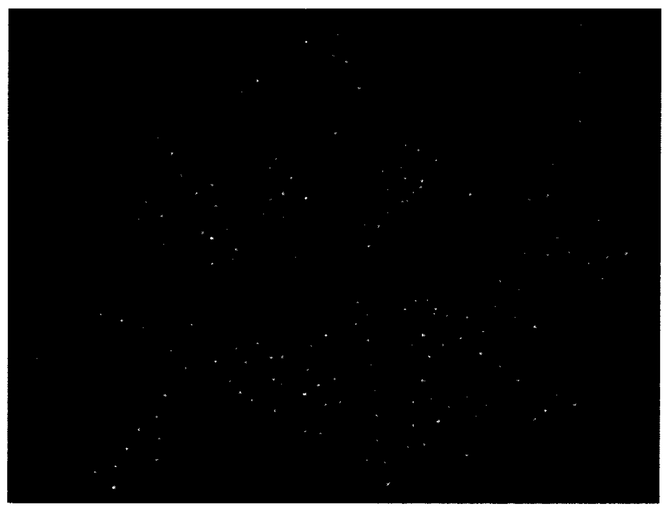

Figure 5. A seamlessly combined semantic structure of 304 papers from three sources, including 169 CHI papers, 127 ACM HTC papers and panels, and 8 papers from the first author.

The combined virtual structures allow us to visually analyse cross-domain interconnections. Neighbouring documents in the space should be of particular relevance to the user. One could use software agents to import other papers into their current personalised digital library automatically.

Route Knowledge

Links preserved by the Pathfinder network are explicitly displayed in our current visualisation techniques. A route from one paper to another has the minimum cost, or the strongest connecting strength. The presence of a route in the virtual environment therefore suggests to the user that papers on the route between two relevant papers may be worth browsing.

Papers from different years were coloured differently. This colouring scheme was designated to detect emerging trends in research questions and application domains addressed by papers in consecutive years of conferences. For example, if we see a group of papers gathered together in blue (i.e., papers from the latest conference), it suggests that new topics were introduced into the conference series. If a group of papers clustered in the network includes every colour but blue, then this may suggest that a particular area was not addressed by papers accepted for that conference.

Self-organised node placement in our approach is based on the spring embedder model, which belongs to a class of graph drawing heuristics known as force-directed placement [15]. The positions of nodes are guided by forces in the dynamic systems. The satisfactory placement is normally obtained when the spring energy in the entire system reaches the global minimal.

General aesthetic layout criteria include minimising the number of link crossing and overlapping, symmetrical displays and closeness of related nodes. We use the term *self-organisation* in this paper to emphasis the role of these heuristics in satisfying several potentially contradicting aesthetic requirements. Although the spring embedder algorithm does not explicitly support the detection of symmetries, it turns out that in many cases the resulting layout demonstrates a significant degree of symmetry.

In addition to the layout heuristics, a good navigation map should allow users to move back and forth between local details and the global context, to zoom in and out the visual display at will, to search across the entire graph. More advanced features may include simulation and animation through consecutive views. Our initial studies show that many of these requirements can be readily met by Virtual Reality Modeling Language (VRML), especially VRML 2.0.

Survey Knowledge

Visual navigation in our virtual environment starts with an overview from a distance. Users then approach the centre of the virtual world for further details. Users have a number of options, such as *walk*, *spin* and *point*. In the next section, we will provide an overview of how an underlying information structure is presented to the user who is visually navigating in our virtual environment.

We will also discuss some preliminary findings from our empirical study in the context of the overall design experience.

SEARCH PATTERNS AND SPATIAL ABILITY

Previous studies in hypertext suggested that spatial ability may be a significant factor affecting users' satisfaction and performance with spatial hypertext systems. We have recently conducted an empirical study to investigate the interaction between users' spatial ability and their search patterns with the spatial hypertext [4]. Here we will summarise some interesting findings of our empirical study as the focus of this paper is on design principles and how these principles can be transformed into spatial hypertext with a set of combined theories and techniques. Readers are referred to [4] for a more detailed report of the empirical study.

Subjects were postgraduate students enrolled in an M.Sc. in Information Systems. They were not familiar with the underlying theories and algorithms used in the design. Subjects were asked to find papers related to particular topics within a 30-minute session. For example, in one task, subjects were asked to find as many papers as they could on information visualisation. In particular, the recall and precision measures were used based on our own relevance ratings. Recall was positively correlated with spatial ability based on a spatial pretest's paper folding scores in two search tasks ($r = 0.42$ and 0.37, respectively). Retrieval accuracy, or precision, was strongly negatively correlated with spatial ability in these tasks ($r = -0.53$ and -0.18, respectively). See [4] for further discussions of this interesting pattern of findings. The important point is that spatial ability strongly influences users' search patterns in these virtual spaces. Individual differences should be considered when designing information visualisations such as ours, and perhaps adapting to the users' abilities over time would be ideal.

Navigation Strategies

In order to study navigational patterns in the spatial semantic space, we superimposed the frequencies of accessing papers that are judged relevant in the first search task, according to a pre-determined relevance judgement, over the visualised semantic structure (see Figure 6). Relevant papers are marked as boxes and the number of dots beside each box indicates how many different individuals successfully found that target.

Task performance (recall and precision) scores suggested that subjects did reasonably well if targets were located in some structurally significant positions in the spatial hypertext. However, if task-relevant papers were located in outskirts of the structure in the user interface, performance hindered. In addition, subjects seemed to be affected by the varying visibility of topical keywords (i.e., whether a search word appears in the title, or is hidden in the abstract, etc.) This could be a serious issue, as it affects whether or not one can easily recognise the relevance of a paper. This is especially significant when papers are located in a key position, such as a gateway or a branching point. If relevant papers are not explored in key positions, we observed that users were more likely to navigate to another key position, rather than continue to explore locally [4]. We found that these key positions, or hotspots, were typically examined by subjects in their first few moves.

The videotapes we recorded of our user sessions revealed that the majority of the subjects regarded the central circle structure as a natural starting point. They tended to aim at the central circle as an initial user interface action and zoom into the virtual world in order to bring this circular area into focus. Outskirts of the central circle tended to be ignored during the initial search. Next, subjects would check a number of positions on the circle, especially points connecting to branches. Over time, subjects would gradually expand their search space outwards to reach nodes farther away from the central area. An example of a good strategy observed was that one subject sampled a single node in each cluster and moved on to other clusters quickly during the initial stage. This strategy maximised the likelihood of not becoming lost in a local minimum.

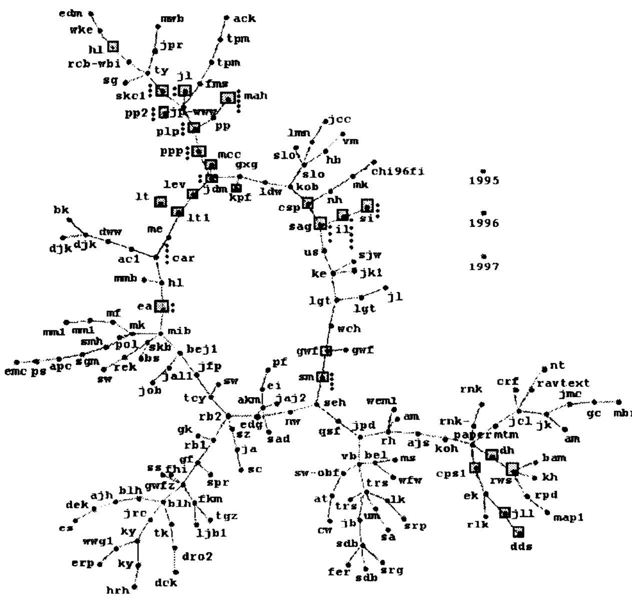

Figure 6. The locations of search targets.

Some subjects hopped from one cluster to another in long jumps, whereas other subjects carefully examined each node along a path according to the virtual semantic structure. Subjects who made longer jumps apparently realised that they might be able to rely on the structural patterns to help with their navigation. Navigational patterns also highlighted the special role of distinctive structural patterns such as circles, stars, and long spikes as we expected. We will be analysing the video more thoroughly to gather more detailed data about navigation strategies and report our findings in the near future.

Spatial Memory

We included a spatial memory test as an alternative viewpoint to look at the interaction between visualised semantic structures and individuals' understanding of how the semantic space is organised. By identifying what subjects learned about the structure and how their memorised user interface details varied from one area to another, we were able to understand more about various characteristics of our visual semantic structure.

Figure 7 shows the sketches of the semantic space from two subjects. These sketches show not only that these subjects focused on different areas in the semantic space, but also that subjects could remember the semantic structures inherent in the user interface quite vividly. These figures are partially related to the differences in interactions between subjects' navigation strategies and their emerging cognitive maps. One interesting question that awaits future research is whether subjects' maps would converge over repeated exposure and use of the information space.

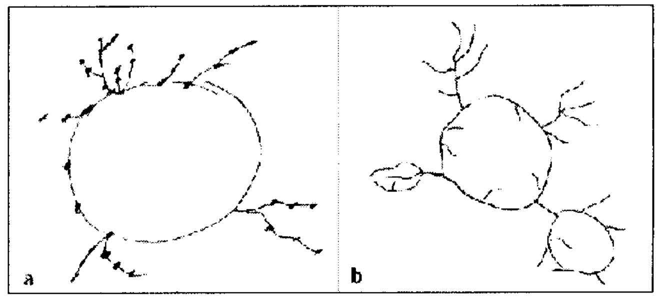

Figure 7. Subjects' sketches of the semantic information space searched during the study.

Most subjects clearly remembered the shape of the central circle. In (a), the subject highlighted the central circle and three sub-areas around the circle. The video analysis confirmed that these had been the most often visited areas in his search. In (b), the subject was able to remember more details about the branches surrounding the central circle. In addition, he added some strokes inside the circle, although they were not as accurate as other structural patterns in his sketch. While this provides an brief hint of how subjects' spatial memory may be influenced by this information visualisation, as well as their individual differences in ability and strategy, we will continue to analyse these structures for meaningful implications for 3D user interface design.

CONCLUSION

In this paper, we have introduced an integrated approach to the development of spatial hypertext. We have highlighted the integral parts played by the theory of cognitive maps, Latent Semantic Indexing, Pathfinder networking scaling and virtual reality modelling. A number of powerful techniques have been incorporated into a generic, extensible and fully automated approach. The use of virtual structures transcends the boundaries of the source data originally stored — they leave all the original data intact. We have demonstrated that searching and browsing can be accommodated within the same semantic space.

The design practice and our preliminary empirical evaluation have provided some valuable experience and insights into the spatial hyperspace [4]. We are planning to conduct more studies in related areas, such as evaluating the usability of such virtual environments and investigating the role of individual differences in the use of spatial user interfaces, especially spatial ability and cognitive styles.

Interpretive use of spatial hypertext has been recognised as an important and fruitful area for future research [22, 28]. Research in this area is increasingly associated with the need for a better understanding of social and ecological dimensions of hypermedia and digital libraries. Information seeking, for example, should be understood in a wider context of intellectual work and from social perspectives of knowledge, communication and collaboration. Our next step is to explore these issues concerning the design of a collaborative digital library. We will be investigating these issues through the development of a domain-specific virtual community environment based on the British Computer Society's HCI conference proceedings (1985—1997). We will continue to explore practical issues in these ongoing projects. We intend to investigate dynamic space transformation in response to usage patterns of users. In addition, we would like to explore more opportunities of applying these techniques to real world situations as a part of the iterative development of the approach.

ACKNOWLEDGMENTS

This work is currently supported by EPSRC research grant GR/L61088. Bell Communication Research kindly provided the software for Latent Semantic Indexing. We would like to thank the anonymous reviewers for their helpful comments.

REFERENCES

1. ACM *The ACM Hypertext Compendium™*, ACM Press, 1991.

2. Anderson, J. *Cognitive psychology and its implications*. W. H. Freeman, San Francisco, 1980.

3. Chalmers, M. and Chitson, P. Bead: Explorations in information visualisation. In *Proceedings* of *SIGIR'92* (Copenhagen, Denmark, June 1992), pp. 330-337.

4. Chen, C. and Czerwinski, M. Spatial ability and visual navigation: An empirical study. *New Review of Hypermedia and Multimedia*, (1998). Available at http://www.brunel.ac.uk/~cssrccc2/papers/nrhm.ps.gz

5. Chen, C. Structuring and visualising the WWW by generalised similarity analysis. In *Proceedings of Hypertext'97* (Southampton, England, April 1997), pp. 177-186.

6. Chimera, R. and Shneiderman, B. An exploratory evaluation of three interfaces for browsing large hierarchical tables of contents. *ACM Transactions on Information Systems*, 12, 4, (1994), 383-406.

7. Czerwinski, M. and Larson, K. The new Web browsers: They're cool but are they useful? In *Proceedings of HCI'97* (Bristol, England, 1997).

8. Darken, R. P. and Sibert, J. L. *Wayfinding strategies and behaviors in large virtual worlds*. In *Proceedings of CHI'96*, (Vancouver, B.C., April 1996). Available at http://www.acm.org/sigs/sigchi/chi96/proceedings/papers/Darken/Rpd_txt.htm

9. Dillon, A., McKnight, C. and Richardson, J. Navigating in hypertext: A critical review of the concept. In *Human-Computer Interaction — INTERACT'90* (D Diaper et al. eds). Elsevier Science Publishers, 1990, pp. 587-592.

10. Deerwester, S., Dumais, S. T., Landauer, T. K., Furnas, G. W. and Harshman, R. A. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41, 6, (1990), 391-407.

11. Furnas, G. Generalised fisheye views. In *Proceedings of CHI'86*, (Boston, MA, April 1986), pp. 16-23.

12. Fairchild, K., Poltrok, S. and Furnas, G. Semnet: Three-dimensional graphic representations of large knowledge bases. In R. Guindon (Ed.), *Cognitive Science and its Applications for Human-Computer Interaction*, Lawrence Erlbaum, 1988, pp. 201-233.

13. Haake, J. M., Neuwirth, C., and Streitz, N. A. Coexistence and transformation of informal and formal structures: Requirements for more flexible hypermedia systems. In *Proceedings of ECHT'94*, (Edinburgh, Scotland, September 1994), pp. 1-12.

14. Hill, G. J. and Hall, W. Extending the Microcosm model to a distributed environment. In *Proceedings of ECHT'94*, (Edinburgh, Scotland, September 1994), pp. 32-40.

15. Kamada, T. and Kawai, S. An algorithm for drawing general undirected graphs. *Information Processing Letters*, *31*, 1, (1989), 7-15.

16. Kellogg, R. B. and Subhas, M. Text to hypertext: can clustering solve the problem in digital libraries? In *Proceedings of DL'96*, (Bethesda, MD, March 1996), pp. 144-150.

17. Kruskal, J. B. Multidimensional scaling and other methods for discovering structure. In K. Enslein, A. Ralston, and H. Wilf (Eds.), *Statistical methods for digital computers*. Wiley, New York, 1977.

18. Lin, X., Soergel, D., and Marchionini, G. A self-organizing semantic map. In *Proceedings of SIGIR'91*, (Chicago, Illinois, October 1991), pp. 262-269.

19. Lokuge, I., Gilbert, S. A., and Richards, W. *Structuring information with mental models: A tour of Boston*. In *Proceedings of CHI'96*, (Vancouver, B.C., April 1996). Available at http://www.acm.org/sigchi/chi96/proceedings/papers/Lokuge/sag_txt.html

20. Marshall, C. C. and Shipman, F. M. Searching for the missing link: Discovering implicit structure in spatial hypertext. In *Proceedings of Hypertext'93*, (Seattle, WA, November 1993), pp. 217-230.

21. Marshall, C. C., Shipman, F. M. and Coombs, J. H. VIKI: Spatial hypertext supporting emergent structure. In *Proceedings of ECHT'94*, (Edinburgh, Scotland, September 1994), 1994, pp. 13-23.

22. Marshall, C. C., Shipman, F. M. Spatial hypertext: Designing for change. *Communications of the ACM*, 38,

8, (August 1995), 88-97.

23. Orendorf, J. and Kacmar, C. A spatial approach to organizing and locating digital libraries and their content. In *Proceedings of DL'96*, (Bethesda, MD, March 1996), pp. 83-89.

24. Robertson, G. G., Mackinlay, J. D., and Card, S. K. Cone Trees: Animated 3D visualisations of hierarchical information. In *Proceedings of CHI'91,* (New Orleans, LA. 1991), pp. 189-194.

25. Salton, G., Singhal, A, Buckley, C., and Mitra, M. Automatic text decomposition using text segments and text themes. In *Proceedings of Hypertext'96*, (Washington, D.C., March 1996), pp. 53-65.

26. Salton, G., Allan, J., and Buckley, C. Automatic structuring and retrieval of large text files. *Communications of the ACM, 17*, 2, (1994), 97-108.

27. Schvaneveldt, R. W., Durso, F. T., and Dearholt, D. W. Network structures in proximity data. *The Psychology of Learning and Motivation: Advances in Research and Theory*, 24, (1989), 249-284.

28. Streitz, N. A., Hannemann, J., and Thüring, M. From ideas and arguments to hyperdocuments: Travelling through activity spaces. In *Proceedings of Hypertext'89*, (Pittsburgh, PA, November 1989), pp. 343-364.

29. Thorndyke, P. and Hayes-Roth, B. Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14, (1982), 560-589.

30. Tolman, E. C. Cognitive maps in rats and men. *Psychological Review*, 55, (1948), 189-208.

31. Whitney, V. K. M. Minimal spanning tree: Algorithm 422. *Communications of the ACM* 15, 4, (1972), 273-274.

32. Zizi, M. and Beaudouin-Lafon, M. Accessing hyperdocuments through interactive dynamic maps. In *Proceedings of ECHT'94*, (Edinburgh, Scotland, September 1994), pp.126-135.