Journal of Network and Computer Applications (1997) **20,** 135–149

3D visualization of WWW semantic content for browsing and query formulation

Richard H. Fowler, Wendy A. L. Fowler and Jorge L. Williams

Department of Computer Science, University of Texas, Pan American, 1201 W. University, Edinburg, TX 78539-2999, USA

Visualization is a promising technique for both enhancing users' perception of structure in the Internet and providing navigation facilities for its large information spaces. This paper describes an application of the Document Explorer to the visualization of WWW content structure. The system provides visualization, browsing, and query formulation mechanisms based on documents' semantic content. These mechanisms complement text and link-based search by supplying a visual search and query formulation environment using semantic associations among documents. The user can view and interact with visual representations of WWW document relations to traverse this derived document space. The relationships among individual keywords in the documents are also represented visually to support query formulation by direct manipulation of content words in the document set. A suite of navigation and orientation tools is provided which focuses on orientation and navigation using the visual representations of document set and term collection.

1. Introduction

The explosive growth of the Internet has added to the need to filter and organize information so that users can efficiently and effectively identify relevant documents. Browsers such as Mosaic [1] for the World Wide Web and Harmony [2] for Hyper-G allow users to traverse the information space of documents through connections provided by document authors. Though this may account for much of the popularity of the Internet, it hides many of the difficulties of effective information access. Searching Internet resources is difficult due to size, diversity of data, and lack of a common indexing scheme. The most widely-used search tools are automatically generated search systems like Lycos and Harvest and manually organized systems like Yahoo and Internet Yellow pages. Though such tools are valuable assets to the Internet searcher, it seems likely that these tools alone will not solve the current problems of information access.

The challenges of information access on the Internet are issues common to all forms of information retrieval. These issues include difficulties in using indexing vocabularies, indexing indeterminacy, and the user's inability to completely specify information needs [3]. Retrieving information that meets users' information needs is an iterative process, and techniques which explicitly incorporate users' judgments, such as relevance feedback [4], provide means to automate some aspects of user-guided retrieval. It is also clear that mechanisms providing alternative paths of access to information can enhance retrieval effectiveness [5].

1084–8045/97/020135 + 15 $25.00/0 ma970042

1.1 *Visualization and the Internet*

One promising approach for enhancing information retrieval through the Internet is visualization to facilitate users' perception of document relation structure. A number of systems have been developed to provide visually-based browsing mechanisms for traversing the link structure of Internet documents. McCahill and Erickson [6] describe designs for three-dimensional (3D) spatial interfaces for Internet Gopher which employ icons with various shapes and textures to represent documents arranged by file structure or search results. The Harmony browser for Hyper-G [7] provides 2D structure maps of a document's link neighborhood, as well the 3D information landscape visualization of link structure hierarchy. The Narcissus system [8] creates a 3D visual representation of WWW link connectivity with documents' positions determined by user manipulable attractive and repulsive forces among document nodes. Munzner and Burchard [9] describe a system which displays WWW documents positioned in a hyperbolic space to increase display density compared to conventional display spaces.

Visualizations of document connectivity based on existing, author-created links provide an important component in facilitating orientation and navigation in the large WWW information space. However, for users it is typically documents' content relationships which are of most interest, rather than link-based relationships, and such semantic relationships are only partially reflected by link structure. Documents with similar content can be found throughout the WWW information space. Though it is likely that a user would find documents dealing with similar topics close together within a site, it is also likely that other documents dealing with similar topics are at other sites, and those related documents may be buried deep within the link structure of a site.

Systems which provide facilities for organizing and viewing document organizations based on semantic content are less common than systems which focus on visualizing document link structure. Mukherjea, Foley and Hudson [10] describe a system which operates on the semantic content of WWW documents to form visually displayed hierarchies. In this system users can specify attributes which affect the organization of documents derived by the system and reconfigure the display to suit their needs. VR-VIBE [11] is a system which fixes WWW documents at locations in 3D space by allowing users to interactively position keywords on a pyramid. Document locations are determined by calculating the distance for each document from keywords' locations on the pyramid. In this way users create visually distinct clusters of documents for further inspection. LyberWorld's Relevance Sphere [12] also uses this general approach for forming 3D document cluster displays. Gershon *et al.* [13] describe a system which allows users to view an organization of documents visited as a hierarchy of links and also to construct a separate hierarchy based on their own needs. Additionally, the system creates a co-occurrence map of terms in documents visited to help users identify word patterns useful in a more fine-grained search.

1.2 *Components of information visualization*

Information visualization shares with scientific visualization the goal of providing understanding and insight using visual representations, yet is often faced with challenges not encountered in scientific visualization. For scientific visualization the visual representation is derived from a *physical* phenomenon, hence, there typically exists a

Table 1. *Components of information visualization*

- Spatialization of abstract data:
  - Information organization
  - Derivation of a spatial-visual representation:
    - Structure extraction
    - Mapping into visually realizable form
- Display space
- User interaction

natural visual representation based on spatial, temporal, or other properties of the phenomenon. For information visualization elements to be visualized typically have no physical component which might naturally supply a basis for visualization [14]. Rather, the elements to be visualized often only have *semantic* properties with no inherent spatial analog from which to create a visual representation. Any spatial ordering for semantic relations must be created as part of the information visualization process. From this perspective, visualization of semantic information spaces must provide: (1) spatialization of the abstract data, which may entail both (i) data organization and (ii) derivation of a visual spatial representation of the data, (2) presentation of the spatial representation in a display space for user interaction and viewing, such as Euclidean 2- or 3D space or a distorted space, and (3) techniques and tools for user interaction with the visual representation. This characterization of information visualization is summarized in Table 1.

One aspect of spatialization is the process of *organizing the information*, or data, which is to be visualized. For example, extracting keywords to provide a multi-dimensional vector space. For some systems the data organization is already available in the data to be displayed, e.g. a file system hierarchy. Spatialization also creates a viewable, or displayable, representation, *a spatial-visual representation*. It typically is at this step that the final spatial nature of a particular data organization is derived by transforming it to a representation which is useful for information visualization. The term spatial-visual is chosen to emphasize that it is a form which is at once spatial, and also visual, i.e. viewable in three dimensions. This distinguishes it from representations which may be spatial, e.g. vector space, which is spatial in n dimensions, but which are not viewable without a reduction of dimensionality through a mapping into a visualizable, typically three dimensional, space. The derivation of a spatial–visual representation is at the core of information visualization and much of the utility of information visualization in the future will come from the development of useful mappings of data organizations to visual forms.

As used here, *display space* is the visual space used for presentation to the user. The separation is in part motivated by the observation that when investigators name a system or technique it is really the combination of spatialization methods, together with a particular display space and set of interaction methods which create the unique, named technique. Most frequently, display spaces are simply three dimension perspective projection, e.g. the Xerox PARC cone tree [15] is radial tree layout in three dimensions presented in Euclidean 3D space.

The third set of components in this cahracterization are *user interaction mechanisms*. These are the ways of interacting with the spatial display and sometimes interactively changing the spatial-visual representation, e.g. changing the fisheye view filtering or expanding and collapsing nodes in a network. Most systems include such user-manipulable filtering mechanisms which affect the spatial–visual representation to varying degrees, e.g. recalculating degree of interest or just hiding or showing elements.

2. Visualizing WWW semantic content with the Document Explorer

This paper describes an application of the Document Explorer system to the visualization of WWW content structure. Visualization, browsing, and query formulation mechanisms are based on documents' semantic content. These representations complement text and link-based search by supplying a visual search environment using semantic associations among WWW documents. The Document Explorer provides mechanisms for each of the three components of information visualization outlined above to display a semantic space of WWW documents and keywords. WWW document relationships and keyword relationships are organized as associative networks based on content similarity among documents and frequency of keyword co-occurrence. The networks are presented in 3D display space with network nodes positioned using a layout algorithm based on a mechanical spring metaphor. A variety of interaction and display tools are used to supply orientation cues and navigation mechanisms to assist users in browsing the document and keyword networks, as well as in the formulation of queries by direct manipulation of system objects.

The content based network of WWW documents is shown in Fig. 1 above as the main window of the screen display. In this view documents are labeled by their content. Alternatively, the display can present the HTML title. Below are overview diagrams showing the location of the detailed view in the main window in the complete network. To the left and right of the overview is a series of visual bookmarks set while browsing which can be used to return to a previous viewpoint. The leftmost shows a view of the complete network. Other navigation and orientation tools are also available, such as anchors and signposts, which the user can leave and revisit at points traversed in the network. Change of viewpoint using navigation aids is always done by zooming to new viewpoints to maintain fluid motion and minimize disorientation. In the upper right of the screen is a natural language query which has been transformed to a user-manipulable query graph. The query can be used to supply an entry point in the document network or be used for conventional weighted vector search to provide a list of documents from which to initiate browsing.

2.1 *Extracting and organizing WWW semantic content*

The system's principal visualizations are network displays based on documents' keyword lists. The lists can be provided by automatic content extraction tools, such as Harvest [16] or derived within the system. Keyword lists for each document are used to determine the associations among documents and among terms. The statistical text analyses rely on recovering conceptual information from natural language by considering the frequency and co-occurrence of words. This basic approach has been used in a wide range of contexts and its utility and limitations are well known [17].

The visual environment for exploration and direct manipulation uses the same visual

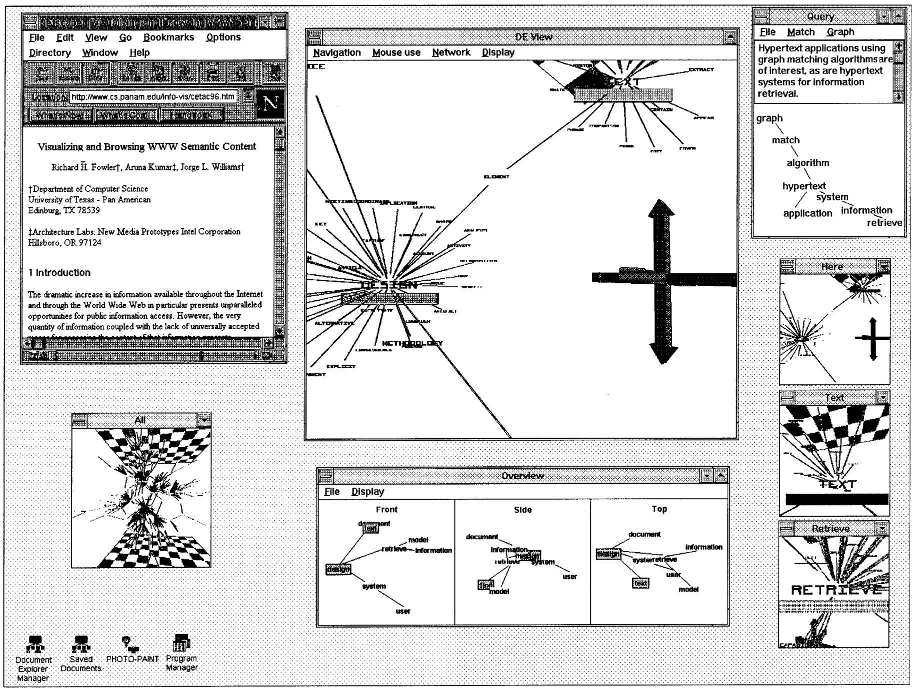

Figure 1. Document Explorer screen. The content based document network is shown in the main window with overview diagrams below it. Windows to the left and right of the overview show visual bookmarks the user has set. At the upper right is the user's natural language query and its visual representation.

representation for query, associative thesaurus, and document content to facilitate query revision by direct manipulation of system objects. The representations underlying the visual displays are minimum cost networks derived from measures of term and document associations. The associations are derived from natural language text for queries, single documents, and associative term thesauri. The document collection is represented as a network of documents based on interdocument similarity.

2.1.1 *Deriving the term associations.* For each set of WWW documents the system uses a separate set of terms formed from the most frequently occurring word stems, excluding function words. For some forms of retrieval this simple procedure suffers from the limitation that frequently-occurring terms have relatively little value for discriminating among documents [18]. However, one function of the associative thesaurus is to give a picture of all of the concepts in a document set. The most frequently-occurring terms tend to be general terms that provide useful information about the domain of the document collection.

To derive the distances between terms used to construct networks, text is analysed by first finding the sequence of term stems present in the text. This sequence is used to assign distances between terms based on lexical distance and co-occurrence in syntactic units with a metric similar to that used by Belkin and Kwasnik [19]. Term pair similarity is calculated as the sum of values added when terms are adjacent, or occur in the same sentence, paragraph or document. These similarities provide the associations used in deriving the networks displayed by the system.

There are three reasons for using statistically-based associative structures in an interactive browsing system. One reason follows from the view that information retrieval systems should supply the user with a variety of tools and retrieval techniques. Statistically-based associative information structures provide one class of retrieval tools that can complement other retrieval aids. For example, an associative thesaurus based on term co-occurrence in documents presents a structure of term relationships quite different than presented in a thesaurus showing term hierarchies. The associative thesaurus can encourage browsing and exploration, as well as bring the user's own associations into play. For information needs in which the user is not familiar with the domain, and indeed may not even know what his or her information needs are, the associative structures provide one means to explore and gain information to better define the information need.

A second reason for using statistically-based associative structures is the desire to have a representation that can be derived automatically in an interactive system, rather than through knowledge-engineering efforts such as are required for most deep representations. Associative structures have been used effectively as one component of hybrid systems incorporating both deep and shallow representations [20]. The final reason is the desire to provide a common visual representation for retrieval tools. Networks are naturally represented visually and can provide a common representation for the system's several visualization components.

2.1.2 *Pathfinder networks: a technique for network abstraction.* The associative networks used in the system are Pathfinder (PfNets) [21]. The Pathfinder algorithm was developed to model semantic memory in humans and to provide a paradigm for scaling psychological similarity data [22]. A number of psychological and design studies have compared PfNets with other scaling techniques and found that they provide a useful tool for revealing conceptual structure [23,24].

The PfNet representations underlying the system's network displays are minimum cost networks derived from measures of term and document associations. The network of documents is based on interdocument similarity, as measured by co-occurrence of keywords between document pairs. For the network of terms, or associative term thesaurus, the visual representation of the user's query, and single document representations the associations are derived from text with association measured by keyword co-occurrence and lexical distance within documents. PfNets can be conceptualized as path-length-limited minimum cost networks. Algorithms to derive minimum cost spanning trees (MCSTs) have only the constraints that the network is connected and cost, as measured by the sum of link weights, is a minimum. For PfNets, an additional constraint is added: not only must the graph be connected and minimum cost, but also the longest path length to connect node pairs, as measured by number of links, is less than some criterion. To derive a PfNet direct distances between each pair of nodes are compared with indirect distances, and a direct link between two nodes is included in the PfNet unless the data contain a shorter path satisfying the constraint of maximum path length.

In constructing a PfNet two parameters are incorporated: r determines path weight according to the Minkowski r-metric and q specifies the maximum number of edges considered in finding a minimum cost path between entities. As either parameter is manipulated, edges in a less complex network form a subset of the edges in a more complex network. Thus, the algorithm generates two families of networks, controlled by r and q. The least complex network is obtained with $r=\infty$ and $q=\mathrm{n}-1$, where n is the total number of nodes in the network. The containment property has in practice provided a particularly useful technique for systematically varying network density to provide both relatively sparse networks (the union of MCSTs with $r=\infty$ and $q=\mathrm{n}-1$) for global navigation, as well as more dense networks for local inspection.

Reducing of the complexity of network representations is a central objective in many efforts in visualizing Internet information structures. Complexity can be reduced by transforming a network of document connectivity to a hierarchy by removing links. Compared to the more general problem of representing directed graphs, visualization of hierarchies is relatively well developed. This sort of transformation is used by Mukherjea *et al.* [10] for WWW documents. In the Document Explorer, nodes in the complete network are identified for display and navigation in a fashion similar to the identification of cluster centroids in single link clustering. These document nodes are continuously displayed in the overview diagram and the three-dimensional space within which the user navigates.

The system uses networks of differing densities to provide separate views of document and term interrelationships. The least dense network, which allows paths of any length in satisfying the minimum cost criterion, is useful for global navigation and orientation. Conceptually, it shows the strongest relations among elements. As such, it is effective in supplying views for global navigation and structure perception. To provide a more detailed view of relationships the system also maintains a PfNet of elements in which the maximum path length is relatively small, creating a complementary network with many more links. This more dense network display is most beneficial when viewing a small set of elements and serves as a sort of magnifying glass for revealing relations among elements not shown in the sparse network.

2.2 *Spatial representation of the system's networks*

A number of data presentation problems require the drawing or display of graphs, and interest in computer-based visualization has increased attention to methodologies for graph display in three dimensions. For visual presentations graph layout criteria centre on how quickly and clearly the meaning of the diagram is conveyed to the viewer, the readability of the graph. Graph drawing algorithms have as their goal the layout and presentation of an inherently mathematical entity in a manner which meets various criteria for human observation. The aesthetics of a layout determine its readability and can be formulated as optimization goals for the drawing algorithm [25]. For example, the display of symmetry and minimization of the number of edge crossings in two-dimensional drawings are fundamental aesthetics for visual presentations.

Considering the wide application of graph structures in display, there are relatively few algorithms for drawing general undirected graphs [26]. This is due in part to the inability to specify the aesthetic criteria individuals use in understanding graphs [27]. Nonetheless, the certain restricted classes of graphs in which graph-theoretic expressions of aesthetic criteria can be specified, satisfactory algorithms have been developed [28–30]. Some of the aesthetics for drawings of general undirected graphs are symmetry, minimization of edge crossings and bends in edges, uniform edge lengths, and uniform node distribution. These aesthetics are such that optimality of one may prevent optimality in others. Additionally, graph layout algorithms in general can be viewed as optimization problems and are typically NP-complete or NP-hard. These two observations suggest a heuristic approach to general graph drawing for many applications.

The spring embedder algorithm [31] is a heuristic approach to graph drawing based on a physical system. This algorithm stimulates a mechanical system in which a graph's edges are replaced by springs and nodes are replaced by rings connecting edges, or springs, incident on a node. From the initial configuration of ring positions, the system oscillates until it stabilizes at a minimum-energy configuration. Among the parameters that control the forces acting on the rings and causing their movement are spring length, spring stiffness, spring type, and initial configuration. This is a very general heuristic which can be combined with other algorithms [32] to provide approximate solutions for competing aesthetics.

The spatial representations of the Document Explorer's networks are designed to facilitate users' perception of network structure. Network nodes are positioned in three dimensions using a graph layout algorithm [33] based on the spring metaphor which is similar to Kamada and Kawai's [34] 2D network layout algorithm. As with other spring embedder algorithms, nodes are treated as connectors and spring length and strength among connectors is derived from network link distances. Nodes are allowed to vary in three dimensions and iteratively positioned at the points which minimize energy in the system of springs. Varying spring length and strength allows layouts which are useful for user interaction and visually reveal clustering and connectively among nodes.

2.3 *Display and interaction mechanisms*

One of the central challenges in the display of large information spaces is to overcome users' feelings of disorientation, the feeling of being 'lost in hyperspace'. In the Document Explorer feelings of disorientation are attenuated in part of the layout of network nodes to facilitate perception of global structure coupled with overview diagrams which track the users current viewing position. Additionally, facilities to interactively vary network density provide global orientation while examining local detail. Finally, navigation tools including bookmarks, anchors, and signposts supply mechanisms allowing users to control the course of navigation and facilitate way finding in the large document and term spaces.

2.3.1 *Overviews, orientation, and network abstraction.* The size and density of the Internet document network requires viewing and navigation tools that allow users to perceive the overall structure of document relations, explore smaller regions in detail, and select and view individual documents. Display and interaction mechanisms in the Document Explorer supply orientation and overview of the global structure of document associations, together with navigation and retrieval tools for exploring local detail. Overview diagrams [35] display a small number of nodes selected to provide information about the organization of the complete networks. These nodes provide orientation by serving as landmarks to assist the user in knowing what part of the network is currently being viewed. Additionally, the nodes provide entry points for traversing the network.

In the Internet document collections we have used, PfNets derived for associative thesauri and networks of documents have a characteristic structure. There tend to be a small number of nodes that have many nodes directly connected and there are relatively short paths between these highly connected nodes, i.e. there are relatively few high degree nodes and network diameter is small. This network form suggests a criterion for selecting nodes to include in overview diagrams. The system's overview diagrams include those nodes of highest degree in the complete network. The term nodes selected for associative thesauri overviews tend to be general terms that provide a guide to the content of the database, and document nodes selected to appear in overviews tend to be general articles about a content area. The overview nodes are landmarks in that they supply information about both the content and structure of the database. As the user changes viewpoint in the main viewing window, and thus the portion of the network which is viewed in detail, the overview diagram tracks the overview nodes which are visible in the view volume of the detailed view. This helps attenuate disorientation by providing the context for the individual network nodes which are in view.

Other display mechanisms for changing network views are also designed to facilitate orientation to the overall structure while examining local detail. The availability of two separate networks which differ in density, or number of links, provides a form of network abstraction enhancing user's perceptual extraction of network structure, as well as supplying a mechanism for navigation. The most sparse network displayed by the system is essentially a tree, thus having the fewest links necessary to have all nodes connected. It is this sparse network which is used to provide the spring analogues used to position nodes in 3D space. The arrangement of nodes on this basis in practice supplies a characteristic clustering of nodes in which structure is relatively easy to perceive, such as shown in the leftmost bookmark of Fig. 1. Yet, much of the utility of browsing in the derived semantic space of documents comes from exploring relations which are relatively weak. These relations are captured in the more dense network. The user typically identifies a single document or document region by browsing or traversing links from a known document in the sparse network, and then selects to have the weaker links of a node or node set become visible. These links are identified by a differing colour and the point at which the selection is made is identified by a marker, or anchor, in the network which is relatively large and can be easily zoomed back to during the course of exploration.

2.3.2 *Other navigation and orientation tools.* Several navigational tools supply mechanisms to return users to previous viewing points or provide information about the current viewing position's context. Visual bookmarks, shown in Fig. 1, can be set by the user at any point while moving through the network. The bookmark is displayed as an annotated snapshot of the user's view. Selecting the bookmark returns the user to the viewpoint at that time the bookmark was set. Visual anchors, displayed as coloured arrows, can be set at any time and remain visible from long viewing distances. Selecting an anchor also returns the user to the viewpoint at the point the anchor was set.

Signposts provide a technique to supply context information about a network region when it is being viewed in detail. As mentioned above, users can select nodes in the overview windows to navigate to a viewpoint near the selected overview node, and this facility is often used as the initial method of exploring the large network. In practice users quickly become familiar with the relatively small set of overview document labels and terms. Signposts exploit the users familiarity with the overview nodes in mechanism which provides global context while viewing a detailed local region. Creating a signpost displays a three-dimensional arrow, similar to the visual form of an anchor, but with each point of the arrow labelled with the nearest overview node. The labelled points of the arrow show the direction of the closest overview node. The same orientation information about the location of the detailed view in relation to the complete network is, of course, available to the user by inspecting the overview window to see which portion of the network is being viewed in detail, i.e. discerning the location of the local detail by examining the global context. Yet, the signpost provides a complementary type of information. Information about the global context, the location of the overview nodes, presented within the detailed view.

Additional display and orientation facilities are available. Link and node colourings can be set to indicate parts of the network of high similarity to the user's query. A backtracking facility allows users to retrace the path followed to arrive at a particular viewpoint. Users can select a node which represents a single document or term and expand or collapse connected nodes. Again, colour is used to mark the course of browsing and network exploration. The user can enter the associative thesaurus from any node displayed, e.g. from the query or an individual document. Similarly, any document, e.g. from the results of a search, can serve as an entry point into the network of documents. All navigation maintains fluid movement in the space, always zooming, rather than jumping, to new viewpoints.

2.3.3 *Head-tracked stereoscopic display.* We have also explored display and interaction using head-tracked stereoscopic display of the networks. For the sorts of network displays in the Document Explorer there is an increase in users' abilities to perceive structure and perform 3D interaction tasks, as found in a number of recent studies of stereoscopic network viewing, both with and without head position tracking [36,37]. Recent attention to immersive interfaces has focused on head mounted displays (HMDs). Such displays typically provide the user a wide field of stereoscopic view and use head position tracking to determine view. Though HMDs can be quite effective in evoking the feeling of presence necessary for total immersion, the apparatus and restricted view preclude their use in most environments.

Head-tracked stereoscopic display (HTS) using LCD shutter glasses afford a compromise in which many of the perceptual and performance advantages of fully immersive interfaces obtain, but the user is free to interact less encumbered by the display apparatus, and information can be displayed with resolution sufficient for text-based tasks. It supplies a relatively immersive interface, yet allows the user to work comfortably at a desk and switch among tasks. In HTS 3D objects appear to be positioned in a volume bounded by the screen perimeter and extending from somewhat in front of the screen to somewhat behind the screen. Much of the utility of 3D display is obtained from the user's active interaction which changes the viewpoint and allows the user to confirm or disconfirm perceptual hypotheses. In HTS simple head movements easily and naturally change viewpoint to 'look around' the objects. This viewing paradigm has been called 'fish-tank virtual reality' and human factors experiments indicate enhanced perception and interaction in three dimensional tasks for HTS viewing compared to flat screen presentations [37].

The Document Explorer provides steroscopic viewing together with untethered head position tracking for HTS viewing. The enhanced 3D perception and interaction supplied by this display addresses the needs in our system to aid the user in perceiving structure and interacting with the system's large networks. The system also provides sufficient display resolution and user-controlled switching from three to two dimensions for text-based tasks typical of document retrieval. In using the various system displays the user moves between stereoscopic 3D displays for some tasks, e.g. viewing the document collections and term networks, and 2D viewing, such as reading a document abstract, without removing any apparatus.

3. Query formulation

Although the system's most novel features centre on its visual representations to support browsing and structure perception, it is useful for information retrieval systems to supply multiple paths of access to information items. For the type of minimum cost networks used in this system, there is close relation to various clustering algorithms. Using the visual document space representation to browse documents can be characterized as a form of user directed, cluster-based search. The system also supplies conventional vector space retrieval as an adjunct, supported by direct manipulation techniques for forming queries.

In a typical use of the system the user enters a natural language statement of information need to begin the browsing and retrieval process. The system converts the natural language to a weighted vector representation and uses conventional weighted vector matching to form a sequence of documents matching a vector representation of the query. The user can then select a document from this list to serve as the entry point in the network of documents, i.e. the viewpoint is positioned near that document and the document is distinguished by colour.

The availability of an association map of keyword relations, a visual representation of term space, assists the user in formulating queries. The relations among terms as portrayed in an association map are typically quite different than found in a conventional thesaurus and serve as an alternative ordering of term relations. To support query formulation the system provides direct manipulation facilities for constructing queries. A visual, manipulable representation of the user's query is displayed together with the natural language from which it was originally derived. This visual representation can be manipulated by the user by moving terms from the term network into the query window and connecting them to the query graph. The vector space retrieval is based on the query graph and is performed by converting it to a weighted vector representation. This functionality could certainly be provided by other interaction mechanisms, yet the visual graph manipulation technique encourages active interaction generally, and allows the use of network representations throughout the system.

4. A sample session

Perhaps the easiest way to get a feel for the system is to describe a sample session using the Document Explorer. The user's initial display might only have the query window and a WWW browser present. The user enters their request for information in the query window using natural language. At this point a number of alternatives are available for finding WWW documents. In addition to visualization mechanisms, the system also provides conventional retrieval using weighted vector matching. The results of a search are typically used to provide a starting point for visually-based browsing of the document and term networks. In this example the user selects to perform a search using a vector representation of the natural language query. The ranked list of documents based on the comparison of the query keyword vector and document vectors is presented to the user in another window as icons of the documents' titles. Now the user can examine WWW documents as in conventional text-based system by selecting a title, and the system sends the documents' URL to the WWW browser for display.

The user can also use the documents found in this search to begin use of the system's visualization capabilities. In addition to the examining of a document's text in the browser, the user can examine a visual network representation of a document. This single document network representation is constructed from the document's text to emphasize the content through the visual form of the network displayed and facilitate a rapid scanning of the ranked document set through successive views of document networks. Using either the document scan facility or the text representation the user might find a document of interest and move it to the saved document window.

After viewing a number of the documents from the ranked document set the user finds that in addition to the words they used to initially state their information needs, there are some other words that are also relevant to finding all the information they need. The user decides to revise the initial query using the visual direct manipulation facilities of the system. To do this they indicate to the system to convert the initial query to a network representation and display that visual representation in addition to the text version. The user then uses the currently-displayed single document network to copy a term node into the query network window and connect it to the query to form a new network used for retrieval.

A second search is performed using the modified query, and the new ranked document set can be examined. The process of scanning documents is continued and more documents are added to the saved document set. After gaining some familiarity with the keywords which are relevant to the user's information needs through examining document texts and networks, the user decides to learn more about the terms by consulting the associative thesaurus. The user double clicks on a term in the query network and the large network representing term space is displayed with the entry term, or node, near the viewpoint. Within this visual network display the user finds that some of the terms close to the entry term also are of interest in their search, so they copy them to the query network.

An overview of the entire term set is also presented to show which small part of the network is currently being viewed in detail and the other landmark terms in the network. After creating a visual bookmark which can be used to return to the entry point, the user uses the overview to zoom to another region. As they are flown to the new region, many terms are briefly in view. The user sets a bookmark when they arrive at the selected term, and continues searches around the term, using the bookmark to return to a familiar point and search alternative paths.

During the user's travels they copied several terms to the query and now connect them to revise the query network. A new search is initiated and they see that the documents retrieved are mostly those already explored. At this point the user decides to enter the visual network of documents, and selects the document which seems to be the best they have found so far to be the entry point into this large network. Again, an overview of this large network is displayed showing overview documents, the entry point document is just in front of the viewpoint, and the user sets a bookmark so they can return to familiar territory. The document network nodes are labelled with abbreviated titles, and they explore documents close to the entry point by viewing their text in the browser and their visual network representations in Document Explorer. Through this exploration the user identifies documents which are relevant to the search and saves them, as well as finding more terms in the single documents' networks which are useful in further modifying the query network.

The user performs another search with the modified query network, noting that some new documents are found. These are examined, and some are saved. Further explorations into the term and document networks are made, using both new results and returning to the familiar bookmarks. In this process the user has explored both relations among terms and documents to gain a measure of knowledge about the resources available through the Internet. Satisfied that they have found the most relevant documents, or that it is now time to move to another task, the user ends the session.

5. Conclusion and future directions

The Document Explorer supplies visualization, browsing, and query formulation mechanisms based on the semantic content of WWW documents. Users can view and interact with a visually displayed network of documents based on content similarity in a WWW information space that is an alternative to link based representations. Relationships among individual keywords in the documents are also displayed visually to support query formulation by direct manipulation and convey information about the keyword set. Navigation and orientation tools facilitate interaction and enhance perception of the document set and term collection structures through visual representations which facilitate the exploration of local detail while remaining oriented to the global context. These facilities can serve as useful additions to textual representations and visual representations based on document link structure.

One of the tenets of information retrieval is that users' information needs are most likely to be met when multiple paths of access to information are provided. In practice the suite of visualization and query formulation tools provided by the Document Explorer are of most utility in users' search when coupled with other tools, e.g. link based visualizations and conventional search mechanisms. Toward the end of supplying an integrated set of tools for retrieval we are currently developing in integrated suite of tools, the Semantic Space Viewer (SSV), which seamlessly combines the use of varying information visualization and content extraction techniques. Its architecture follows from the delineation of information visualization components suggested in Section 2. In using the SSV users select among techniques for structure derivation, e.g. network or hierarchical clustering, to be presented in a display space, e.g. 2- or 3D perspective, or 3D parabolic, and make use of various display and orientation mechanisms which suit their needs, e.g. stereoscopic or monoscopic display, overviews of differing complexity, etc. The goal of the SSV is to provide a testbed for exploring the design space of information visualization through a domain independent suite of techniques and tools.

6. Acknowledgements

Work on this project has been supported by NASA grants NAG9-551 and NAG9-842.

References

1. B. R. Schatz and J. B. Hardin 1994. NCSA Mosaic and the World-Wide Web: Global hypermedia protocols for the Internet. *Science*, **265**, 895–901.
2. K. Andrews and E. Kappe 1994. Soaring through hyperspace: a snapshot of Hyper-G and its Harmony client. *Proceedings of Eurographics Symposium of Multimedia/Hypermedia in Open Distributed Environments*, 181–191. Berlin: Springer Verlag.
3. P. Ingwerson and I. Wormwell 1986. Improved subject access, browsing and scanning mechanisms in modern online ir. *Proceedings of ACM SIGIR*, 68–76. New York: ACM.
4. M. E. Maron and J. L. Kuhn 1960. On relevance, probabilistic indexing, and information retrieval. *Journal of the Association for Computing Machinery*, **7**, 216–244.
5. M. J. Bates 1986. Subject access in online catalogs: a design model. *Journal of the American Society for Information Science*, **37**, 357–386.
6. M. P. McCahill and T. Erickson 1995. Design for a 3D spatial user interface for Internet Gopher. *Proceedings of ED-MDIA 95*, 39–44. Charlottesville, VA: AACE.
7. K. Andrews 1995. Visualizing cyberspace: Information visualisation in the Harmony Internet browser. *Proceedings of Information Visualization*, 97–104. Los Alamitos, CA: IEEE.
8. R. J. Hendley, N. S. Drew, A. M. Wood and R. Beale 1995. Narcissus: Visualizing information. *Proceedings of Information Visualization*, 90–97. Los Alamitos, CA: IEEE.
9. T. Munzner and P. Burchard 1996. Visualizing the structure of the World Wide Web in 3D hyperbolic space. Available at http://www.geom.umn.edu/docs/weboogl/, The Geometry Center, University of Minnesota.
10. S. Mukherjea, J. D. Foley and S. Hudson 1995. Visualizing complex hypermedia networks through multiple hierarchical views. *Proceedings of CHI '95*, 331–337. New York: ACM.
11. S. Benford, S. Snowden, C. Greenhalgh, R. Ingram, I. Knox and C. Brown 1995. VR-VIBE: a virtual environment for co-operative information retrieval. *Eurographics '95*, 349–360.
12. M. Hemmje, C. Kunkel and A. Willett 1994. LyberWorld—a visualization user interface supporting fulltext retrieval. *Proceedings of ACM SIGIR*, 249–259. New York: ACM.
13. N. D. Gershon, J. LeVesseur, J. Winstead, J. Croall, A. Pernick and W. Ruh 1995. Visualizing Internet sources. *Proceedings of Information Visualization*, 122–128. Los Alamitos, CA: IEEE.
14. N. D. Gershon 1994. (Panel chair.) Information visualization: The next frontier. *ACM SIGGRAPH '94 Conference Proceedings*, 485–486. New York: ACM.
15. G. G. Robertson, S. K. Card and J. D. Mackinlay 1993. Information visualization using 3D interactive animation. *Communications of the ACM*, **36(4)**, 57–71.
16. C. M. Bowman 1994. The Harvest information and discovery system. *Proceedings of the Second International WWW Conference, WWW '94*, 763–771.
17. G. Salton 1989. *Automatic text processing: the transformation, analysis, and retrieval of information by computer*. New York: Addison-Wesley.
18. K. Sparck Jones 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation*, **28(1)**, 11–20.
19. N. J. Belkin and B. H. Kwasnik 1986. Using structural representations of anomalous states of knowledge for choosing document retrieval strategies. *Proceedings of ACM SIGIR*, 11–22. New York: ACM.
20. W. B. Croft and R. H. Thompson 1987. I3R: A new approach to the design of document retrieval systems. *Journal of the American Society for Information Science*, **38**, 236–261.
21. D. W. Dearholt and R. W. Schvaneveldt 1990. Properties of Pathfinder networks. In R. W.

Schvaneveldt (Ed.), *Pathfinder associative networks: Studies in knowledge organization*, 1–30. Norwood, NJ: Ablex.

22. R. W. Schvaneveldt, D. W. Dearholt and F. T. Durso 1989. Network structures in proximity data. In G. H. Bower (Ed.), *The psychology of learning and motivation: advances in research and theory*, 249–284. New York: Academic Press.
23. N. M. Cooke, F. T. Durso and R. W. Schvaneveldt 1986. Recall and measures of memory organization. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **12**, 538–549.
24. R. W. Schvaneveldt (Ed.) 1990. *Pathfinder associative networks: Studies in knowledge organization*. Norwood, NJ: Ablex.
25. G. Di Battista, P. Eades, R. Tamassia and I. G. Tollis 1994. Algorithms for drawing graphs: an annotated bibliography. *Computational Geometry*, **4**, 235–282.
26. R. Tamassia, G. Di Battista and C. Battini 1988. Automatic graph drawing and readability of diagrams. *IEEE Transactions on Systems, Man, and Cybernetics*, **18(1)**, 61–79.
27. P. Eades and L. Xuemin 1989. How to draw a directed graph. *IEEE Workshop on Visual Languages*, 13–17. Los Alamitos, CA: IEEE.
28. C. Batini, E. Nardelli and R. Tamassia 1986. A layout algorithm for data-flow diagrams. *IEEE Transactions on Software Engineering*, **12**, 538–546.
29. E. Makinen 1988. On circular layouts. *International Journal of Computing Mathematics*, **24**, 29–37.
30. L. A. Rowe, M. Davis, E. Messinger, C. Meyer, C. Spirakis and A. Tuan 1987. A browser for directed graphs. *Software Practice and Experience*, **17(1)**, 61–76.
31. P. Eades 1984. A heuristic for graph drawing. *Congressus Numerantium*, **42**, 149–160.
32. C. Esposito 1988. Graph graphics: Theory and practice. *Computer Mathematics Applications*, **15**, 247–253.
33. A. Kumar and R. H. Fowler 1994. *A spring modeling algorithm to position nodes of an undirected graph in three dimensions*, NAG9-551-4, Department of Computer Science, University of Texas – Pan American, Edinburg. Available at http://www.cs.panam.edu/info_vis.
34. T. Kamada and S. Kawai 1989. An algorithm for drawing general undirected graphs. *Information Processing Letters*, **31**, 7–15.
35. M. D. Apperly, I. Tzavaras and R. Spence 1982. A bifocal display technique for data presentation. *Proceedings of Eurographics '82*, 27–43.
36. C. Ware and G. Franck 1996. Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics*, **15**, 121–140.
37. K. W. Arthur, K. S. Booth and C. Ware 1993. Evaluating 3D task performance for fish tank virtual worlds. *ACM Transactions on Information Systems*, **11**, 239–265.