Information Navigator: An information system using associative networks for display and retrieval*

Richard H. Fowler, Bradley A. Wilson, and Wendy A. L. Fowler

Report NAG9-551, No. 92-1

Department of Computer Science
University of Texas - Pan American
Edinburg, TX 78539-2999
email: fowler@panam.edu

Abstract

Document retrieval is a highly interactive process dealing with large amounts of information. Visual representations can provide both a means for managing the complexity of large information structures and an interface style well suited to interactive manipulation. The system we have designed utilizes visually displayed graphic structures and a direct manipulation interface style to supply an integrated environment for retrieval. A common visually displayed network structure is used for query, document content, and term relations. A query can be modified through direct manipulation of its visual form by incorporating terms from any other information structure the system displays. An associative thesaurus of terms and an inter-document network provide information about a document collection that can complement other retrieval aids. Visualization of these large data structures makes use of fisheye views and overview diagrams to help overcome some of the inherent difficulties of orientation and navigation in large information structures.

* This research was supported by NASA grant NAG9-551 and a University of Texas - Pan American Faculty Research Council grant.

Table of Contents

1. Introduction

1.1. Visualization

Visualization is emerging as a central concern for a wide range of tasks. Visualization brings the human visual system with its pattern recognition capabilities to bear in either discovery or portrayal of information (McCormick, DeFanti, & Brown, 1987).

In scientific visualization large complex data sets are displayed in ways allowing the investigator to view the global nature of numerical solutions and visually explore analyses (Nielson, 1989).Visual programming systems express algorithms in ways that might improve clarity and simplicity of expression (Chang, 1990; Shu, 1988), as well as manage complexity in large systems (Graf, 1990). Knowledge engineering environments often supply graphic representations of rules and facts (Mettrey, 1987). Even hypertext systems with entirely textual content, typically rely on visual representations for navigation and orientation (Conklin, 1987; Nielsen, 1990).

One of the common goals of visualization is to help manage and understand large amounts of data or information. From this perspective information retrieval systems for large text databases are natural candidates for visualization techniques (Card, Mackinlay, & Robertson, 1991; Crouch & Korfhage, 1989). Indeed, some thirty years ago Doyle (1961) explored the use of visual representations for several components of an information retrieval system at the same time that Sutherland (1963) pioneered work in graphic display and manipulation. Recent work on visualization in information retrieval has tended to focus on individual elements of document retrieval systems. A number of systems have included graphic thesaurus displays (Bertrand- Gestaldy, 1986; Frei & Jauslin, 1983; McMath, Tamaru, & Rada 1989), and visual representations of inter-document relations (Crouch, Crouch, & Andreas, 1989; Sammon, 1969). Work on query formation and modification has explored visual interfaces for Boolean queries (Anick et al., 1990; Burgess & Swigger, 1986) and relational database queries (Angelaccio, Catarci, & Santucci, 1990; Chang, 1989), as well as queries constructed as term vectors (Crouch, 1986) and networks (Consens & Mendelzon, 1989).

However, in text retrieval systems the *integration* of visualization techniques supporting search is yet to be fully exploited. This is perhaps surprising - today's widespread end-user searching of bibliographic databases provides a population of naive users that might particularly benefit from visual representations. The integration of a common visual representation and interaction style for retrieval aids throughout the retrieval process has been one of our principal goals (Fowler & Dearholt, 1990).

1.2. Different Information Needs Require Different Tools

Retrieval systems must supply mechanisms for meeting quite different information needs as well as support the complex behavior of user information seeking (Rouse & Rouse, 1984). Information needs can be categorized as those arising when 1) direct bibliographic access is required, 2) the domain is well-known to the user, and 3) the domain is not known to the user (Ingwerson & Wormell, 1986). Due in part to the need to provide the user information about indexing and the subject domain, different retrieval techniques and aids are appropriate for meeting different types of information needs. One means to overcome the limitations of a single technique is to enhance a system's flexibility by providing the user a variety of tools and retrieval techniques (Jones & Furnas, 1987) that provide multiple paths of access to information (Bates, 1986).

Various methods for retrieval have been investigated to overcome some of the well-known shortcomings of techniques based on a single best match of a user's query and document representations (Belkin & Croft, 1987). Automatic query refinement based on relevance feedback (Maron & Kuhn, 1960; van Rijsbergen, 1979) and the use of expert systems (Fidel, 1986; Vickery, 1984; Watters, Sheperd, & Robertson, 1987) can automate some of tasks a user or search intermediary would perform. More often, indexing aids such as a thesaurus together with the feedback provided by search results play the largest roles in allowing the searcher to refine his or her query. Regardless of the system aids available to the user, query refinement is an iterative process. The iterative nature of information retrieval, together with the need for system flexibility, require an environment in which interaction is natural and straightforward. Another goal of our system is to develop such an environment supplying a range of alternative retrieval aids.

1.3. Direct Manipulation and Visual Representation

Accounts of human-computer interaction provide a number of insights into the design and operation of visual systems (Carroll, 1987; Laurel, 1990; Norman & Draper, 1986) - systems incorporating direct manipulation (Shneiderman, 1983) in the interface. The essential characteristic of direct manipulation systems is that changes in the underlying system components are reflected in visual changes in the interface objects representing those objects. The user operates on the visual interface objects to effect change in the system state. In the ideal interface the user feels directly engaged with the underlying task, rather than with an interface which in turn directs the change of state. Among factors contributing to a feeling of direct engagement with systems objects are the cognitive effort required to change the system state, and evaluate a resulting state. To provide an environment for the visual manipulation of system objects it is necessary that the output of one process serve as the input to another process: inter-referential input/output (Draper, 1986). In a visual environment for information retrieval, components for query formulation, retrieval aids, document representations, and as many system components as possible, should share a common representation. This shared visual representation would then allow inter-referential input/output to support the direct manipulation of information objects.

2. A system Integrating Visual Representations

The system we have developed provides the user a visual environment for direct manipulation during information retrieval. A common visual representation is used for query, associative thesaurus, conventional thesaurus and document content.

Perhaps the easiest way to get a feel for the way the system operates as a whole is to look at a series of displays and describe the interactions the user might have. One component is an associative thesaurus formed as a network of the terms used in matching documents to users queries. An overview diagram of the complete network is displayed in the upper left of the screen, as shown in Figure 1. When parts of the thesaurus are viewed, the overview diagram shows where the smaller part being viewed is located. The overview diagram also serves to give the user a feel for the most frequently occurring terms in the database which is currently selected.

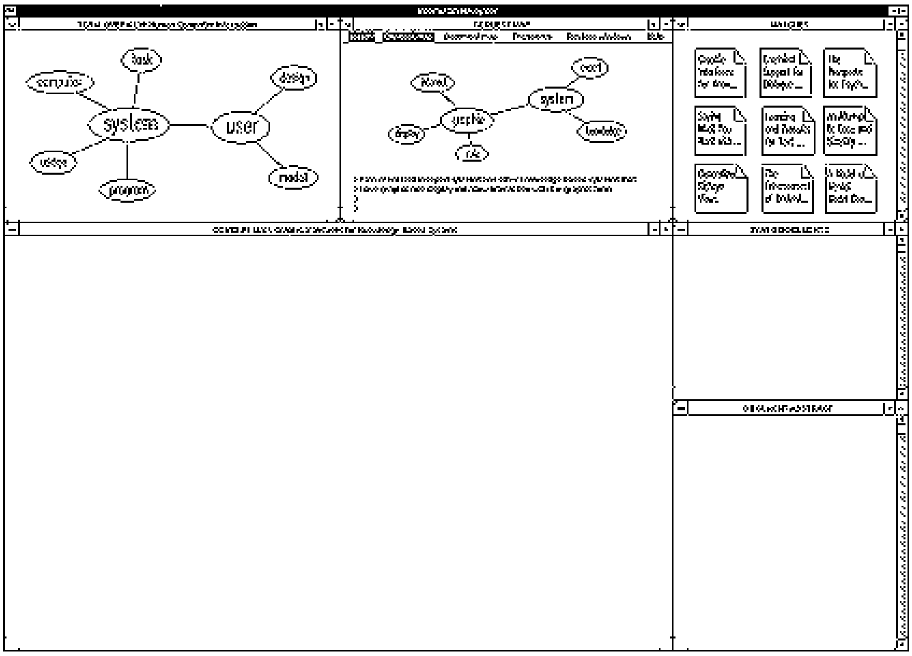

Figure 1. Display after entering a natural language query and performing a search. The upper left window displays an overview diagram of the associative term thesaurus. The text of a natural language query with its visual representation is displayed in the upper middle window. The upper right window contains an ordered list of documents retrieved in the search.

The query process can begin with the user’s entry of a natural language request for information. Term analysis (detailed below) is used to transform the text to a manipulable visual representation, as displayed in the upper middle window. Search results are shown in the upper right window as icons representing documents. The icons are ordered by how closely they match the query.

The query can be revised by several means as the user views the results of successive searches. The user might enter more text to refine the query. The system will then reanalyze the text and display a revised query network. Alternatively, the user can interact directly with the system’s visual representations and add nodes to the query by "dragging" any term the system displays into the query window and connecting it to the query or the user can delete nodes from the query network.

Selecting and dragging a document to the Saved Document window places it on a list of documents. A document icon can also be dragged to the Document Abstract window to display the text of the document’s abstract, as shown in Figure 2.

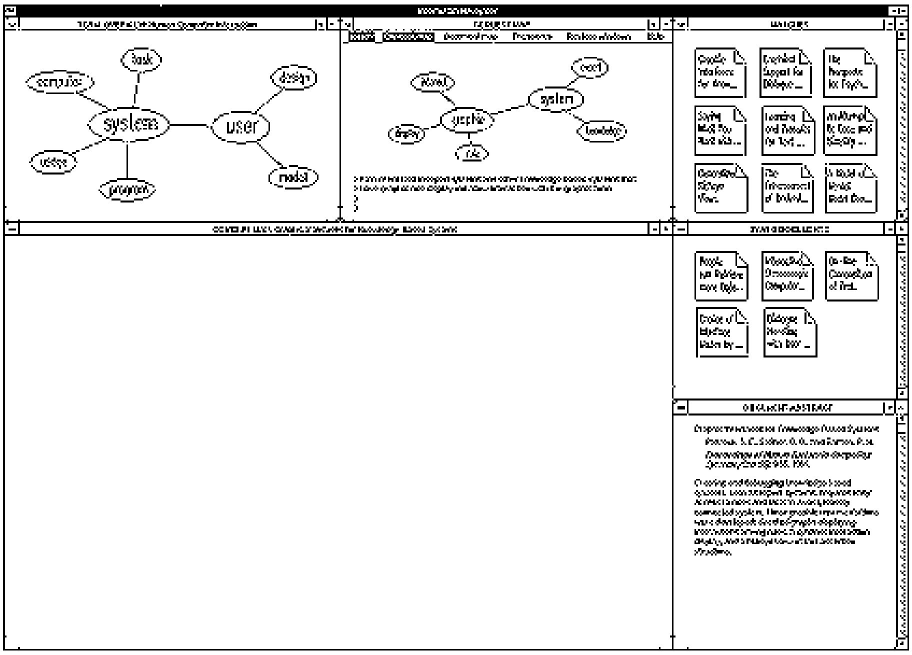

Figure 2. Display after user saves documents in the right middle window selected from search results. Dragging a document icon into the right bottom window displays the text of the document's abstract.

A document icon can also be dragged into the large window to display the abstract's Concept Map, or term network, as shown in Figure 3. These terms from the document abstract network can also be incorporated into the query. Terms in the query that match terms in the document abstract are shaded to give a visual presentation of the degree of match.

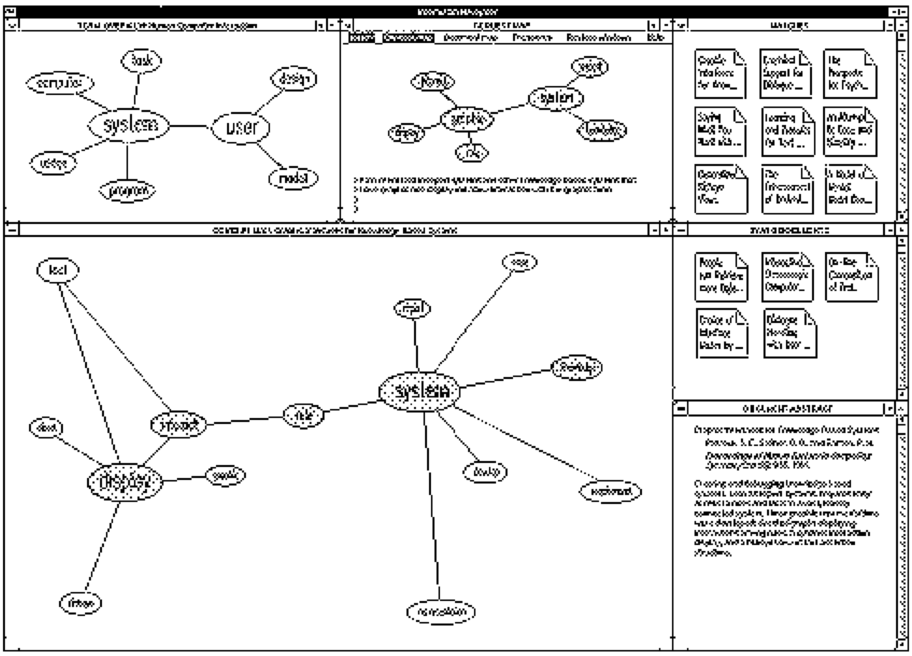

Figure 3. Dragging a document icon into the right bottom window displays the document's Concept Map, or term network. Terms occurring in both the query and abstract networks are shaded.

An associative network of the terms in the database, shown in Figure 4, is also a source of terms for query revision. To display a view of the associative thesaurus, the Concept Map button is selected and a term is dragged to the concept map window. A part of the larger associative term network is then displayed centered, or focused on that term. Here, the term "display" was dragged into the Concept Map window and a fisheye view of the associative thesaurus, focused on "display", has been be displayed

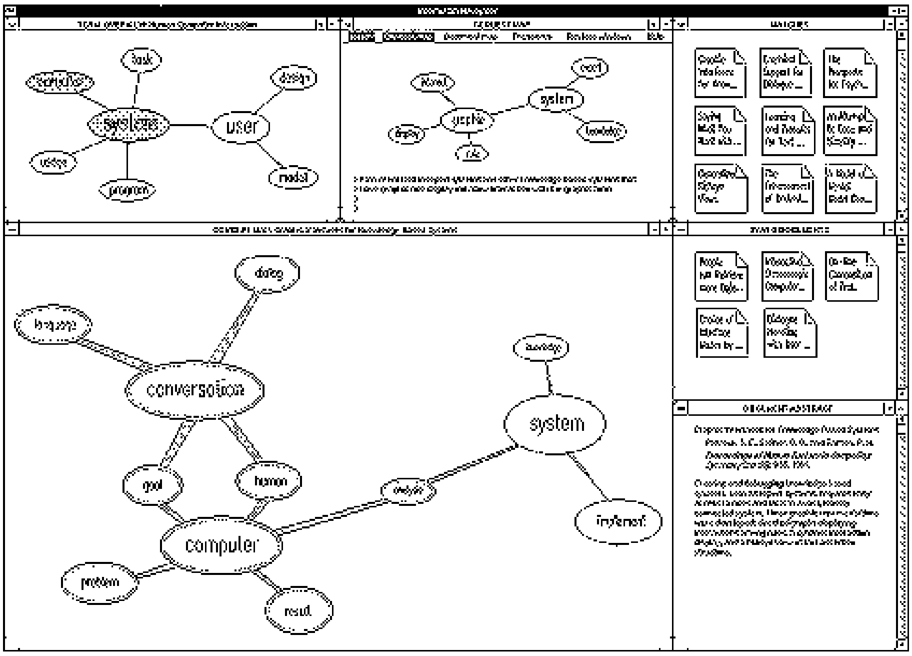

Figure 4. A fisheye view of the associative thesaurus is shown in the large window at the lower left. The overview in the upper left window is shaded to show the location of the portion being viewed.

The overview of the thesaurus shown in the upper left window is constructed using the nodes with highest degree in the complete network. The nodes serve to give a quick overview of the terms in the database, and also serve as navigation aids. The location of the small portion of the associative thesaurus in the complete network is shown by shading the overview where it appears. The user can move about in the network by selecting a node. A new view of the thesaurus is then displayed with the selected node as the focus.

Finally, a network reflecting interdocument relations is also available for browsing and retrieval (Figure 5). To display the network of documents, the user selects the Document Map button. Then, a document icon can be dragged to this window, and a view of the network of documents is displayed focused on that title. The same navigation mechanisms are used as with the network of terms. The overview diagram now is based on the documents with highest degree and again shows the relation of the subgraph displayed.

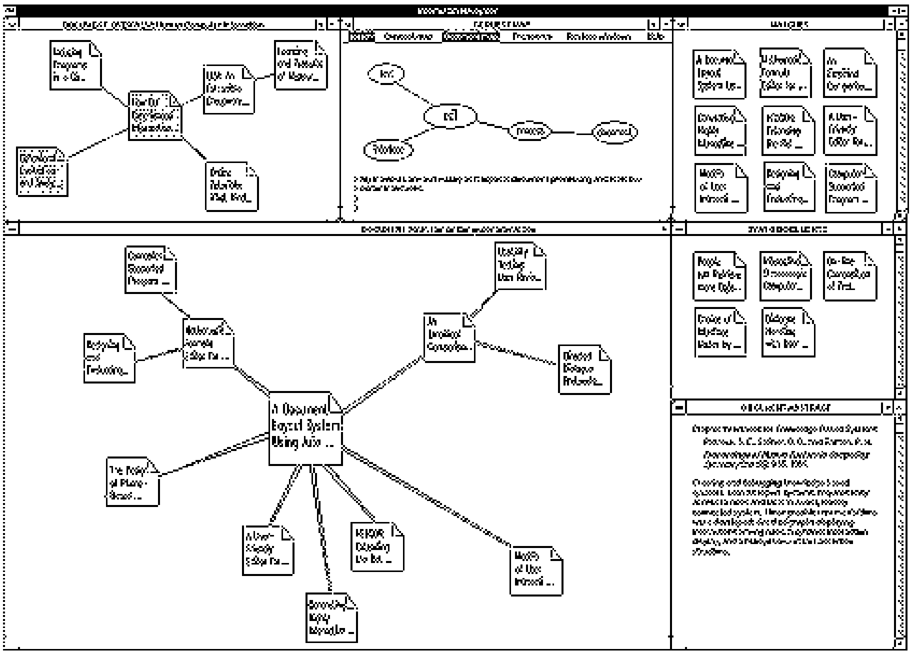

Figure 5. A fisheye view of the network of documents is shown, centered on the document which was dragged into the window. The overview orients the subgraph being viewed to the complete document collection.

2.1. User Interaction

One of the goals of a direct manipulation interface is to reduce the user’s cognitive load to enhance a feeling of engagement in the task. In the system all interaction is performed either by selecting one of the buttons displayed in the top middle window or dragging one of the information items displayed as an icon. For example, a document icon can be dragged to the SAVED DOCUMENT window to place the document in a list of titles, to the CONCEPT MAP window to display the network of terms in the document, or to the DOCUMENT ABSTRACT window to view the text of the abstract. As mentioned above, a term icon is similarly manipulated to form a query and can also be used to define a view of the associative thesaurus or hierarchical thesaurus.

Display of the large networks reflecting document interrelations and term associations is based on fisheye views and the user navigates by selecting document or term nodes to define the center of the fisheye. With the display of Figure 4 the user might select the Concept Map button and then one of the terms from the query, document, or overview. The document network would then be replaced with a view of the associative thesaurus centered on the selected term.

2.2. The System’s Associative Networks

The representations underlying the visual displays are minimum cost networks derived from measures of term and document associations. For queries, document abstracts, and associative term thesauri the associations are derived from natural language text. The network of documents is based on interdocument similarity. 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 relatively well-known (Salton, 1983).

2.2.1. Why associative structures?

For the system we have implemented there are three reasons for using statistically-based associative structures. 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 the 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 can also serve as one component of a hybrid system incorporating both deep and shallow representations (Croft & Thompson, 1987). Longer term, we intend to explore a range of visual representations for information structures. 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 several information retrieval system components.

2.2.2. Deriving the term associations

For each database the system uses a separate set of terms that includes 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 (Sparck Jones, 1972). However, one function of the associative thesaurus is to give a picture of the all 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 analyzed 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 (Belkin & Kwasnik, 1986). 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.

3. Pathfinder Networks: A Class of Minimum Cost Networks

The associative networks used in the system are Pathfinder networks (PFNETs). The Pathfinder algorithm was developed to model semantic memory in humans and to provide a paradigm for scaling psychological similarity data (Schvaneveldt, Dearholt, & Durso, 1989). 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 (Cooke, Durso, & Schvaneveldt, 1986; Schvaneveldt, 1990).

PFNETs are derived by identifying the proximities that provide the most efficient connections between entities. This is accomplished by considering the indirect connections provided by paths through other entities. In this respect it is similar to other minimum cost methods for deriving network structures (Hutchinson, 1989; Knoke & Kuklinski, 1983). It is, however, unique in the generality of the family of networks that can be generated (Dearholt & Schvaneveldt, 1990). Computer programs for generating PFNETs are available both as source code (Gerber et al., 1987) and in executable form (Interlink, 1989).

Pathfinder requires as input a measure of distance between each pair of entities in the target domain. These proximity matrices may be symmetrical (distanceij = distanceji) or asymmetrical. Several different methods have been used to obtain these distance estimates in the assessment of PFNETs' effectiveness in revealing conceptual structure. In cognitive modeling tasks the most common has been to have experts judge the similarity or relatedness of all pairs of concepts (Roske-Hofstrand & Paap, 1986). Other methods include sorting the concepts into categories (McDonald, Dearholt, Paap, & Schvaneveldt, 1986), controlled associations (Miyamoto, Oi, Katsuya, & Nakayama, 1986), and various sequence transformations. The sequences have been obtained from user transactions (Anderson, McDonald, & Schvaneveldt, 1987), recall orders (Cooke, Durso, & Schvaneveldt, 1986), event records, and natural language text (McDonald, Plate, & Schvaneveldt, 1990).

3.1. Deriving PFNETs

Conceptually the algorithm is quite simple. Entities in a domain are represented by nodes, and links connecting entities are assigned: (1) weights (from distances in the data matrix) according to their strengths and (2) labels, reflecting their role in constructing networks. The link membership rule assures that links which are a part of some minimum distance path are preserved between each pair of nodes. To derive a PFNET the direct distances between each pair of nodes in the data matrix are compared with indirect distances. A direct link between two nodes is included in the PFNET unless the data contain a shorter path having two or more links. In constructing a PFNET two parameters are incorporated: *r* determines how path weight is computed and *q* specifies the maximum number of links considered in finding a minimum cost path between entities.

Path weight, r, is computed according to the Minkowski r-metric. It is the rth root of the sum of each distance raised to the rth power for all links in a path between two nodes. Although the r-metric is continuously variable, simple interpretations exist only for r = 1 (path weight is the sum of the link weights in the path), r = 2 (path weight is the Euclidean distance), and r = infinity (path weight equals the maximum link weight in the path). One advantage of r = infinity is that one need only assume that the original distance estimates have ordinal properties. Another advantage is that the link structure will be preserved for any monotonic transformation of the data. The second parameter, q, determines the maximum number of links which will be included in a path. Consequently, q also determines the dimensionality of triangle inequalities (Tversky & Gati, 1982) which are not violated. The larger the value of q, the fewer triangle inequalities are violated. If q is one less than the number of nodes, then no triangle inequality is violated.

Pathfinder allows for systematic variation in the complexity (number of links) in the resulting networks as the two parameters are varied. Complexity decreases as either r or q increases. As either parameter is manipulated, links in a less complex network form a subset of the links in a more complex network. Thus, the algorithm generates two orthogonal families of networks, controlled by r and q. The least complex network displays only the most salient relationships and is obtained with r = ¥ and q = n-1, where n is the total number of nodes in the network. This is the PFNET used in the system. It is the union of all minimum cost spanning trees that can be derived for a data set and can be constructed efficiently.

3.2. An Alternative Network Structure

Pathfinder provides an alternative procedure for automatically deriving network representations for information retrieval systems. Proximities can be used to derive both PFNETs and the threshold networks often used in information retrieval, but there are significant differences. Threshold networks include a link in the final network if the magnitude of internode association is above some criterion. Link membership is determined by examining only local proximities, proximities of nodes adjacent in the original data. PFNETs differ in a fundamental way. A link is included in a PFNET depending not on a fixed magnitude of association, but on the role it plays in determining minimum cost paths between nodes. Link membership is determined using a global, path oriented approach that considers other connections, potentially across the entire network. The two approaches lead to quite different link structures, even when the number of links is the same.

Another consideration is the extent of empirical evaluation of networks used in information retrieval in revealing psychologically salient relations in link structure. For uses in which psychological salience is important, such as our desire to have a "natural" representation suited to visual display and interaction that minimizes cognitive load, empirical assessment is important. Though statistically based graph representations used in information retrieval have sometimes assessed psychological salience (Palmquist & Eisenburg, 1984), PFNETs have been developed from the outset as a representational scheme for human conceptual structure based on psychological theory. Work across a number of domains indicates that PFNETs can provide a representation effective in elucidating conceptual structure (Cooke, Durso, & Schvaneveldt, 1986; Schvaneveldt, 1990).

4. Visual Display of Networks

Graph display is an important issue for a number of tasks. Considering the wide application of graphs structures in display, there are relatively few algorithms for drawing general undirected networks (Tamassia, Battista, & Battini, 1988). This is due in part to the difficulty in precisely specifying the perceptual and aesthetic criteria individuals use in understanding graphs (Eades & Xuemin, 1990). Nonetheless, when it is possible to specify graph theoretic expressions of criteria that can be used to guide the viewer’s extraction of information, such as for trees, satisfactory display algorithms can be developed (Batini, Nardelli, & Tamassia, 1986).

The system’s network displays center on visually conveying information about the networks’ edge structure and weights. Several graph theoretic criteria are employed by the display algorithms. *Node degree* is the number of edges incident to a node. This descriptor is used to provide a measure of the density of the network’s link structure. *Path weight* is the weight associated with a series of edges connecting two nodes. *Path distance* is used here to refer to the number of edges connecting two ends of a path in. It is the distance in the edge structure without considering weights.

The system provides displays of both of relatively small networks for query and document abstract terms, and much larger networks for the associative thesaurus and network of documents. In the displays shown earlier of a document abstract network, several visual cues were incorporated to guide the user in viewing the network.

We considered several layout algorithms before arriving at the final display. The simplest was a radial positioning of nodes: starting at the center, nodes are simply laid out on a circle around the node to which they are connected. This is a very efficient method and allows relatively dense displays, though line (or edge) crossings are a problem.

We also looked at multi-dimensional scaling (MDS) as a way to position nodes. The original distance measures used to derive the network were used to find a scaling solution in two dimensions, and this two dimensional MDS solution provided the node positions. This seemed a promising line in that the MDS solution could provide information about global term relations that might complement the local relationships reflected by a minimum cost network. Unfortunately, again line crossings obscured the network relations.

We settled on an algorithm developed by Kamada and Kawai (1989) that considers the nodes of a graph to be connected by virtual springs between all nodes. The algorithm minimizes the overall tension in the system of springs by iteratively repositioning nodes. The strength of the spring reflects the number of edges separating nodes - the network’s edge structure, but not the edge weights.

With spring weight determined as the number of edges separating nodes, interconnections among all nodes are taken into account, and consideration of the overall edge structure leads to improvement over the first two techniques by essentially eliminating line crossings. Figure 6 shows the document abstract of Figure 4 with nodes positioned using this technique based on path distance.

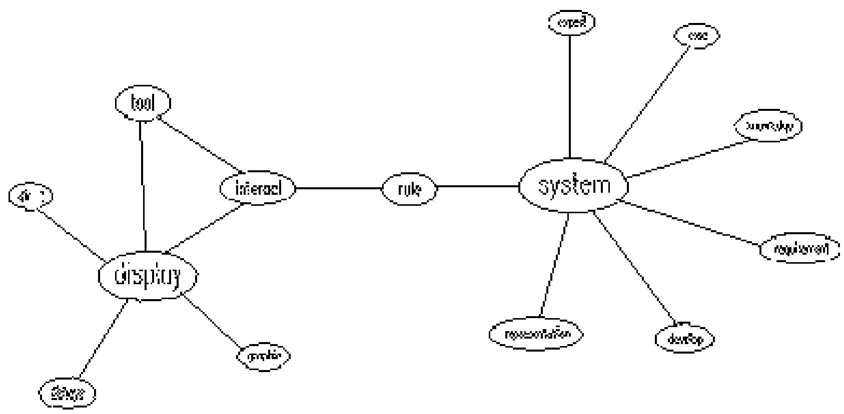

Figure 6. Document abstract network. Nodes positioned by Kamada & Kawai’s spring algorithm with spring strength determined by path distance.

Still, there is additional information in the underlying network: the links’ weights. In the final layout scheme, spring tension is determined as the sum of edge weight, rather than number of edges, between nodes. The placement of nodes in Figure 7 is determined only indirectly through structure, as represented in the sum of network path weights, and incorporates information about strength of association among nodes.

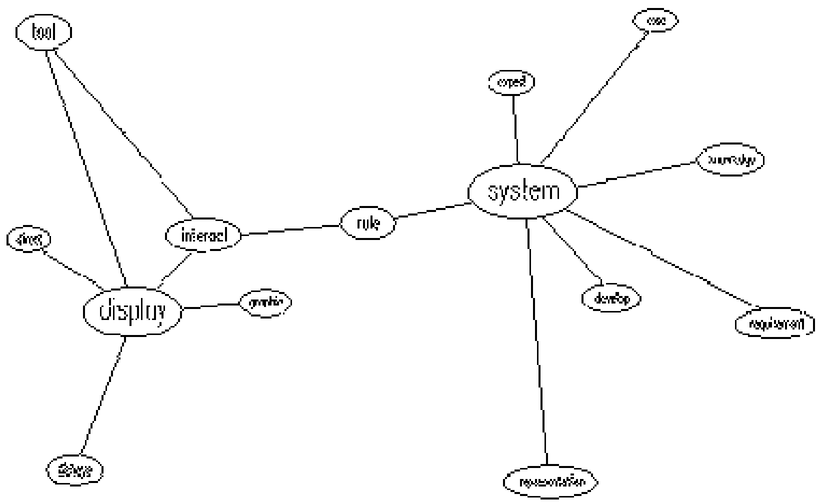

Figure 7. Nodes positioned in document abstract network with spring strength determined by path weight. This layout reflects strength of association as well as edge structure.

Finally, to direct the users attention in scanning the network nodes with high degree are made larger. High degree nodes are visually emphasized by their size. In the context of the complete system, one function of document abstract networks is to allow the evaluation of a document's relevance to the query. To facilitate this judgment, terms that are in both the query and the abstract network are shaded. Both differing node size and node shading are shown in Figure 4.

5. Large Information Structures: Associative Thesaurus and Network of Documents

In addition to the query and document term displays the user can access two other visually displayed network structures: an associative thesaurus of terms, and a network of documents. The associative thesaurus is based on a PFNET of all terms in the database. The distances for deriving this network are found using the same weighted co-occurrence measure used in assigning term distances in documents and queries. All documents are analyzed and an additional value is added to term pair similarity is for terms co-occurring in the same document. For the network of documents, distances between documents are calculated using the same matching algorithm used to assess query-document similarity. Network similarity is calculated by combining the number of commons terms with a measure of structural similarity for these common terms (Goldsmith & Davenport, 1990).

5.1. Orientation and Navigation in Large Information Structures

With the relatively small networks for queries and document abstracts it is possible to display the complete networks. For the much larger networks of all terms and all documents only part of the complete networks can be displayed, and mechanisms allowing the user to browse, or navigate, in the large network are required. It is also useful to relate the small portion of the network being viewed to the complete network to provide the user orientation within the overall structure. Orientation and navigation are challenges shared by hypertext and other large information spaces. The navigation and orientation mechanisms used in this system are based on overview diagrams of the complete network structure and fisheye views of the detailed network view.

5.2. Overview Diagrams

Overview diagrams are a common means of supplying a user with (1) knowledge about the organization of the complete network, (2) a means for navigating the network, and (3) orientation within the complete network. In overview diagrams a small number of nodes, selected to provide information about the organization of the complete network, are displayed to the user. Additionally, the nodes typically provide entry points for traversing the network. These nodes provide orientation by serving as landmarks to assist the user in knowing what part of the network is currently being viewed.

In the 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. There are relatively few nodes of high degree and the diameter of the network is small. This form of network suggests criteria for selecting nodes to include in overview diagrams. Overview diagrams displayed by the system include those nodes of highest degree in the complete network. The overview is displayed using the same techniques employed for query and document term networks.

To display a view of the complete network of terms the user selects a term and then selects the Concept Map button. Figure 8 shows the screen when node labeled "conversation" is selected. Entry points can be selected from the query, overview diagram, or a document's graph. Here, the subgraph of the complete term network is displayed centered, or focused, on "conversation".

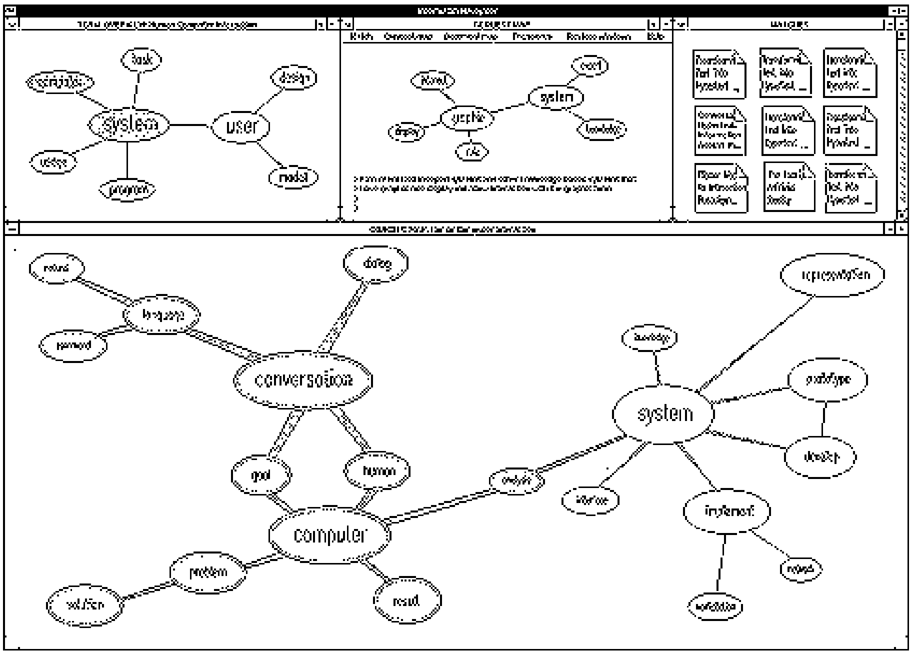

Figure 8. Information Navigator display focused on "conversation" in the associative thesaurus displayed in the central Network window. The aerial view orients the subgraph to the complete thesaurus by tracking aerial view nodes present in the subgraph.

The terms selected for overview diagrams of associative thesauri tend to be general terms that provide a guide to the content of the database. They are landmarks in that they supply information about both the content and structure of the database. For example, in Figure 1 "system" is a central term in that it has the most connections to other terms. Its' high degree reflects frequent co-occurrence with many other terms in the document collection.

The overview diagram can also be used for navigation. By selecting one of the overview terms a fisheye view centered on the selected node can be displayed. In this example a term from the query was selected to serve as the focus of the fisheye view displayed in the large window. The location of the subgraph in relation to the overview diagram is indicated by shading any node in the overview diagram that is displayed in the fisheye view. If no term from the overview diagram is present in the detailed view, the overview term closest to the center of the fisheye as measured by path distance is shaded.

5.3. Fisheye Views

Furnas' description of fisheye views in display (Furnas, 1986) provides a framework that is general enough to account for a number of factors important for navigation and orientation in large information spaces. A fisheye view displays objects close to the current viewing point, or focus, in more detail than things farther away. The display of elements depends on both the viewer's distance from an element and the *a priori* importance assigned element. The *a priori* importance for each element together with the distance of each element from the viewing point to each element allows the size (or any factor) of all elements to be related to the view point. View point and *a priori* interest are related by a degree of interest function. For example, the degree of interest function might simply compute the product of *a priori* interest and distance. The degree of interest function supplies a general mechanism for algorithmically providing orientation landmarks within information structures (Fairchild, Poltrock, & Furnas, 1988; Godin, Gecsei, & Pichet, 1989) that has been applied effectively in large information spaces (Valdez & Chignell, 1988).

Though the account of fisheye views is usually considered when dealing with changing viewpoints, overview diagrams might also be characterized as fisheye views. For example, in a hypertext overview diagram the author includes those nodes that the he or she determines are of high interest in that they are useful as navigational entry points and landmarks for orientation in the complete network of text. They are nodes of high a priori interest and should always be displayed in the same way. In our system the overview diagrams is constructed automatically by selecting nodes in the complete networks of highest degree. They are of high interest as landmarks in that they supply 1) a set of the general terms in a database and 2) navigation points that identify the most dense parts of the network. For overview diagrams the user's viewing point can be considered to be outside the network and remain constant, so that the size or detail of the overview nodes are the same and do not change.

5.4. The user's view of network detail

The system's display of network detail within the complete network is based on a fisheye view. To display a fisheye view the user first selects a term to define the focus and then selects the Concept Map button. Figure 8 shows the screen when the node labeled "conversation" is selected from the query graph. Entry points can also be selected from the overview diagram or a document's graph.

For the fisheye display a node's *a priori* importance is its degree. Viewing distance is calculated as the path weight of the focus to other nodes. Using these two measures, the system's degree of interest function yields a value for each node in the network. Degree of interest values are computed previously and stored as a table. A threshold degree of interest value is used and nodes with values above this criterion are displayed. The value of the threshold reflects the size of the display window, so that an appropriate number of nodes will be shown. Having determined which nodes to display, the node layout algorithm for small networks is used to position nodes.

The visual form of nodes and edges conveys additional information about the relations among nodes in the fisheye view. The degree of a node is reflected in its size. Nodes that have high values are larger, and so are more prominent in the display. Additionally, the focus node is drawn as large as the largest node in the view. The large, high degree, nodes are the nodes the user can explore to find the densest parts of the network.

The shape of the edges also conveys information about the structure of the network. The edges are widest at the focus node and narrow as they connect nodes that are farther away from the focus. However, the measure of distance that this narrowing reflects is not path weight, which was used in computing the degree of interest function. Instead, the narrowing of edges reflects path distance, the number of edges connecting nodes. This is useful information for navigation because not all nodes on the path will be displayed due to the degree of interest threshold criterion for display. As users browse the network by selecting new focus nodes from the fisheye view, they can use these cues about network structure to guide their exploration.

When a node in view is selected and made the new focus, a number of nodes from the earlier view will be included in the new fisheye view. The high degree nodes remain displayed and supply useful orientation information. The high degree nodes of the previous display provide a context that is elaborated by the change of focus and the recalculation of degree of interest values based on the new distances. Additionally, the network overview reflects the new focus position.

5.5. Network of documents

Browsing among shelves in a library is exploration and search among documents guided by the classification system used in the library. Within an information system more search and organizational flexibility is possible. Access to documents can be based on a network of documents derived from interdocument relations. From some entry point the network can be traversed and documents selected. This technique has the advantage for some types of information needs of requiring little query formulation and knowledge of the subject area.

The final visual structure the user interacts with is a network reflecting interdocument relations, shown in Figure 9. To construct the PFNET of document titles distances are calculated between all document pairs using the same metric used for query-document matching. In this display nodes from the network of documents are shown as document icons labeled with abbreviated titles. An overview diagram is constructed and displayed in the same manner as the overview of the associative thesaurus. The network of documents can serve the same browsing function as the associative thesaurus. The same mechanisms for navigation and fisheye display are used. Manipulating the nodes of the document network, as with any document icon displayed in the system, allows the user to save the title or view the document’s abstract text and term network. Functions using the PFNET of documents provide an additional means of gaining domain knowledge and moving among visual structures.

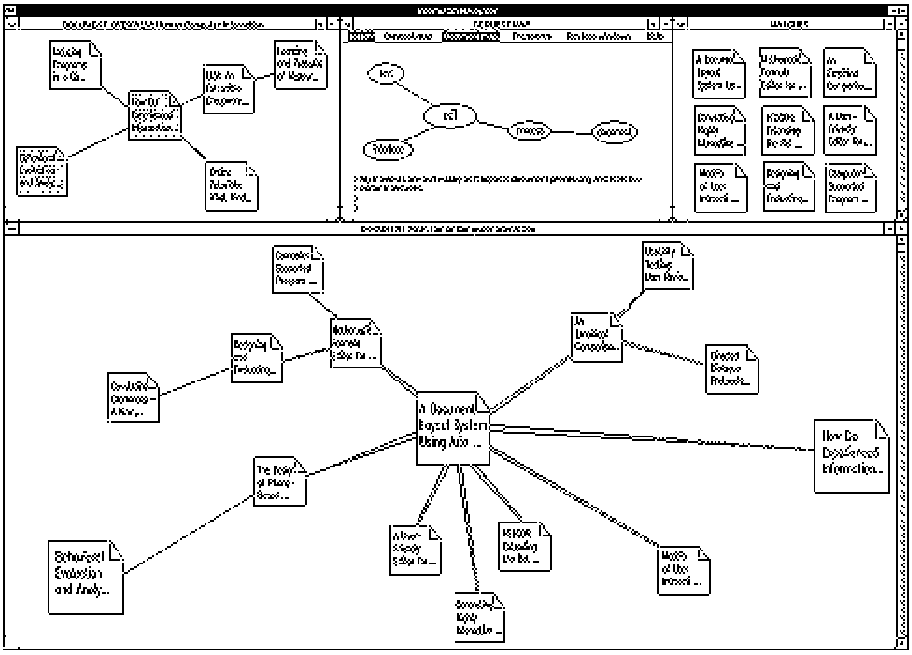

Figure 9. System with fisheye view and overview of the document collection network displayed. Nodes in the network are selectable document icons with abbreviated titles. The result of a search performed by traversing the document collection PFNET is displayed in the upper right window.

In addition to the retrieval mechanism ordering documents by similarity to a query, a second form of retrieval is available. Documents can be retrieved using the network of documents by traversing the network starting at some entry point document. The entry point can be directly provided by the user by selecting a document icon, or determined by the system as the document that best matches the query. Additional documents are then retrieved by following the edges from the starting point in the order of a breadth first search. The sequence of retrieved documents displayed to the user is ordered by the number of edges from the entry point document.

In addition to the retrieval mechanism ordering documents by similarity to a query, a second form of retrieval is available. Cluster based retrieval uses the PFNET of documents and is based on traversing the network beginning at a particular document. The entry point can be directly provided by the user by specifying a title, or determined by finding the document that best matches the query. Additional documents are then retrieved by following the links from the starting point. The sequence of retrieved documents displayed to the user is ordered by the number of links from the entry point document.

6. Conclusion

Our principal goal has been to provide an environment for information retrieval integrating system components through a visual representation allowing direct manipulation. The system focuses on interaction techniques to facilitate query modification and browsing in large information structures. In deriving visual displays of network structures for the user, a number of issues concerned with graph layout, navigation, and orientation were addressed. Furnas' account of fisheye viewing supplied a general orientation to display and a technique for managing large structures.

7. Acknowledgments

This research was supported by NASA grant NAG9-551 and a University of Texas - Pan American Faculty Research Council grant. John Spicer programmed many of the text analysis and network generation programs.

8. References

Anderson, M. P. McDonald, J. E. & Schvaneveldt, R. W. (1987). Empirical user modeling: Command usage analyses for deriving models of users. *Proceedings of the Human Factors Society 31st Annual Meeting*, 41-45, Santa Monica, CA: Human Factors Society.

Angelaccio, M., Catarci, T. & Santucci, G. (1990). QBD: A fully visual system for E-R databases. *IEEE Workshop on Visual Languages*, 56-61.

Anick, P. G., Brennan, J. D., Flynn, R. A., Hanssen, D. R., Alvey, B. & Robbins, J. M. (1990). A direct manipulation interface for Boolean information retrieval via natural language query. *Proceedings of the ACM SIGIR International Conference on Research and Development in Information Retrieval*, 135-150.

Bates, M. J. (1986). Subject access in online catalogs: A design model. *Journal of the American Society for Information Science, 37(6),* 357-386.

Batini, C., Nardelli, E. & Tamassia, R. (1986). A layout algorithm for data-flow diagrams. *IEEE Transactions on Software Engineering*, 12(4), 538-546.

Belkin, N. J., & Croft, W. B. (1987). Retrieval techniques. *Annual Review of Information Science and Technology, 22,* 109-146. Amsterdam: Elsevier.

Belkin, N. J., & Kwasnik, B. H. (1986). Using structural representations of anomalous states of knowledge for choosing document retrieval strategies. *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*, 11- 22.

Bertrand-Gestaldy, S., & Davidson, C. H. (1986). Improved design of graphic displays in thesauri - through technology and ergonomics. *Journal of Documentation, 42,* 225-251.

Burgess, C., & Swigger, K. (1986). A graphical interface for casual, naive users. *Information Processing and Management, 22(6),* 511-521.

Card, S. K., Mackinlay, J. D., & Robertson, G. D. (forthcoming). The information visualizer: An information workspace. *Proceedings of Human Factors in Computing Systems (CHI '91),* New York: ACM.

Carroll, J. M. (Ed.). (1987). *Interfacing thought.* Cambridge, MA: MIT Press.

Chang, S. K. (1989). Visual reasoning for information retrieval from very large databases. *IEEE Workshop on Visual Languages*, 1-6.

Conklin, J. C. (1987). Hypertext: An introduction. *IEEE Computer*, *7*, 17-41.

Consens, M. P. & Mendelzon, A. O. (1989). Expressing structural hypertext queries in Graphlog. *Hypertext ’89 Proceedings*, NY: ACM Press. 269-292.

Cooke, N. M., Durso, F. T., & Schvaneveldt, R. W. (1986). Recall and measures of memory organization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12(4),* 538-549.

Croft, W. B., & Thompson, R. H. (1987). I3R: A new approach to the design of document retrieval systems. *Journal of the American Society for Information Science, 38,* 236-261.

Crouch, D. (1986). The visual display of information in an information retrieval environment. *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*, NY: ACM Press, 58-67.

Crouch, D. B., Crouch, C. J., & Andreas, G. (1989). The use of cluster hierarchies in hypertext information retrieval. *Hypertext ’89 Proceedings*, NY: ACM Press. 225-237.

Crouch, D. B. & Korfhage, R. (1989). The use of visual representations in information retrieval applications. In R. Korfhage (Ed.), *Visual languages and applications*. NY: Pergamon Press.

Dearholt, D. W. & Schvaneveldt, R. W. (1990). Properties of pathfinder networks. In R. W. Schvaneveldt (Ed.), *Pathfinder associative networks: Studies in knowledge organization*, (pp. 1-30).Norwood, NJ: Ablex.

Doyle, L. (1961). Semantic roadmaps for literature searchers. *Journal of the Association for Computing Machinery, 8(4),*367-391.

Draper, S. W. (1986). Display managers as the basis for user-machine communication. In Norman, D. A, & Draper, S. W. (Eds.). (1986). *User centered system design. 339-352.* Hillsdale, NJ: Erlbaum.

Eades, P. & Xuemin, L. (1989). How to draw a directed graph. *IEEE Workshop on Visual Languages,* 13-17.

Fairchild, K. M., Poltrock, S. E., & Furnas, G. W. (1988). SemNet: Three-dimensional graphic representations of large knowledge bases. In R. Guindon (Ed.), *Cognitive Science and its applications to human-computer interaction* (pp. 201-233). Hillsdale, NJ: Erlbaum.

Fidel, R. (1986). Towards expert systems for the selection of search keys. *Journal of the American Society for Information Science, 37(1)*, 37-44.

Frei, H. P., & Jauslin, J.-F. (1983). Graphical presentation of information and services: A user- oriented interface. *Information Technology: Research and Development, 2*, 23- 42.

Fowler, R. H. & Dearholt, D. W. (1990). Information retrieval using pathfinder networks. In R. W. Schvaneveldt (Ed.), *Pathfinder associative networks: Studies in knowledge organization*, (pp. 165-178).Norwood, NJ: Ablex.

Furnas, G. W. (1986). Generalized fisheye views. *Proceedings of Human Factors in Computing Systems (CHI ’86), 1*6-23. New York: ACM.

Gerber, L. M., Dearholt, D. W., Schvaneveldt, R. W., Sachania, V. & Esposito, C. (1987). *Documentation for PATHHFINDER: A program to generate PFNETs*. Computing Research Laboratory Memorandum in Computer and Cognitive Science, MCCS-87-47, New Mexico State University.

Godin, R, Gecsei, J., & Pichet, C. (1989). Design of a browsing interface for information retrieval. *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*, NY: ACM Press, 32-39.

Goldsmith, T. E. & Davenport, D. M. (1990). Assessing structural similarity of graphs. In R. W. Schvaneveldt (Ed.), *Pathfinder associative networks: Studies in knowledge organization*, (pp. 75-88). Norwood, NJ: Ablex.

Graf, M. (1990). Building a visual designer's environment. In S. Chang (Ed.), *Principles of visual programming systems*, (pp. 291-325). Englewood Cliffs, NJ: Prentice- Hall.

Hutchinson, J. W. (1989). NETSCAL: A network scaling algorithm for nonsymmetric proximity data. *Psychometrika, 54,* 25-51.

Ingwerson, P., & Wormell, I. (1986). Improved subject access, browsing and scanning mechanisms in modern online ir. *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*, 68-76.

Interlink. (1989). KNOT. Interlink, Inc., P. O. Box 4086 UPB, Las Cruces, NM 88003-4086.

Jones, W. P., & Furnas, G. W. (1987). Pictures of relevance: A geometric analysis of similarity measures. *Information Processing and Management, 38(6),* 420-442.

Kamada, T. & Kawai, S. (1989). An algorithm for drawing general undirected graphs. *Information processing letters, 31,* 7-15.

Knoke, D., & Kuklinski, J. (1982). *Network Analysis*. Beverly Hills, CA: Sage.

Laurel, B. (Ed.). (1990). *The art of human-computer interface design*. Reading, MA: Addison-Wesley.

Maron, M. E., & Kuhn, J. L. (1960). On relevance, probabilistic indexing, and information retrieval. *Journal of the Association for Computing Machinery, 7(3),* 216-244.

McDonald, J. E., Dearholt, D. W., Paap, K. R., & Schvaneveldt, R. W. (1986). A formal interface design methodology based on user knowledge. *Proceedings of Human Factors in Computing Systems (CHI '86),* 285-290. New York: ACM.

McDonald, J. E., Plate, T. A., & Schvaneveldt, R. W. (1990). Using pathfinder to extract semantic information from text. In Schvaneveldt, R. W. (Ed.), *Pathfinder associative networks: Studies in knowledge organization* (pp. 149-164). Norwood, NJ: Ablex.

McMath, C. F., Tamaru, R. S., & Rada, R. (1989). A graphical thesaurus-based information retrieval system. *International Journal of Man-Machine Studies, 31*, 121-147.

McCormick, B. H., DeFanti, T. A., & Brown, M. D. (1987). Visualization in Scientific Computing. (Eds.), *Computer Graphics, 21(6).*

Mettrey, J. (1987). An assessment of tools for building large KB systems. *AI Magazine, 8(4),* 81-89.

Miyamoto, S., Oi, K., Osamu, A., Katsuya, A., & Nakayama, K. (1986). Directed graph representations of association structures: A systematic approach. *IEEE Transactions on Systems, Man, and Cybernetics, 16,* 53-61.

Nielsen, J. (1990). *Hypertext and hypermedia.* San Diego, CA: Academic Press.

Nielson, G. M. (1989). Visualization in scientific computing, *IEEE Computer. 22(8),* 10- 11.

Norman, D. A, & Draper, S. W. (Eds.). (1986). *User centered system design.* Hillsdale, NJ: Erlbaum.

Palmquist, R. A., & Eisenburg, M. (1984). Testing a text analysis metric using magnitude estimation. *Proceedings of the American Society for Information Science Meeting,* 21, 231-236.

Roske-Hofstrand, R. J., & Paap, K. R. (1986). Cognitive networks as a guide to menu organization: An application in the automated cockpit. *Ergonomics, 29(11),* 1301-1312.

Rouse, W. B., & Rouse, S. H. (1984). Human information seeking and design of information systems. *Information Processing and Management, 20(1),* 129-138.

Salton, G. & McGill, M. J. (1983). *Modern information retrieval*. New York, McGraw- Hill.

Sammon, J. W. (1969). A nonlinear mapping for data structure analysis. *IEEE Transactions on Computers, 18(5),* 401-409.

Schvaneveldt, R. W. (Ed.), (1990). *Pathfinder associative networks: Studies in knowledge organization.* Norwood, NJ: Ablex.

Schvaneveldt, R. W., Dearholt, D. W., & Durso, F. T. (1989). Network structures in proximity data. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 24, pp. 249-284). New York: Academic Press.

Shneiderman, B. (1983). Direct manipulation: A step beyond programming languages. *IEEE Computer, 16(8),* 57-69.

Shu, N. C. (1988). *Visual programming*. NY: Van Nostrand Reinhold.

Sparck Jones, K. (1972). A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation, 28(1),* 11-20.

Sutherland, I. E. (1963). Sketchpad: A man-machine graphical communication system. *Proceedings of the Spring Joint Computer Conference,* 329-346.

Tamassia, R., Di Battista, G & Battini, C. (1988). Automatic graph drawing and readability of diagrams.

IEEE Transactions on Systems, Man, and Cybernetics, 18(1), 61-79.

Tversky, A., & Gati, I. (1982). Similarity, separability, and the triangle inequality. *Psychological Review*, *89*, 123-154.

Van Rijsbergen, C. J. (1979). *Information Retrieval (2nd Edition)*. London: Butterworth.

Valdez, F. & Chignell, M. (1988). *Browsing models for hypermedia databases.* Proceedings of the Human Factors 32nd Annual Meeting, 318-322.

Vickery, A. (1984). An intelligent interface for online interaction. *Journal of Information Science*, *9*, 7-18.

Watters, C. R., Sheperd, M. A., & Robertson, W. (1987). Towards an expert system for bibliographic retrieval. *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*, 272-281.