

Visualizing Digital Library Search Results with Categorical and Hierarchical Axes

Ben Shneiderman*, David Feldman, and Anne Rose

Human-Computer Interaction Laboratory, Institute for Advanced Computer Studies

*Department of Computer Science

and *Institute for Systems Research

University of Maryland, College Park, MD 20742

<http://www.cs.umd.edu/hcil>

ben@cs.umd.edu

Abstract

Digital library search results are usually shown as a textual list, with 10-20 items per page. Viewing several thousand search results at once on a two-dimensional display with continuous variables is a promising alternative. Since these displays can overwhelm some users, we created a simplified two-dimensional display that uses categorical and hierarchical axes, called *hieraxes*. Users appreciate the meaningful and limited number of terms on each hieraxis. At each grid point of the display we show a cluster of color-coded dots or a bar chart. Users see the entire result set and can then click on labels to move down a level in the hierarchy. Handling broad hierarchies and arranging for imposed hierarchies led to additional design innovations. We applied hieraxes to a digital video library used by middle school teachers and a legal information system.

1 Introduction

Digital library researchers face numerous challenges in helping users find what they want. Among them are designing meaningful overviews so that patterns can be easily recognized, creating comprehensible interfaces to specify what they want, and providing effective displays of search results. The visual information seeking mantra: overview first, zoom and filter, then details on demand provides a good starting point, but interpreting the mantra has produced a remarkably diverse set of research and commercial interfaces (Shneiderman, 1998; Card, Stuart, Mackinlay, Jock, Shneiderman, 1999; <http://otal.umd.edu/Olive>). Within this body of work, two important strategies are (1) two-dimensional and (2) browsers for hierarchical data sets.

Two-dimensional visualizations with meaningful axes have been an effective visualization strategy because thousands of items can be shown at once. Typical axes make use of continuous variables (for example, relevance or publication date) or categorical variables (for example, language or name of journal). The labels on such axes are cues to searchers and may contain valuable information if they are ordinal variables. Two-dimensional visualizations with non-meaningful axes have been tried in systems such as BEAD (Chalmers and Chitson, 1992), BIRD (Kim and

Korfhage, 1994), and ThemeScape (Wise, Thomas, Pennock, Lantrip, Pottier, Schur, and Crow, 1995) (commercial version is at www.cartia.com). Clusters can be meaningful, but users can make mistakes in assuming that distances between clusters have strong significance.

Hierarchical browsers, such as Yahoo!, have met with success and are common, as hierarchical data abound. Hierarchies have the potential to reduce complexity by organizing related information into comprehensible structures. Some hierarchies, such as the Dewey Decimal System, are fairly straightforward trees. Others, such as the National Library of Medicine's Medical Subject Headings (MeSH), the Universal Medical Language System (UMLS) (Pratt, 1997), and the ACM's Computing Reviews Categories are partially hierarchical, multi-hierarchical (organized simultaneously into multiple hierarchies), or otherwise complicated. Hierarchies are typically shown by textual lists of one level at a time, by node-link diagrams, cone trees (Hearst and Karadi, 1997) or by an outliner in which levels and branches can be expanded and contracted (Nation, Plaisant, Marchionini, and Komlodi, 1997).

In this paper we apply hierarchical axes, also referred to as *hieraxes*, to combine the power of a hierarchical browse tool with a two-dimensional visualization. This combination preserves visual overviews and enables users to rapidly comprehend search results. We focus our attention to search results that may include 100 to 10,000 items, but we anticipate that hieraxes may be useful for viewing much larger sets of items or entire digital libraries. We describe our initial work on a digital video library for middle school teachers, our current implementation for legal information, and discuss design alternatives for hieraxes.

2 Previous Research

Two-dimensional displays with meaningful axes have been used successfully to show document spaces, such as in Envision (Fox, Hix, Nowell, Brueni, Wake, Heath, Lenwood, and Rao, 1993; Nowell, France, Hix, Heath, Lenwood, and Fox, 1996) and FilmFinder (Ahlberg and Shneiderman, 1993). FilmFinder's x-axis was fixed as the year of a film's production, and its y-axis was the film's popularity (number of video rentals, on a scale of 1-9). Each spot on the graph was

color-coded by film type (drama, action, musical, etc.). The commercial expansion of this tool (Ahlberg and Wistrand, 1995), Spotfire (<http://www.spotfire.com>) allows any spreadsheet data set to be visualized. Spotfire allows both continuous and categorical variables to be put on either axis or used for color and size coding (continuous variables are color-coded using a spectrum, categorical variables with a set of discrete colors). Shape and rotation are available options, but these codings are less easily perceived. Changes to the axes or codings can be made in a few seconds.

Our experience and Spotfire’s commercial success shows that it is powerful, but that some training is needed. Moreover, some users find the rich display and extensive controls (including control panels, sliders, and zooming) confusing. Additionally, the potential clutter of color-coded, size-coded, overlapping spots requires users to plan their visualizations carefully. Labels and lines connecting related items add further challenges for novice users.

Tree structures are commonly presented as indented outlines that can be scrolled and opened or closed to expose or hide levels – as in Microsoft’s Windows Explorer, Allen’s “Dewey GUI” (Allen, 1995a) or WebTOC (Nation et al., 1997). Trees represented as node-link diagrams are useful, but as they grow to include thousands of nodes and many levels, layout and navigation problems become serious. Treemaps ensure a fixed-size presentation, but are complex, and many users require 15-20 minutes of training (Johnson and Shneiderman, 1991). Allen (1995b) described two “facet space” browsers for exploring multi-hierarchical (multi-dimensionally hierarchical) data sets. However, these tools did not allow users to view their data along categorical axes.

On an axis, a hierarchy must be represented one-dimensionally, so we chose the metaphor of opening and closing folders to move up and down the hierarchy, displaying one level at a time.

3 The Baltimore Learning Community

The need for hieraxes arose in designing user interfaces for the Baltimore Learning Community. This digital library consists of approximately 2000 educational resources, including digital video, websites, and images, for a community of thirty middle school teachers. In addition to several categorical and continuous fields, each resource was also indexed according to a three-level curriculum hierarchy (subject/topic/subtopic) to make searching easier. Teachers needed to be able to find resources of interest quickly, without the burden of complex controls. In response to usability testing with 14 teachers (Rose, Ding, Marchionini, Beale Jr., and Nolet, 1998) we created the first version of Dotfire, a subset of Spotfire that supports limited hierarchies (Figure 1).

Dotfire 1.0 uses categorical variables to organize search results on a two-dimensional grid. Categorical axis variables reduce the number of XY regions and increase overlap (i.e., many data points are mapped to the same XY region). In Dotfire 1.0, displaying each region as a cluster of dots (one for each data point in the region) solves the overlap problem. If there are more than 49 items in a region, a bar chart showing the frequency of each type of object replaces the cluster. Users can click on a cluster or bar chart to reveal a list of resource titles. Selecting an item in this list provides a more detailed description.




Figure 1: Dotfire, displaying a single hierarchical y-axis at the top level

Dotfire 1.0 supports basic hierarchical axes. In Figure 1, the y-axis shows science topics (the second level of the curriculum hierarchy). Users can refine their searches by clicking on the folder icon to the left of an axis value, which expands the selected category and filters the display. For example, clicking on the “Animals” folder changes the y-axis to display only Animal subtopics (Birds, Mammals, etc.) and filters the result set so only Animal resources are displayed (Figure 2). Users can click on the “up” icon at the top of the axis to return to the science topic level.

However, Dotfire 1.0 only handles the three-level curriculum hierarchy, and does not generalize to other hierarchies. It displays the immediate supercategory of the user’s current location, but not the complete browse history. And it does not handle axis variables with large value sets (more than 15). Even with a relatively small set, all the values are added to the axis, thereby forcing the potential truncation of long names (Figure 3).




Figure 2: Expanding the “Animals” Category on the y-axis in Dotfire




Figure 3: An Overcrowded Axis

4 Legal information digital library

Our interest in hieraxes grew in conducting research for WestGroup Publishers. Their digital library provides access to 12 million online cases, numerous statutes, and other valuable resources for the legal community. However, searching this digital library requires careful thought; a search result set could easily include tens of thousands of items, and users have pressing demands for comprehensiveness in their legal research. Finding a few relevant cases is rarely sufficient. Many search situations require review of every potentially relevant case and an understanding of the range of existing precedents.

WestGroup maintains an elaborate hierarchy of terms called the Key Number system. As this hierarchy is used as a

topical index of cases and headnotes, its inclusion in our legal research system was crucial. The first three levels (473 items) are meaningful and familiar to many lawyers, but still complex. The entire 85,000-item hierarchy is many more levels deep, extremely complex, and specialists become familiar only with branches relevant to their work. A portion of the hierarchy is displayed below with the total number of branches for each item shown in parentheses:

- Persons (5)
 - Natural Persons (16)
 - ...
 - Drugs and Narcotics
 - Drugs and Druggists in General (31)
 - Narcotics and Dangerous Drugs (145)
 - ...
 - Search under warrant
 - In general
 - Issuance, requisites, and validity of warrants in general (2)
 - Affidavits, complaints, and evidence for issuance of warrants (5)
 - In general
 - Probable cause in general
 - Error or ambiguity; omissions
 - Informants
 - In general
 - Personal knowledge
 - Reliability; corroboration
 - Confidential or unnamed informants
 - ...
 - Execution and return of warrants in general (5)
 - Disposition of property
 - ...
 - Food
 - ...
- Property (4)
- Contracts (4)
- Torts (0 at level 2, 16 directly at level 3)
- Crimes (0 at level 2, 56 directly at level 3)
- Remedies (13)
- Government (6)

Both depth and breadth are variable: a particular item may have as few as 2 branches or over 100, and the tree is as shallow as 2 levels in some places and as deep as 7 in others. Items are not always ordered alphabetically or numerically.

Other hierarchies relevant to legal research are geographical region and court district. In addition, some ordinal non-hierarchical fields (such as dates, relevance, numbers of citations, and names) can have hierarchies imposed on them to improve ease of use (see Imposed Hierarchies below).

In trying to use Dotfire with this data set, we realized that more comprehensive hierarchy support would be necessary, and began development on the general purpose hieraxis widget. We made some modifications to Dotfire itself, and then wrote the Hieraxis widgets in Java, including an API to allow them to work with diverse applications.

5 Hierarchical Axes: Hieraxes

In response to the challenge of serving legal information searchers, we developed a general purpose hierarchical axis widget for Dotfire and other applications, termed a hieraxis. Our goal was to support hieraxes on both x and y axes in a convenient manner that preserved screen space for viewing result sets, and to provide full-featured widgets for applications which meet the following criteria:

- Includes a two-dimensional visualization for which it is desirable that the user be able to change axis variables
- Data must include some categorical fields

As mentioned above, one shortcoming of Dotfire's original hierarchy support was the potential to get lost in a confusing hierarchy. To alleviate this, we added a simple browse history to the hieraxes. For each axis, users see their current location in the hierarchy in a drop-down list and as an indented series of labeled folder icons. Users can zoom out to any supercategory of their current location by selecting from the drop-down list or clicking on a folder (see Figure 4). Users can also zoom out one level by clicking on a standard "Up Level" icon near the browse history drop-down.

Any feature that allows users to view a subset of the entire visualization risks confusing them, causing them to lose their place. The browse history can help users remain comfortable with their location in the data set despite the fact that they may not be viewing the entire data set. By starting at the top, making all data exclusions user-driven, and providing an easy-to-use history we hope to reduce disorientation.

5.3 Hieraxes and Filtering

Some visualization tools (Dotfire, Spotfire) have dynamic axes, for which the variable can be changed rapidly. With

- Data set must be large enough to necessitate some sort of management strategy

Note that we do *not* include hierarchical fields in this list, as the categorical features of the Hieraxis widget are useful without the hierarchical ones, and the hierarchical functionality need not manifest itself.

5.1 Categorical Zooming

Visualizations have been tried with both continuous and categorical/discrete axis variables. Though zooming on one or both axes has been a familiar concept for some time with continuous variables, it has not been used with categorical variables. Partly this is because "zooming" generally applies to continuous spaces. However, drilling down into a hierarchy is, in essence, categorical zooming. Such a process is already familiar to users in browse tools such as Yahoo, but not in conjunction with visualizations. With hieraxes, we have applied it to the axes of a two-dimensional visualization and explored the results (Figure 4).

5.2 Browse History: Maintaining Context

hieraxes, application developers can employ a powerful data filtering mechanism that allows a visualization to serve simultaneously as search tool and results display.

If users filter such a display using categorical zooming, then select a new axis variable, the new variable's values will be displayed, but the displayed data will not increase; whatever data has been excluded by previous operations will continue to be excluded. Through successive category expansions (exclusions by zooming) and re-categorizations (changes of axis variable), users can graphically locate the information they desire. In this way we can continue to blur the unnecessary line between query and results. However, this approach might eventually need to be coupled with a more sophisticated search history, since the user's actions can no longer be represented accurately as a linear series of category expansions.

Dotfire does not employ this filtering mechanism, as such functionality is somewhat outside its scope: if the user selects a new axis variable, the axis returns to the entire data set. However, our Hieraxis widget supports this filtering technique.




Figure 4a: Dotfire 2.0 with Hieraxes widgets, showing hierarchical data on the y-axis



Figure 4b: Dotfire 2.0 with hieraxes after two category expansions on the x-axis




Figure 4c: Dotfire 2.0 showing PreviewPane after user clicks on a cluster

5.4 Overcrowded Axes (Broad Hierarchies)

Often, categorical axis data is extensive: hundreds, even thousands of discrete axis points, either as a single level of a hierarchy or a non-hierarchical axis variable. If the data is large enough, no labeling scheme can possibly enumerate all data points or distinguish among regions of the visualization (see Figure 3). We considered several solutions and implemented several of them (Figure 5).

5.4.1 Tool Tips (Figure 5a)

For slightly crowded axes, a solution akin to Microsoft's tool tips is effective. Each axis label can be truncated to fit in the available space; moving the mouse pointer over a label reveals its full text, temporarily obscuring adjacent labels. However, if space becomes even more limited and no more than a few letters of a label can be displayed, users will not have adequate information to determine which labels interest them. If the axis becomes crowded enough, space may not be available for even a single letter.

A y-axis variant on the tool tip approach involves showing all tool tips at once when the user mouses over the axis, or temporarily expanding the axis width to accommodate the labels. Dotfire implements standard tool tips on the x-axis, and temporary axis expansion on the y-axis.

Additional solutions exist when axes are only slightly overcrowded. Angling the axis labels, labeling in two or three staggered rows, and similar strategies can increase the number of readable labels significantly. But, as with tool tips, these techniques are ineffective when the number of items becomes too large, and users are once again faced with overcrowding.

5.4.2 Imposed Hierarchies (Figure 5b)

There may be an appropriate hierarchical *grouping* – alphabetical, chronological, numerical, etc. – which can be *imposed* on the data before transfer to an axis, creating a hierarchy on the fly. If the grouping can be done in a meaningful manner, the groups will be comprehensible to users, who can then expand a group as they would a category in a semantic hierarchy. In particular, categorical but *ordinal* variables work well with this method: fields like date, age, rank, etc. which can easily be ordered, and therefore grouped, in a meaningful way (see Figure 5b).

5.4.3 Paging (Figure 5c)

Another possibility is *paging*. If a display has room for 16 columns and the axis field has 80 items, the first 15 columns can be displayed as usual, with the 16th representing the remaining 65 items. The user will be able to see from the visualization how much data is contained in those 65 columns, and can move to the next “page” of data by clicking on a folder or arrow icon next to the 16th item in the x-axis. The distinction between the first 15 columns and the 16th column can be further highlighted by a different background color. If users are on a page other than the first or last, the first and last columns will contain the contents of all previous and remaining pages, respectively. This strategy has the advantage of maintaining a display of all data at all times, but a potential disadvantage of user disorientation due to too much paging. Dotfire implements paging.

5.4.4 Scroll/Zoom (Figure 5d)

A fourth strategy is a combination scroll/zoom technique, using range (double-box) sliders as Spotfire does. All items are placed on the axis, but compressed together with labels minimized or removed when the entire display is visible. Users can zoom in on an area of the axis by decreasing the range of the slider and scroll along the axis by sliding the slider. At any time they can easily return to the full visualization by increasing the slider's range to its maximum. This strategy minimizes disorientation by maintaining users' connection with the full visualization and tightly linking the potentially confusing scrolling with zooming. However, with somewhat smaller data sets (16-100 items) this method is awkward as both scrolling and zooming become "chunky" and confusing.


5.4.5 Implementation Strategy

Our primary design for handling overcrowded axes uses paging, supplemented by tool tips/axis expansion. Though the

other two strategies seem more compelling when the data fits them – scroll/zoom when the data is large enough, imposed hierarchies when the data is well-suited to grouping – paging seems to be the most scalable. It behaves in much the same manner with 17 items as it does with 17,000. Since moving sequentially through thousands of pages is tedious, users can jump to a specified page (see Figure 5c). Users control the strategy by selecting from a pull-down menu.

A scenario for a legal information seeker might deal with a topic such as "user interface copyright infringement." The search would produce a 4,000 items in the result set. Displaying these on a grid with year along the x-axis and Key Number along the y-axis immediately reveals the dramatic growth of cases on this topic in recent years. The results fall into the high level categories of Property and Contracts. Clicking on Property yields the next level of detail and a clarification that Intellectual Property is the major heading, but some interesting special cases are also visible.

Figure 5: Managing Overcrowded Axes



6 Conclusions and Future Directions

Hieraxes offer a simple approach to browsing search result sets. Users can see an overview by color coded dots or bar charts arranged in a grid and organized by familiar labeled categories. They can probe further by zooming in on desired categories or switching to another hierarchical variable. Our early success with middle school teachers has encouraged our development of a general-purpose tool. Usability tests with legal information searchers will be conducted in late spring 1999.

Although the basic idea of hieraxes is appealing to many users, not every user is visually oriented. For example, some legal researchers will prefer to see 20 titles of cases at a time, rather than 2000 color-coded dots. For digital libraries that do not have existing hierarchical categorizations, the imposed hierarchies may provide some benefit, but paging through long lists of authors or other unique (or relatively unique variables) could be tedious. Rapid visual exploration may be useful for some users and tasks, but it might undermine the willingness of users to reconsider their initial search terms. These and the following issues will be examined in future work:

6.1 Making sure the user doesn't get lost.

Though much of our effort has focused on avoiding user disorientation, it may continue to be a problem. Prominent indications of the user's location, carefully chosen icons, and animation may help with this, as well as a more extensive and intuitive search/browse history.

6.2 Multiple-category expansion

In many cases, users may want to expand several (adjacent) categories at once. This can be accomplished by allowing users to band-select several categories at once, expanding them onto the axis together. This could prove especially useful for ordinal axis fields and well-constructed hierarchies in which proximity is meaningful.

6.3 Category expansion through direct user-visualization interaction

Forcing the user to click on the axes in order to expand the visualization is somewhat indirect. Instead, we could allow the user to select portions of the visualization itself to expand, on one or both axes (depending on the shape of the selection).

6.4 Three-dimensional hieraxes

Our current work extends Dotfire 1.0's original functionality by including robust x- and y-hieraxes, with support for greater depth and breadth, allowing users to visualize complex, multi-hierarchical data sets. However, addition of a third dimension, through an additional hieraxis or another method, might yield an even more powerful tool.

Acknowledgements

We acknowledge the support of WestGroup Publishers and the U.S. Department of Education Technology Grant (#R303A50051) to the Baltimore City Public Schools. We

thank Robert B. Allen for his thoughtful comments on a draft of this paper.

References

Ahlberg, Christopher and Shneiderman, Ben, Visual information seeking: Tight coupling of dynamic query filters with starfield displays, *Proc. CHI'94 Conference: Human Factors in Computing Systems*, ACM, New York, NY (1994), 313-321 + color plates.

Ahlberg, Christopher and Wistrand, Erik, IVEE: An information visualization & exploration environment, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 66-73.

Allen, Robert B., Two digital library interfaces that exploit hierarchical structure, *Proceedings of DAGS95: Electronic Publishing and the Information Superhighway* (1995).

Allen, R.B., Retrieval from facet spaces, *Electronic Publishing*, 247-257, 8, 1995. (presented at Electronic Publishing, 1996, Palo Alto CA).

Card, Stuart, Mackinlay, Jock, and Shneiderman, Ben, *Readings in Information Visualization: Using Vision to Think*, Morgan Kaufman Publishers, San Francisco, CA (1999).

Chalmers, M. and Chitson, P, BEAD: Exploration in information visualization, *Proc. ACM SIGIR Conference*, ACM, New York (1992), 330-337.

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

Fox, Edward A., Hix, Deborah, Nowell, Lucy T., Brueni, Dennis J., Wake, William C., Heath, S. Lenwood, and Rao, Durgesh, Users, user interfaces, and objects: Envision, a digital library, *Journal of the American Society for Information Science* 44, 8 (1993), 480-491.

Hearst, Marti and Karadi, Chandhu, Cat-a-Cone: An interactive interface for specifying searchers and viewing retrieval results using a large category hierarchy, *Proceedings of the 20th Annual International ACM SIGIR Conference*, ACM, New York (1997).

Johnson, Brian, and Shneiderman, Ben, Tree-maps: A space-filling approach to the visualization of hierarchical information structures, *Proc. IEEE Visualization'91*, IEEE, Piscataway, NJ (1991), 284-291.

Kim, H. and Korfhage, R., BIRD: Browsing Interface for the Retrieval of Documents, *Proc. of IEEE Symposium on Visual Language'94*, IEEE (1994), 176-177.

Nation, David A., Plaisant, Catherine, Marchionini, Gary, and

Komlodi, Anita, Visualizing websites using a hierarchical table of contents browser: WebTOC, Online *Proceedings of 3rd Conference on Human Factors and the Web*, (June 1997). <http://www.uswest.com/webconference/index.html>

Nowell, Lucy Terry, France, Robert K., Hix, Deborah, Heath, Lenwood S., and Fox, Edward A. Visualizing search results: Some alternatives to query-document similarity, *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Zurich, Switzerland, August 1996, 67-75.

Pratt, Wanda, Dynamic organization of search results using the UMLS, *Proceedings of the American Medical Informatics Association Symposium* (1997).

Rose, A., Ding, W., Marchionini, G., Beale Jr., J., Nolet, V., Building an electronic learning community: From design to implementation, *Proceedings of CHI'98: Human Factors in Computing Systems*, ACM, New York (1998), 203-210.

Shneiderman, Ben, *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, Addison Wesley Longman Publishers, Reading, MA (1998), Chapter 15.

Wise, James A., Thomas, James, J., Pennock, Kelly, Lantrip, David, Pottier, Marc, Schur, Anne, and Crow, Vern, Visualizing the non-visual: Spatial analysis and interaction with information from text documents, *Proc. IEEE Information Visualization '95*, IEEE Computer Press, Los Alamitos, CA (1995), 51-58.