Visualizations for Digital Libraries

Gang Wan

The concept of digital libraries is familiar to both librarians and library patrons today. These new libraries have broken the limits of space and distance by delivering information in various formats via the Internet. Since most digital libraries contain a colossal amount of information, it is critical to design more user-friendly interfaces to explore, understand, and manage their content. One important technique for designing such interfaces is information visualization. Although computer-aided information visualization is a relatively new research area, numerous visualization applications already exist in various fields today. Furthermore, many library professionals are also starting to realize that combining information visualization techniques and current library technologies, such as digital libraries, can help library users find information more effectively and efficiently. This article first discusses three major tasks that most visualization for digital libraries emphasize, and then introduces several current applications of information visualization for digital libraries.

A good understanding of user tasks is the foundation of designing useful visualizations. Rao et al. defined several specific user tasks of digital libraries and illustrated some existing information-visualization techniques that can be used to enhance these tasks, such as TileBar, Cone Tree, and Document Lens.[1] The tasks were browsing subsets of sources iteratively, viewing context-of-query match, visualizing passages within documents, rendering sources and results, reflecting time costs of interaction, managing multiple-search processes, integrating multiple search and browsing techniques, and visualizing large information sets. Moreover, Zaphiris et al. generalized these tasks into three essential ones: searching, navigation, and browsing.[2] Indeed, most information-visualization projects for digital libraries have emphasized these three tasks.

In terms of searching, Shneiderman et al. proposed the use of a two-dimensional display with continuous variables to view several thousand search results simultaneously.[3] This visualization included two strategies: two-dimensional visualizations, and browsers for hierarchical data sets (implemented by using categorical and hierarchical axes). In combination with a grid display, this visualization let users see an overview by color-coded dots or bar charts arranged on a grid and organized by familiar labeled categories. Users could probe further by zooming in on desired categories or switching to another hierarchical variable. A language-independent document-classification system, completed by Liu et al., provided a search aid in a digital-library environment and helped users analyze the search query results visually.[4] This system used a vector model to calculate the similarities between documents and a Java applet to display the classification to the user.

In terms of navigation, there are also a variety of information-visualization applications. The previous example of two-dimensional display developed by Shniederman et al. also contained navigation functions.[5] Another example is Hascoet's map interface applied to a digital library.[6] This prototype was associated with summary views in the form of navigation trees and neighbor trees that showed documents related to one focus document. The user interface was composed of maps automatically generated based on the characteristics of documents retrieved and a default configuration. Users could also modify the configuration of maps and edit maps (classical operations such as cut, paste, move, delete, save and load a view, and expand a view).

As for browsing, the use of dynamic queries is a technique that has been employed for some time. Ahlberg and Shneiderman's (1994) FilmFinder is an early example.[7] Users can move several sliders to select query parameters, and the search results change with the movement of the sliders. This tool can help users browse movie records more easily and cognitively. Another technique is Query Previews, proposed by Doan et al.[8] Query Previews allows users to rapidly gain an understanding of the content and scope of a digital collection. Users are presented with generalized previews of the entire database using only the most salient attributes. When they select rough ranges, they will immediately learn the availability of the data for their proposed query.

All these applications provide good examples and paradigms to some recent projects. This paper's discussion of visualization techniques will be based on these three essential tasks—searching, navigation, and browsing.

Techniques and applications

This section presents several recent information-visualization projects applied to digital libraries. All these applications emphasize searching, navigation, and browsing.

Gang Wan (wangang11@gmail.com) is Science Librarian, Texas A&M University, College Station.

LVis—Digital Library Visualizer

Indiana University's (IU's) LVis (Digital Library Visualizer) project aims to aid users' navigation and comprehension of large document sets retrieved from digital libraries. Borner et al. developed a prototype of LVis based on the data set in the Dido Image Bank, provided by IU's department of art history.[9] LVis is a good combination of information-retrieval algorithms and visualized-search interface. In the information retrieval and classification stage, it adopts Latent Semantic Analysis (LSA) to automatically extract semantic relationships between images. The LSA output feeds into a clustering algorithm that groups images into several classes sharing semantically similar descriptors. A modified Boltzman algorithm is then used to lay out images in space. This section will focus on the interface metaphors used to display the results of this classification.

Two interfaces have been implemented for LVis. A 2D Java applet was used on a desktop computer for details, and a 3D immersible environment for the CAVE (CAVE Automatic Virtual Environment). CAVE is a virtual reality 10' x 10' x 10' theater made up of three rear-projection screens for walls and a down-projection screen for the floor. Projectors throw full-color workstation fields (1,280 x 512 stereo) at 120Hz onto the screens, giving between 2,000 and 4,000 linear pixel resolution to the surrounding composite image.[10]

Both 2D and 3D interfaces give users access to three levels of detail: they provide an overview about document clusters and their relations; they show how images belonging in the same cluster relate to one another; and they give more detailed information about an image, such as its description or its full-size version.

In the CAVE environment, users can first enter a virtual display theater that stages the digital library as a cyberspace Easter Island, presenting gateways to specific subject categories established by the previous classification process. Borner et al. used 3D icons here to encode subject categories (in this case, they actually used a sculptural form of heads inspired by images of the data set).[11] After users "enter" into these head icons, they are transited to a new 3D spatial metaphor that presents images in the current category. These images, or slides from the digital library, are presented in crystalline structures (figure 1).

In this environment, each crystal represents a set of images with semantically similar image descriptions. Again, physical proximity is used here as a metaphor to encode semantic similarities among images. The formations of the crystalline structures depend on the size of the actual search-result data set. Navigation in this space is easy. Users can also "walk" through this environment and select images of interest to display a larger- and clearer-size version (as two images shown in figure 1). If the larger version is not satisfactory, it can be returned to its previous iconic presentation.

UC system: A Fluid Treemap Interface for digital libraries

The UC system—the acronym "UC" came from its original (but no longer used) internal name "UpLib Client"—was developed by Good et al. at Palo Alto Research Center.[12] It was built on the UpLib personal digital-library platform, which provides an "extensible, format-agnostic, and secure document-repository layer."[13] "Personal" here means that the user already has the right to use all of the data objects in the library, and already has local possession of those objects. However, this visualization can be employed in more general digital libraries.

The UC system uses continuous and quantum Treemap layouts to present collections of documents. Continuous Treemaps are space-filling visualizations of trees that assign an area to tree nodes based on the weighting of the nodes.[14] In continuous Treemaps, the aspect ratio of the cells is not constrained, although square cells are often preferred. Quantum Treemaps extend this idea by making cell dimensions an even multiple of a unit size.[15] The Treemap visualizations provide meaningful overviews of document collections and fast, intuitive navigation among multiple documents in a working set.

An important aspect of the interface is the fluidity of navigation. This allows the user to focus on the documents rather than on interacting with the tool. The interface allows a user to zoom in on an object with a left-click, and to zoom out when the user clicks on the background;
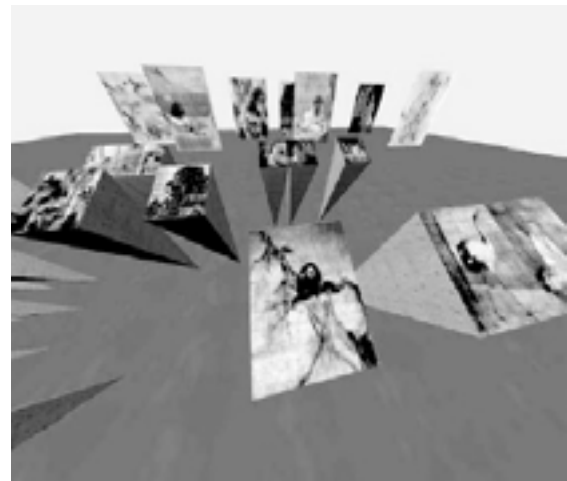

Figure 1. LVis: image crystals and panels

however, the combination of a "zoomable" user interface and continuous Treemaps leads to a problem: conflicts with aspect ratios. To solve this problem, Good et al. proposed to zoom and morph the cell to the window size while leaving the rest of the layout in place.[16] Thus the visual disturbance of the display is minimized since only a single cell moves.

With respect to searching, this system provides several methods to filter results. First, its interface includes a mechanism to search for specific content within the documents. As letters are added to the search query, the system increasingly highlights matching documents to immediately indicate matching documents. Secondly, the user can also choose to update the view to display only those documents that match the current query. Figure 2 gives an example of a user-initiated query process.

In the scenario described in figure 2, the user first enters the search terms (2.1), and interactive highlights then appear for groups with matching documents (2.2). The user presses a button to limit the view to only the matching documents (2.3). Finally, the user zooms in on a document and begins reading with an integrated reading tool. (2.4)

The UC system also offers a mechanism that allows the user to compare multiple documents. After users retrieve a set of documents through a search, they can press a button to "explode" the documents to pages. They can continue zooming in to a portion of a single document, and then select a document page to read with the integrated reading tool.

In short, the UC system uses Treemaps as the primary visual metaphor. It also uses various visualization techniques that enhance user interactions, such as zooming, interactively highlighting, exploding, etc.

ActiveGraph

ActiveGraph is an information-visualization tool designed by Marks et al. (2000) at Los Alamos National Laboratory (LANL).[17] It aims to provide users with a concise, customizable view of documents in a digital library. In this system, a set of digital-library documents is represented as a data set in a 2D or 3D scatter plot. The data set can represent any digital-library objects in various formats including books, journals, papers, images, and Web resources.

Marks et al. used six visual attributes of the scatter plot: the X-, Y-, and Z-axes, color, size, and shape to encode the bibliographic information of documents in a digital library, including title, author, date of publication, and number of citations.[18] The user can select and adjust these attributes from a control panel on the right-hand side of the screen. Thus, ActiveGraph allows users to both view and customize the contents of a digital library.

The main visual representation of this tool is a scatter plot. Scatter plots have been used to represent large sets of data for a long time. They provide an overview of a data set and show the distribution of data points clearly, revealing clusters and statistical information.[19] Hence, these scatter plots make it possible for users to perceive meaningful patterns of the data.

An example of using ActiveGraph scatter plots to visualize citation data for postdoctoral researchers at LANL is given in figure 3. This scatter plot intends to provide users with information, such as the number of times their papers published between 1998 and 2002 were cited. The visualization is based on the metadata in the LANL digital library.

In this scatter plot, the postdoctoral researcher's last name is mapped to the X-axis and the number of citations is mapped to the Y-axis. A pixel of a particular color can provide two pieces of information, for example, by encoding a paper and the subject category of the paper. A group of pixels of a particular size, shape, and color can provide four pieces of information by encoding a paper, the subject category of the paper, whether the paper has been cited, and whether the paper has been read by another user of the collaborative library. From this scatter plot, users can easily learn the citation pattern of these papers.

Unlike some other scatter plot applications such as HomeFinder and FilmFinder, ActiveGraph uses different filters for queries. Instead of filter sliders, it uses filter lists, which consist of selection list boxes, one for each data attribute. These filter lists can provide users with functionality that is important in the context of digital libraries. ActiveGraph allows users to manipulate the display of data in another manner by applying a logarithmic transformation. As some data sets, such as citation data, can frequently have an exponential distribution, the

Figure 2. The search interface of the UC system

logarithmic transformation can spread the clustered dots more evenly across the scatter plot.

Other data transformations and visualizations may be important in some cases as well, such as parallel coordinates for displaying citation statistics for the same group of researchers at different points in time.

The scatter plot is not a new visualization technique. This example, however, demonstrates that by encoding document attributes and designing proper filters, it can be used in a digital library environment effectively and efficiently.

3D Vase Museum

The previous example, LVis, already introduced the applications of 3D representations in digital libraries. The 3D Vase Museum developed by Shiaw et al. at Tufts University is another good example of 3D space metaphor in digital libraries using a variety of visualization techniques.[20]

In this 3D Vase Museum, the user can navigate seamlessly from a high level scatter-plot-like plan view to a perspective overview of a subset of the collection, to a view of an individual item, to retrieval of data associated with that item, all within the same virtual room and without any mode change or special command.

Unlike the traditional digital library, which displays thumbnails and descriptions of vases in the main browser interface, this museum is a 3D virtual environment that presents each vase as a 3D object and places it in a 3D space of a room within a museum. Figure 4 gives a wide-angle view of this 3D museum.

In this view, one wall represents the timeline (year BCE) and the adjacent wall represents the types of wares (e.g., red figures, black figures). The user can "walk" through this virtual room and look at the vases. The wide-angle view pictured in figure 4 will then be transited to an eye-level view so that the user can probe the objects more clearly. When the user continues "walking" toward an object of interest, secondary information about this vase will appear in the virtual scene. If the user looks closer, the text information becomes clearer. As the user moves farther and farther away, the information becomes less and less visible until it eventually disappears from the scene.

If the user clicks on the vase HTML page, a version of the original HTML page will be loaded, from which a 3D model of the vase can be loaded. The user can then rotate this 3D model on the screen using the mouse to see all the aspects of the vase. The 3D Vase Museum is maintained in the background all times.

The user can also navigate the room in a perspective view by switching the view port upward toward the ceiling (watching from upside down). The user can then switch the views between a high-level scatter plot and a 3D perspective view. Similarly, in this application, the X- and Y- axes are mapped to two attributes of the vases: year and ware. With this seamless blend from a high-level data plot to 3D objects, the user can navigate without losing the point of view or context by just moving within the virtual environment.

According to Shiaw et al., a usability test has been administered, in which tasks based on archaeology courses were designed and subjects were asked to perform these tasks in the original traditional digital library and this 3D museum.[21] The results showed subjects who used the 3D Vase Museum performed the tasks 33 percent better and did so nearly three times faster.
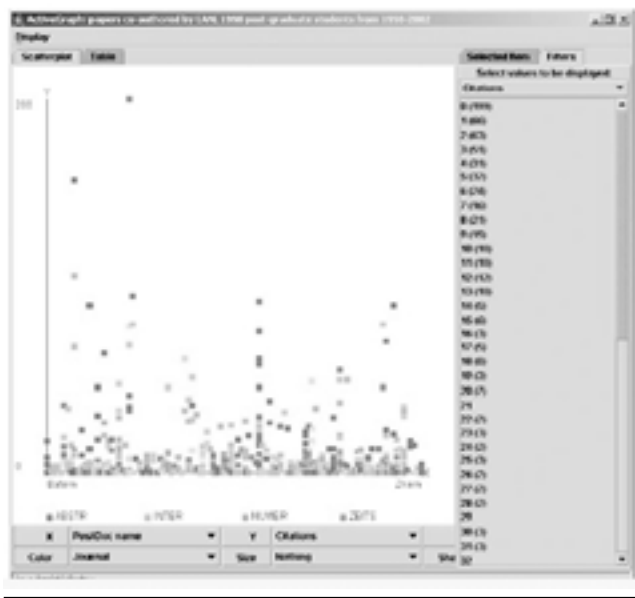
Collaborative Visual Interfaces to digital libraries

Collaborative Visual Interfaces is an ongoing project led by Borner at Indiana University (IU). Borner et al. (2002) proposed the development of a shared 3D document space for a scholarly community—namely faculty, staff, and students at IU's School of Library and Information Science.[22] The space will provide access to a collection of various online documents including text, images, video, and software demonstrations.

A Semantic Treemap algorithm has been developed to layout documents in a 3D space.[23] Semantic Treemaps utilize the original Treemap approach to determine the size (dependent on the number of documents) and layout

Figure 3. ActiveGraph scatter plot of citation data for papers authored by LANL postdoctoral researchers

of document clusters. Subsequently, an algorithm (force directed placement) was applied to the documents in each cluster to place documents spatially, based on their semantic similarity, which was encoded by the physical proximity between two dots. An example of the Semantic Treemaps is shown in figure 5.1.

A 3D space metaphor was then used to display these documents on the desktop interface, as shown in figure 5.2. In this 3D space, each document is represented by a square panel textured by the corresponding Web page's thumbnail image and augmented by a short description such as the Web page title that appears when the user moves the mouse over the panel. As in other 3D environments, users can "walk" through this space to probe documents of interest. Upon clicking the panel, the corresponding Web page is displayed in the Web-browser interface. Users can collaboratively examine, discuss, and modify (add and annotate) documents, thereby converting this document space into an ever-evolving repository of the user community's collective knowledge that members can access, learn from, contribute to, and build upon.

Certain usability studies have been performed to determine the influence of panel size and panel density on retrieval performance. Results showed that subjects were slightly faster and more accurate if Web-page panels are larger and denser.

▉ AquaBrowser

AquaBrowser is a fuzzy visualization tool that shows the high-level description of a conceptual space, hiding irrelevant information and displaying information elements in context.[24] It is a generic Java applet that can be embedded into any Web page. Medialab, the developer of AquaBrowser, claims that users of AquaBrowser can browse through a dynamic conceptual space that is continually reshaped to reflect their interests. Animations make transitions from one state to another appear more fluid, showing users why and how the information is rearranged.

Medialab uses the term "word cloud" as the visual metaphor of the AquaBrowser interface. But in fact, the primary visual representation is a network of linked words

that are distributed in the conceptual space. The search term that the user assigns will display at the center of this network. The physical distance between another term node and this term encodes the relevancy between these terms. The larger and closer the word is to the center of the screen, the greater its relevance to the search term. In contrast, the smaller and more peripherally positioned, the less relevant it is. Each of the user's actions will change and rearrange the distribution and importance of the words, putting those of greater interest closer to the user and those of less interest nearer to the edge of the screen. It also uses colors to encode attributes of terms, such as spelling variations, visited words, and translations. Figure 6 shows an example of a search display.

This tool has been used by a number of libraries to enhance their online catalog search interfaces. It could be a very useful search aid in digital libraries as well.

▉ Summary and trends

The above applications are just a few examples of information visualization in a digital-library environment. Many other metaphors and techniques, such as perspective wall, cone tree, document lens, and hyperbolic browser, have been used or can potentially be used to facilitate searching, browsing, and navigating through the maze of information in a digital library.

The digital library is an interdisciplinary subject involving several research areas such as information retrieval, multimedia information processing, and classification. All these aspects of digital libraries make information visualizations more complicated in this environment. Therefore, the systems described in this paper have integrated various visualization techniques.
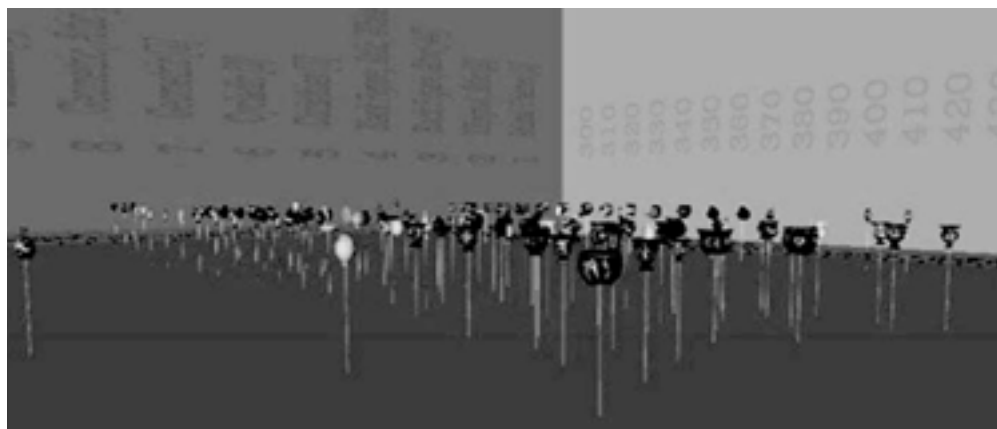

Figure 4. A wide-angle overview of the 3D Vase Museum

The examples in this paper, along with many others, show that the 3D space metaphor has attracted much attention from information-science communities. The combination of 3D space and virtual reality that can be accessed from Web browsers these days is becoming a trend of information visualization for digital libraries. This technique gives the user maximum freedom to walk through the library collections, searching and browsing documents. The 3D visual structures, however, have greater implementations compared with those that are 2D, since they require more processing power and include more parameters.[25] That is partly why many 2D visualizations developed in the 1990s are still widely used. For example, both ActiveGraph and 3D Vase Museum have employed 2D scatter plots; both UC system and Collaborative Visual Interfaces have used Treemaps.
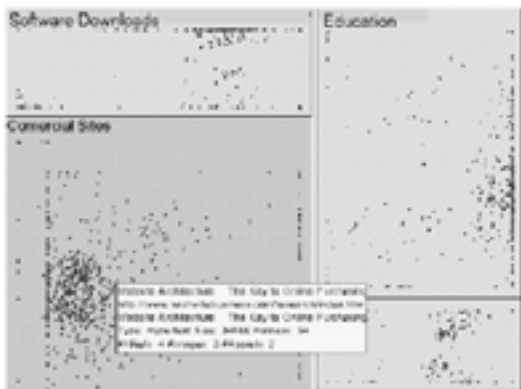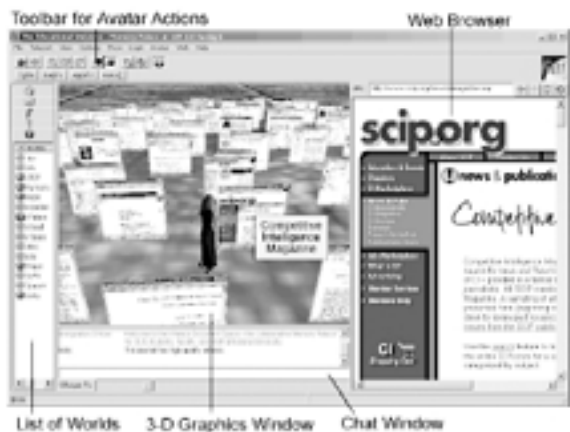
Furthermore, it is very important to focus on the actual needs of users. Research on any visualization for digital libraries should be based on the detailed analysis of users, their information needs, and their tasks.[26] Usability tests have been done for some of the above applications, but not for all of them. Further research and usability tests are required to determine to what extent a visual interface facilitates the user's perception of information.

References and notes

1. R. Rao et al., "Rich Interaction in the Digital Library," *Communications of ACM* 38, no. 4 (1996): 29–39.

2. P. Zaphiris et al, "Exploring the User of Information Visualization for Digital Libraries," *The New Review of Information Networking* 10, no. 1 (2004): 51–69.

3. B. Shneiderman et al., "Digital Library Search Results with Categorical and Hierarchical Axes," *DL-00: 5th ACM Digital Library Conference, San Antonio* (New York: ACM Pr., 1999).

4. Y. Liu et al, "Visualizing Document Classification: A Search Aid for the Digital Library," *Journal of the American Society for Information Science* 51, no. 3 (2000), 216–27.

5. Shneiderman et al., "Digital Library Search Results with Categorical and Hierarchical Axes."

6. M. Hascoet, "Using Maps As a User Interface to a Digital Library," *SIGIR '98, Melbourne, Australia* (New York: ACM Pr., 1998).

7. C. Ahlberg and B. Shneiderman, "Visual Information Seeking Using the FilmFinder," *ACM CHI94 Conference, Boston* (New York: ACM Pr., 1994).

8. K. Doan et al., "Query Previews for Networked Information Services," *Advanced Digital Libraries Conference* (Washington: IEEE, 1996).

9. K. Borner et al., "LVis—Digital Library Visualizer," Proceedings, IEEE International Conference on Information Visual-

Figure 5.1. A semantic Treemap of Web links

Figure 5.2. Interface to the document space

Figure 6. A search display in AquaBrowser

ization, July 19–21, 2000, London, England (Los Alamitos, Calif.: IEEE Computer Society, 2000), 77–81.

10. C. Cruz-Neira et al., "Surround-screen Projection-based Virtual Reality: The Design and Implementation of the CAVE," *Computer Graphics (Proceedings of SIGGRAPH '93)*, vol. 27 (New York: ACM SIGGRAPH, 1993), 135–42.

11. K. Borner et al., "LVis—Digital Library Visualizer."

12. L. E. Good et al., "A Fluid Treemap Interface for Personal Digital Libraries," *JCDL'05, June 7–11, Denver* (New York: ACM Pr., 2005).

13. W. C. Janssen and K. Popat, "Uplib: A Universal Personal Digital Library System," *ACM Symposium on Document Engineering* (New York: ACM Press, 2003), 234.

14. Good et al., "A Fluid Treemap Interface for Personal Digital Libraries."

15. B. B. Bederson et al., "Ordered and Quantum Treemaps: Making Effective Use of 2D Space to Display Hierarchies," *ACM Transactions on Computer Graphics* 21, no. 4 (2002): 833–54.

16. L. E. Good et al., "Zoomable User Interface for In-Depth Reading," *JCDL'04, June 7–11, Tucson, Ariz.* (New York: ACM Pr., 2004)

17. L. Marks et al., "ActiveGraph: A Digital Library Visualization Tool," *International Journal on Digital Libraries* 5, no. 1 (Mar. 2005), 57–69.

18. Ibid.

19. E. R. Tufte, *The Visual Display of Quantitative Information* (Cheshire, Conn.: Graphics Pr., 1983).

20. H. Shiaw et al., "The 3D Vase Museum: A New Approach to Context in a Digital Library," *JCDL'04, June 7–11, Tucson, Ariz. (*New York: ACM Pr., 2004).

21. Ibid.

22. K. Borner et al., "Collaborative Visual Interfaces to Digital Libraries," *JCDL'02, July 13–17, Portland, Ore.* (New York: ACM Pr., 2002).

23. Y. Feng and K. Borner, "Using Semantic Treemaps to Categorize and Visualize Bookmark Files," *Visualization and Data Analysis 2002: 21–22 January 2002, San Jose, USA* (Proceedings of SPIE, v. 4665) (Bellingham, Wash.: SPIE—the International Society for Optical Engineering, 2002), 218–27.

24. A. Veling, "The AquaBrowser—Visualization of Dynamic Concept Spaces," *Journal of AGSI* 6, no. 3 (1997): 136–42.

25. B. Eden, "3D Visualization Techniques: 2D and 3D Information Visualization Resources, Applications, and Future," *Library Technical Reports* 41, no. 1 (2005).

26. E. Bertini et al., "Visualization in Digital Libraries," www.dis.uniromal.it/~delos/docs/ivdls_book_chapter.pdf (accessed Jan. 12, 2006).