Visualizing Retrieved Information: A Survey

S. K. Card

UIR-R-1996-01

Visualizing Retrieved Information: A Survey

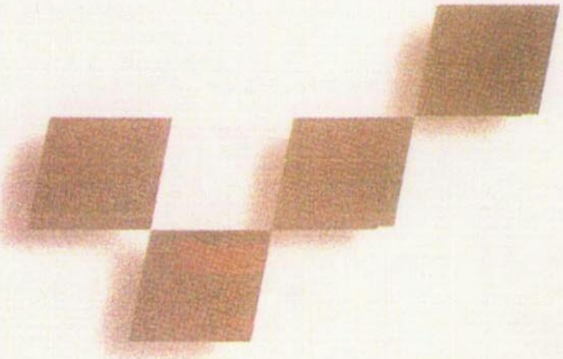

Stuart K. Card
Xerox PARC

Clearly, the presentation method for information retrieved from the global information infrastructure (GII) makes a big difference to users. The NCSA Mosaic interface, for example, with its point-and-click multimedia page presentation, swelled popular interest in the World Wide Web. This experience suggests the possibility of increased usefulness if we apply visualization techniques to information retrieved from the GII.

Note that, although the term "information visualization" is coming into use, the goal is really "information perceptualization." The latter implies a richer use of many senses, including sound and touch, to increase the rate at which people can assimilate and understand information.

In discussing the visualization of retrieved information, it helps to consider four functional levels: (1) the infosphere, (2) the workspace, (3) sensemaking tools, and (4) the document. This simple classification lets us separate the functions served by the visualizations from the techniques themselves, which can be applied across functional levels.

Infosphere

The infosphere (variously called the docuverse or docuspace, or sometimes just cyberspace) is the reachable space of information sources, largely databases and documents. Of course, a visualization of the infosphere does not presuppose that all elements in the space are visible at one time (especially given a space the size of the GII). Logically, the infosphere includes *repositories*, places where collections of information reside. Whether this structure is represented in the visualization is a current research issue.

How might we visualize the infosphere? The galaxy visualization method of Battelle Pacific Northwest National Laboratory[1] uses multidimensional scaling of the similarities between documents (based on the correlation between their vector representations) to plot one point in Euclidean space for each document. These points form a semantic scatterplot. Similar documents tend to lie near each other, although the axes of the space aren't always easy to interpret. The galaxy can also be projected onto a plane and elevation used to code frequency of mention, forming a semantic landscape of doc-

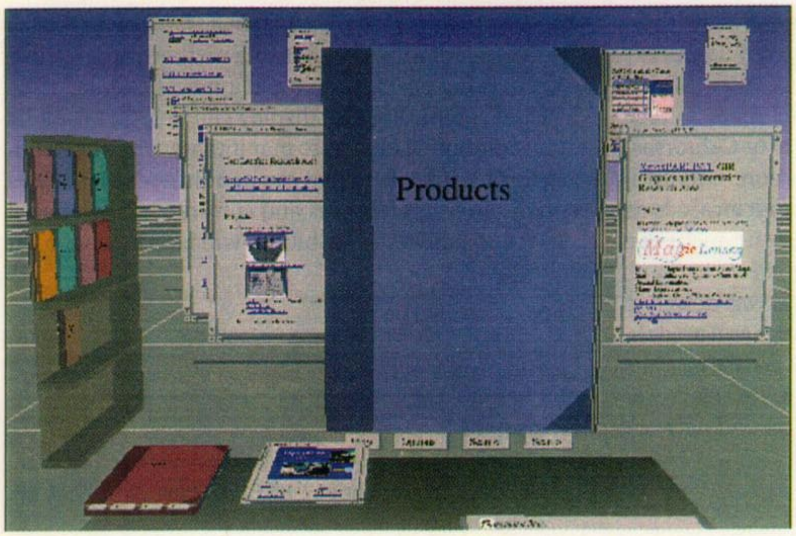

Xerox PARC

1 The Web Forager, which provides a dynamic 3D workspace.

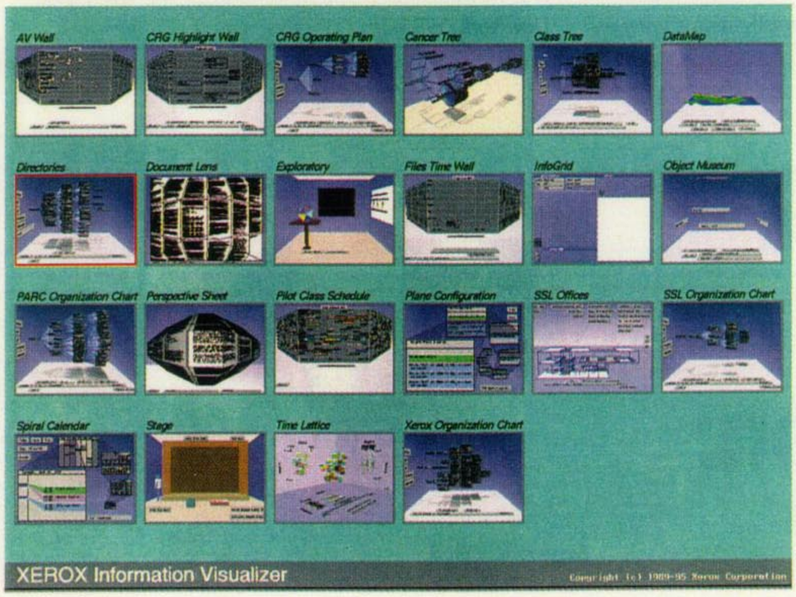

Xerox PARC

2 With 3D rooms, users can either occupy a space or get an overview of all the space.

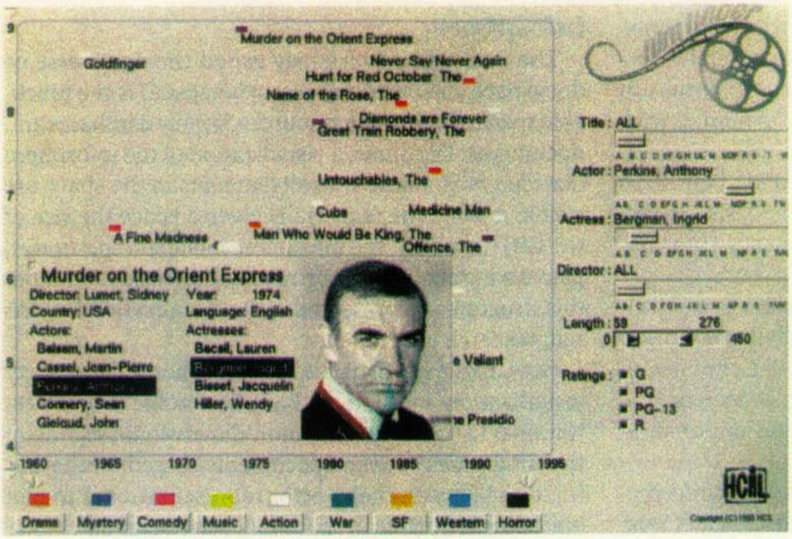

Murder on the Orient Express
Director: Lumet, Sidney
Year: 1974
Country: USA
Language: English
Actors:
Balsam, Martin
Cassel, Jean-Pierre
Perkins, Anthony
Connery, Sean
Gielgud, John
Actresses:
Bacall, Lauren
Bergman, Ingrid
Bisset, Jacquelin
Hiller, Wendy

Title: ALL
Actor: Perkins, Anthony
Actress: Bergman, Ingrid
Director: ALL
Length: 59 276
0 450
Ratings: ■ G ■ PG ■ PG-13 ■ R

By permission of B. Shneiderman

3 The Film Finder.

uments (see Figure 2 in Gershon's section of this report).

Ironically, a system with a similar name uses an almost opposite approach. The Galaxy of News system[2] creates a semantic space of news stories whose visual structure is not fixed, but changes as the user moves. The user sees an arrangement of words derived from the news stories. Zooming in on a particular word causes other words, representing finer classifications of the story, to appear. The user zooms towards the relevant word, and yet more words appear. Words previously rejected may appear in the context of words now selected, taking on new, nuanced meanings. Although the user experiences a space, there is no stable space—what is stable is the path.

Feiner's "worlds within worlds" method[3] for multidimensional databases falls somewhere between these two. In his technique, a six-dimensional object, say, is plotted by having a new coordinate system for the last three dimensions sit inside the coordinate system for the first three. Benedikt[4] proposed representing the entire infosphere in a similar manner.

Other methods tackle how to represent parts of the infosphere. In contrast to the representation of documents as points or clouds in a multidimensional space, Kohonen maps divide the available space into semantic regions.[5] Or, instead of using spatial coordinates, documents can be visualized as linked graphs or trees in two or three dimensions, as in the Narcissus system.[6] Even more radically, the nodes and links can be projected onto a non-Euclidean landscape. Lamping and Rao[7] recast a linked network as a tree, then plot the tree in hyperbolic space. As the user traverses the landscape, parts of it appear to enlarge and shrink according to distance from the user.

Workspace

A *workspace* is a collection of objects whose access is arranged to make relevant tasks efficient. Workspaces allow multiple objects to be compared and grouped or combined. The design of a workspace induces a cost structure over it, that is, a set of times associated with accessing and manipulating its objects. Visualization methods can improve this cost structure by enabling more information in a display, by increasing the information-carrying abilities of periph-

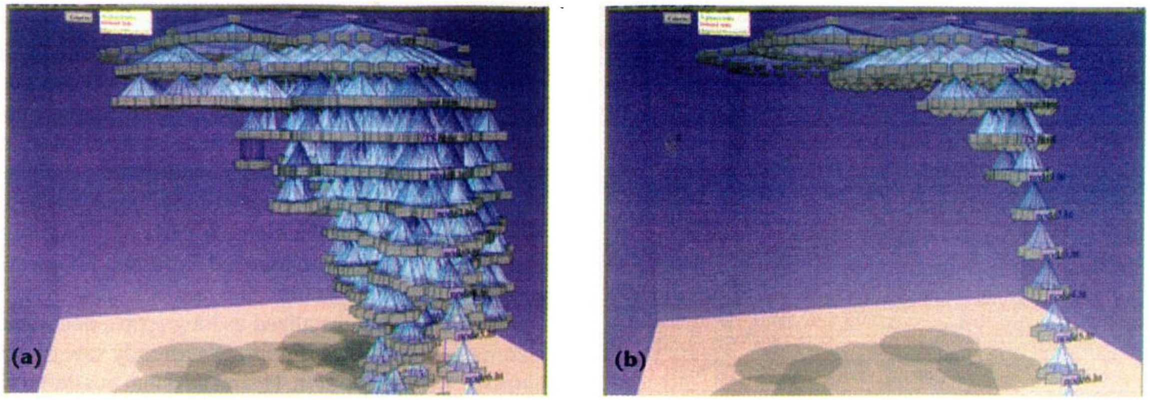

Xerox PARC

4 (a) Cone tree with 20,000 nodes representing part of the WWW and (b) the fish-eye version of the same cone tree.

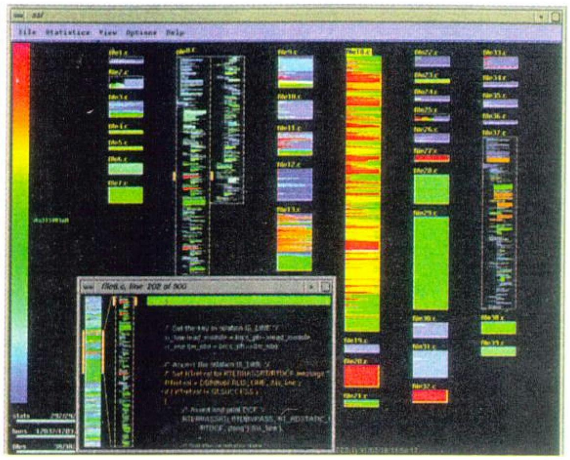

By permission of S Eick

5 SeeSoft display of program code from AT&T Bell Labs.

eral parts of the display, and by shifting load from slower cognitive to faster perceptual processes.

Figure 1 shows one version of a dynamic 3D workspace, the Web Forager.[8] The 3D perspective allows more objects in the space, and a gesture interface permits rapid interaction.

Figure 2 shows another technique for building information-intensive workspaces. The workspace divides into "rooms" corresponding to user information working sets,[9] such as projects or clients. Information can be in more than one room at a time, and the user can either occupy a space or get an overview of all the space.

A workspace can also be based on zooming, as in Pad++.[10] Large amounts of information can be placed in the workspace by making them very small. Visualization techniques help present objects in a way that provides information across large zoom scales.

Yet another technique provides a different scale for the center and the periphery of the workspace. In Spence and Apperly's bifocal lens,[11] the sides of the workspace are horizontally compressed to hold more information.

Finally, considerable research involves basing workspaces on more literal geometric spaces through VRML (the Virtual Reality Modeling Language).

Sensemaking tools

Sensemaking tools help users understand information by associating and combining it. They re-represent retrieved information to make patterns visible, or they allow the construction of new information patterns from old. One type of visualization is a *clever static view*, such as Spoerri's "information crystal."[12] This method, an elaborate form of Venn diagram for information retrieval based on Boolean queries, allows the user to see the patterns of hits for many terms. Another example, the VIBE system,[13] uses word frequencies to position a moderate number of documents (50 to 120) relative to a small number of spatially-positioned, term-defined points of interest.

A second type of visualization, a *dynamic view*, changes in response to user action, either by culling, by movement, by distortion, or by altering visual properties. For example, the Film Finder[14] employs culling (see Figure 3). The user moves sliders to indicate values of interest for actor, director, and length of a movie. Movies not meeting these specifications are visually culled as the slider moves. Those movies that do meet the criteria are plotted in a scattergraph of rating by year The user can click on dots representing movies to get details.

Examples of dynamic movement and distortion include the *cone tree* and the *perspective wall*. The perspective wall is a variant of the bifocal display, but it uses perspective instead of direct scaling to reduce the size of elements outside the focal area. The cone tree is a 3D tree that articulates to bring forward into focus those position nodes the user touches. The tree can be culled manually or automatically to show only items of interest. In the fish-eye version of the cone tree, an algorithm based on Furnas's work[15] expands only those parts of the tree related to some target node.

The cone tree can be applied to a portion of the World Wide Web infosphere by recasting the links as a tree according to a traversal algorithm Figure 4a shows a cone tree of 20,000 nodes representing the internal Xerox WWW. Figure 4b shows the same tree after applying the fish-eye techniques. When the users touches some node in the tree, the parts of the tree rotate, elements around the touched element grow, and less relevant parts of the tree shrink.

Sensemaking often benefits from visualization techniques that let the user both span large-scale factors and use coloring, position, or other visual attributes to present classifying information. In large scale, these visual attributes form patterns quickly detected by the eye. The SeeSoft system[16] (see Figure 5) allows a user to view more than 12,000 lines of code. Colors represent how

6 WebBook being fanned so that the user can rapidly examine a group of related WebPages.

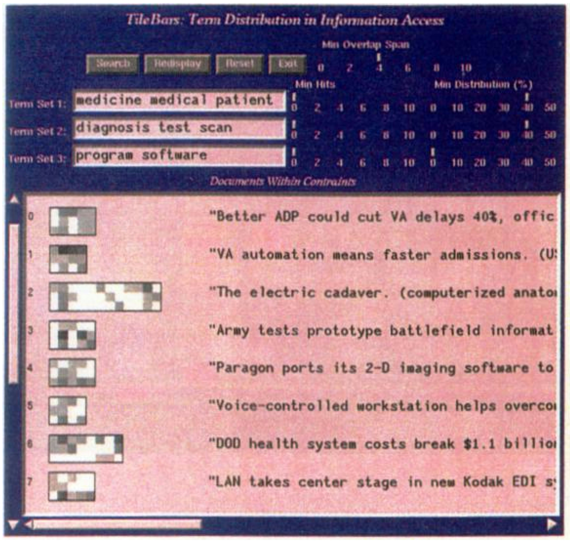

Term Set 1: medicine medical patient

Term Set 2: diagnosis test scan

Term Set 3: program software

0 "Better ADP could cut VA delays 40%, offic

1 "VA automation means faster admissions. (U:

2 "The electric cadaver. (computerized anato

3 "Army tests prototype battlefield informat

4 "Paragon ports its 2-D imaging software to

5 "Voice-controlled workstation helps overco

6 "DOD health system costs break $1.1 billio

7 "LAN takes center stage in new Kodak EDI s

7 TileBars.

By permission of M Hearst

recently the code was modified, but could also represent programmer or department. Brushing the recency scale highlights the appropriate lines in the code text display. The figure also shows a separate inset window of code (a focus display) associated with the selected segment of the overview display.

Rather than having separate focus and context displays, detailed focal information can be integrated with a scaled display, as in the Table Lens system.[17] The Table Lens permits the display of a matrix, such as a spreadsheet, up to 100 times larger than would ordinarily fit on the display. The matrix can be sorted by column and manipulated with other techniques to reveal patterns in the data.

Visual representation of the matrix of correlations between word co-occurrences or other structural parameters of text form the basis for a number of systems. Gershon et al.,[18] for example, count the co-occurrences of words in a set of documents, then rearrange the matrix to give clusters.

Document

Finally, consider visualization of individual documents, which, for our purposes here, are the elementary units of retrieval. Documents may be large and themselves have structure.

The WWW is currently organized as pages. An obvious step is to create entities at a higher level than the page. The WebBook[8] (see Figures 1 and 6) arranges pages into a simulated physical book. This lets users do the most elementary operation of sensemaking: grouping. Instead of waiting for each page on a hotlist to be accessed, users can access pages immediately, even fan them for rapid scanning (as in Figure 6). One of the best-known document visualizations, Superbook,[19] lets users turn existing text into an indexed hypertext book with a fish-eye table of contents.

Text can also be visualized as long strips, decorated in some way to show content. We already saw this in Figure 5. Another version, Hearst's system of TileBars[20] (see Figure 7), divides the text into equal blocks. The user inserts query terms, and each block on a strip darkens according to the frequency of those terms in that block, each term being associated with a different line on the strip.

Visualization and user interfaces

The research summarized here suggests how to improve the human interface to the Web. Whereas browsers currently show only a single local node, they could be given a visualization of the infosphere with landmarks. Instead of a single page, they could be given a visualization of a workspace of multiple pages to compare and combine. Furthermore, they could be given visualization sensemaking tools for analyzing and investigating the information. Finally, instead of a single page, they could be given visualizations of compound documents and other abstractions. Information access is a "killer app" for the 1990s. Information visualization will play an important role in its success. ■

References

1. J.A. Wise et al., "Visualizing the Nonvisual: Spatial Analysis and Interaction with Information from Text Documents," *Proc Info Vis. Symp. 95*, N. Gershon and S.G. Eick, eds., IEEE Computer Society Press, Los Alamitos, Calif., 1995, pp. 51-58.
2. E. Rennison, "Galaxy of News An Approach to Visualizing and Understanding Expansive News Landscapes," in *UIST 94*, ACM Symp. on User Interface Software and Technology, ACM Press, New York, 1994, pp. 3-12
3. S. Feiner and C. Beshers, "Worlds Within Worlds Metaphors for Exploring *n*-dimensional Virtual Worlds." in *UIST 90. ACM Symp on User Interface Software*, ACM Press, New York, 1990. pp. 76-83.
4. M Benedikt, "Cyberspace Some Proposals," in *Cyberspace: First Steps*, M. Benedikt, ed., MIT Press, Cambridge, Mass., 1991, pp. 119-224.
5 X Lin, P. Liebscher, and G. Marchionini, "Graphic Representation of Electronic Search Patterns," *J. Am Society for Info Science*, Vol. 42, No 7, 1991, pp. 469-478.
6. R.J. Hendley et al , "Narcissus Visualizing Information," in *Proc. Info Vis Symp 95*, N Gershon and S.G Eick, eds., IEEE CS Press, Los Alamitos, Calif , 1995, pp 90-96
7. J Lamping, R. Rao, and P. Pirolli, "A Focus + Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies," in *CHI 95*, Proc ACM Conf on Human Fac-

tors in Computing Systems, ACM Press, New York, 1995, pp. 401-408.

8. S.K. Card, G.G. Robertson, and W. York, "The WebBook and the Web Forager: An Information Workspace for the World Wide Web," to appear in *CHI 96*, ACM Conf. on Human Factors in Software, ACM Press, New York, 1996.
9. J.D.A. Henderson and S.K. Card, "Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention in Window-based Graphical User Interfaces," *ACM Trans. on Graphics*, Vol. 5, No. 3, July 1986, pp. 211-241.
10. B. Bederson, "Pad++: Advances in Multiscale Interfaces," in *CHI 94*, Conf. Companion, ACM Press, New York, 1994, p. 315.
11. R. Spence and M. Apperley, "Data Base Navigation: An Office Environment for the Professional," *Behavior and Info. Technology*, Vol. 1, No. 1, 1982, pp. 43-54.
12. A. Spoerri, "Visual Tools for Information Retrieval," in *Proc. IEEE Symp. on Visual Languages*, IEEE CS Press, Los Alamitos, Calif., 1993, pp. 160-168.
13. K.A. Olsen et al., "Visualization of a Document Collection: The VIBE System," *Info. Processing and Management*, Vol. 29, No. 1, 1993, pp. 69-81.
14. C. Ahlberg and B. Shneiderman, "Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays," in *CHI 94*, Proc. ACM Conf. on Human Factors in Software, ACM Press, New York, 1994, pp. 313-317.
15. G. Furnas, "Generalized Fisheye Views," in *CHI 86*, Proc. ACM Conf. on Human Factors in Software, ACM Press, New York, 1986, pp. 16-23.
16. S.G. Eick, J.L. Steffen, and E.E. Sumner, "SeeSoft—A Tool for Visualizing Software," *IEEE Trans. on Software Engineering*, Vol. 18, No. 11, Nov. 1992, pp. 957-968.
17. R. Rao and S.K. Card, "The Table Lens: Merging Graphical and Symbolic Representation in an Interactive Focus + Context Visualization for Tabular Information," in *CHI 94*, Proc. ACM Conf. on Human Factors in Software, ACM Press, New York, 1994, pp.318-322.
18. N. Gershon et al., "Visualizing Internet Resources," in *Proc. Info. Vis. Symp. 95*, N. Gershon and S.G. Eick, eds., IEEE CS Press, Los Alamitos, Calif., 1995, pp. 122-128.
19. J.R. Remde, L.M. Gomez, and T.K. Landauer, "Superbook: An Automatic Tool for Information Exploration," in *ACM Hypertext 87 Proc.*, ACM Press, New York, 1987, pp. 175-188.
20. M.A. Hearst, "TileBars: Visualization of Term Distribution Information in Full-Text Information Access," in *CHI 95*, Proc. ACM Conf. on Human Factors in Software, ACM Press, New York, 1995, pp. 55-66.