

Experiences with starfield visualizations for analysis of library collections

J. Alfredo Sánchez^a, Michael B. Twidale^b, David M. Nichols^c, Nabani N. Silva^{a*}

^a Department of Computer Systems Engineering, Universidad de las Américas-Puebla,
Cholula, Puebla 72820 México;

^b Graduate School of Library and Information Science,
University of Illinois at Urbana-Champaign, Champaign, IL 61820, USA;

^c Department of Computer Science, University of Waikato, Hamilton, New Zealand

ABSTRACT

This paper presents a qualitative and formative study of the uses of a starfield-based visualization interface for analysis of library collections. The evaluation process has produced feedback that suggests ways to significantly improve starfield interfaces and the interaction process to improve their learnability and usability. The study also gave us clear indication of additional potential uses of starfield visualizations that can be exploited by further functionality and interface development. We report on resulting implications for the design and use of starfield visualizations that will impact their graphical interface features, their use for managing data quality and their potential for various forms of visual data mining. Although the current implementation and analysis focuses on the collection of a physical library, the most important contributions of our work will be in digital libraries, in which volume, complexity and dynamism of collections are increasing dramatically and tools are needed for visualization and analysis.

Keywords: Collections, starfields, large information spaces, libraries.

1. INTRODUCTION

Information visualization has produced various alternative interfaces intended to facilitate user interaction with large and complex information spaces. One of the most promising approaches to information visualization is based on the representation referred to as “starfield”, which maps large information spaces to a two-dimensional grid. Typically, various filters may be applied to this representation by manipulating graphical sliders to zoom in and out of areas of interest until relevant data elements are located [Ahlberg et al., 1992; Ahlberg and Shneiderman, 1994]. Starfield-based interfaces were first developed at the University of Maryland’s Human-Computer Interaction Lab (HCIL) and were successfully applied to seeking information in multiple domains, such as real estate databases or movie collections.

As we defined various interfaces for our digital libraries program, we decided to assess whether starfield-based interfaces might be used for exploring book collections in a physical library’s catalogue. In our initial review of the literature we found that information spaces for which starfields had been used were large (up to about 5,000 items) when compared to the number of items in a desktop environment, but still small if compared to the size of a university’s book collection, which typically comprises hundreds of thousands or even millions of items. A relatively small information space makes it possible for a starfield-based application to keep entire collections in main memory and perform recalculations and rendering in real time as the user manipulates sliders and filters in an interactive fashion.

We thus undertook a project to investigate ways to cope with scale problems that would arise if starfield interfaces were used for information spaces that are several orders of magnitude larger than those reported in the literature. The principal challenge with scale-up is to maintain the immediacy of direct manipulation interactions with a very large dataset. One of the results of this project was EVA2D, a visualization environment that implements and extends existing starfield-based interfaces to facilitate the exploration of large collections comprised by digital libraries. It is worth

* Further author information: (Send correspondence to J. Alfredo Sanchez)

J. Alfredo Sánchez: E-mail: alfredo@mail.udlap.mx, Telephone: +52-222-229-2666

Michael B. Twidale: E-mail: twidale@uiuc.edu, Telephone: 1 217 265-0510

David M. Nichols: E-mail: dmn@cs.waikato.ac.nz, Telephone: +64-07-858-5130

Nabani N. Silva: E-mail: nabanisilva@hotmail.com, Telephone: +52-222-229-2666

mentioning that other efforts to enhance visualization techniques for large visualization spaces have been reported recently, including hardware acceleration, non-standard visualization attributes and animation applied to “treemaps” and scatter plots [Fekete and Plaisant, 2002]. EVA2D introduced four main features with respect to existing starfield interfaces: pre-computation of graphical data, direct zooming, precision filtering mechanisms, and quasi-immediate feedback. More details on the results of introducing these features and on the implementation of EVA2D have been reported elsewhere [Silva et al., 2003].

In this paper we describe a formative evaluation study of the uses of EVA2D and discuss lessons learnt and implications for the design of visualization environments intended to support analysis of very large information spaces.

In Section 2 we provide additional context for our study in terms of our organizational setting and salient related work. Section 3 offers a description of the EVA2D visualization environment as used in the study reported in this paper. Section 4 discusses a number of observations on the use of starfields for library collections prior to our experimental design, which is presented in Section 5. Then, Section 6 presents the main results of the study. Section 7 goes further to discuss design implications and Section 8 provides a summary of conclusions and planned future work.

2. BACKGROUND

The research and development program that frames the work reported in this paper has brought physical and digital libraries together by introducing novel services and interfaces for a wide user community of a formerly traditional library. The interface to a library collection has to serve many roles. Conventionally, online public access catalogues (OPACs) have been developed to support known-item searching. That is the situation in which users have a good idea of what they want and just need to know whether it is available and if so, how to obtain it. The user may have imperfect information about the desired item, or even wrong information, and so the OPAC should be designed to support less than ideal queries. Furthermore, the users of an OPAC may not be very familiar with how it works, may be unwilling to invest much effort in learning how to use it and may lack key skills and concepts about how to perform efficient searches. A well designed OPAC should support inexperienced as well as expert use.

However there are other kinds of library use that are poorly supported by current OPACs. OPACs support search far better than they support browsing – looking for a resource that the user is unable to specify precisely in advance, or where their information need co-evolves with what they find as they start looking. The limited interface of an OPAC as a powerful search engine enables users to very rapidly home in on particular books, but fails to give users a qualitative sense of the resources of the library as a whole. A useful comparison is between two kinds of physical library; a closed stack and an open stack collection. In a closed stack collection (often a collection of important, valuable or sensitive materials), the user must ask at the desk for what they want, and after a delay those materials are delivered to them. The user does not need to know how the books are organized behind the scenes, but serendipity is eliminated. In an open stack collection (by far the most prevalent kind) the user can wander around the book stacks themselves, looking for particular books, but also just getting a sense of what is on offer, and noting related books shelved near to a particular book of interest. The desirability of open stack collections for numerous reasons has been noted within professional librarianship for many decades.

It has been in this context that we decided to explore the possibilities of visualization of physical and digital collections, particularly using starfields. We believe that interfaces like EVA2D have the potential to help users gain a qualitative sense of the extent, or coverage of a library, analogous to the impression gained by wandering around the book stacks. Indeed in some ways it might even be better than walking around. As a graphical user interface, it allows users to switch between different kinds of perspectives that would require considerable time, effort or expertise to obtain by walking around. Although the current implementation and analysis focuses on the collection of a physical library, we believe that ultimately the greatest contribution of this approach will be in digital libraries where the option of walking the stacks is not even possible.

Chen and Börner [2002] note several problems associated with visual interfaces to digital libraries; in this paper we are concerned with four in particular: scalability, labeling, individual differences and evaluation. In terms of scalability we are visualizing the entire collection of a library, at least an order of magnitude greater than previous starfield studies reported in the literature [Ahlberg and Shneiderman, 1994]. Fekete and Plaisant [2003] report that 10^4 is a typical number of items for many systems. The labeling issue arises both in labeling individual items and in labeling the axes; an issue also reported by Shneiderman et al. [2000]. Evaluating visual interfaces is a difficult task given their

complexity and users' unfamiliarity with new interaction mechanisms and there appear to be relatively few qualitative studies of starfield/dynamic query interfaces. Two studies that do report qualitative results are [Shneiderman et al., 1992] and [Shneiderman et al., 2000]; which we compare with our findings later in the paper (Section 6).




Figure 1. An overview of the collection of a library using a starfield display in EVA2D.




Figure 2. An EVA2D collection showing the selection of a region of the display.

3. A STARFIELD-BASED INTERFACE TO A LIBRARY CATALOG

EVA2D (a Spanish acronym for “Visual exploration mapped to a 2D representation”) is a layered system designed to experiment with starfield interfaces for datasets of various sizes and characteristics with only minimum reprogramming effort. Currently it is being used as an alternative interface to the catalog of a physical library that comprises around 200,000 books. Figure 1 illustrates an overview of the library as displayed by our visualization environment. In Figure 1 every small dot in the “starfield” represents one or more books in the catalog. Dots (books) are placed in the starfield according to the values of the attributes represented by the two main axes. In the figure, the horizontal axis represents author names ordered alphabetically, whereas the vertical axis represents subject categories ordered according to the scheme used by the US Library of Congress. Sliders (known as “zoom bars” [Ahlberg, 1996]) next to each axis allow the user to adjust the range of values to be displayed. Additionally, sliders on the panel to the right represent other book attributes (e.g. edition, publisher, etc.) and allow for further dynamic filtering of the books being displayed. When the cursor is positioned over a book, its cataloguing data is displayed, as shown in Figure 1.

It is also possible to reconfigure the starfield and obtain alternative views of the library by changing what the main axes represent. Figure 2 shows an alternative view of the library in which the horizontal axis represents the year of publication of each book. In this case, the sliders have been used to restrict the view to books published after 1910. Additionally, a group of 444 books “near” the category identified as “neurology” and published between 1955 and 1980 have been selected by drawing a rectangle around them. Cataloguing data for each of the selected books can be examined by using the bottom right panel. The two horizontal bars at the bottom of the interface indicate the total number of books displayed in the starfield and the number of books selected by the user, respectively.

Selecting an area of the starfield after clicking on the magnifying glass (located at the top right corner of the interface) is equivalent to adjusting simultaneously both the horizontal and vertical sliders, as shown in Figure 3. Direct selection of an area of 2-D visualizations, a distinctive feature of EVA2D, had been identified as a potential future direction for user interaction in this style of interface [Shneiderman et al., 2000]. It should be noted that in addition to restricting the range of values being displayed, the labels that indicate the values on each axis progressively show greater detail as the user gets closer to an area of the library. This is particularly interesting for the case of subject categories: whereas in Figure 2 “neurology” appears as the only category related to medicine, by zooming into an area in that vicinity the result in Figure 3 shows a larger number of related categories.



Figure 3. An EVA2D collection after zooming in.

4. OBSERVATIONS PRECEDING THE USABILITY STUDY

For the largest collection sizes we have visualized, and despite the very careful system design, EVA2D is sometimes slower than we would like. Small delays in response can have a significant impact on acceptability and learnability of a system [Akscyn et al., 1988]. For example, when changing an attribute associated to one of the main axis, the processing delay can lead to a user believing that they have made a mistake and that that piece of the interface is not interactive. Delays in response can also be irritating, degrading the user experience and the way in which the system is used. We were well aware of these problems with the interface and did not need to undertake a usability study to discover or to validate them. We were concerned however that the delay problem would swamp other more subtle problems that we hoped to uncover in the user tests.

In part, EVA2D had been developed as a proof of concept, to investigate whether it was possible to maintain the power of a starfield visualization on a very large collection, both in terms of data representation and responsiveness to user input. The success of the implementation allows us to consider in more depth the kinds of uses that such a visualization can be useful for, bearing in mind that for datasets like library catalogues the OPAC already exists as a powerful tool for supporting search. Our study aimed to understand more about actual and potential uses to inform future redesign.

Figure 2 illustrates how a starfield visualization can help give various kinds of overview qualitative senses of a collection or a part of it. Even a superficial examination of the figure reveals that the collection appears to have increased in depth over time (more dots on the right than on the left), and that there are certain historic specializations (some long horizontal rows of dots). We wanted to see if users could interpret such kinds of data.




Figure 4. Spotting errors in data.

In exploring the interface for the study design phase we paid particular attention to the use of publication date on the x-axis. This way of laying out a collection with a temporal dimension seemed to offer many possibilities for gaining a qualitative sense of the depth and coverage of the collection. Similarly, the use of a temporal filter enables a complex visualization in the main window to consider, for example, patterns in acquisitions of recent books in technical areas. In exploring in more detail the older part of the collection, we noticed that there was a cluster of books located on the y-axis, implying that they were published in 0AD (Figure 4). Clearly that was wrong, and by clicking on them it became obvious that the cause was a null date being interpreted as a value of zero. This led us to realize that the starfield visualization could be used in even more ways than we had expected, and that one was as a tool for exploring aspects of data quality within a collection. Data quality investigation can be done with straightforward queries (such as “show all books with a date field of zero or null”), but a visualization allows for the discovery of anomalous patterns that can

reveal new kinds of error to be investigated in more depth. As such, the tool is acting more like a data mining visualization interface (e.g. [Ahlberg et al., 1992][Beale et al., 2004][Cox et al., 1997])—a function that had not been anticipated in its original design. This functionality is not available in related developments, such as the commercial product described in [Smith, 2004].

5. USABILITY STUDY DESIGN

An opportunity arose to test the system on a group of fourth- and fifth-year computer science students taking a course on interface design. As part of their preparation for running their own usability tests, they were to have the experience of being a usability subject. We developed our usability study to fit within those pedagogic constraints. The students were introduced to the system by a short talk in a lecture a week before the study, accompanied by a paper describing the system's design and functionality ([Silva et al., 2003]), and subsequently by a short demonstration in class. 28 students used the system, working together in pairs. Each session took 50 minutes and consisted of a very brief introduction to the aims of the study (where the occasional slowness of the system and its consequences were explicitly noted), a set of tasks to perform and an open ended discussion at the end. Subjects were encouraged to think aloud while doing the tasks. With each task, if the students had difficulties in using the system, the experimenter would ask questions to initiate discussion about possible reasons for the confusion. The session was videotaped in a usability lab, capturing both the screen and the faces of the subjects.

Test questions involved setting up a particular visualization, interpreting and refining it. For example, subjects were asked to set the starfield to display the books according to categories and year of publication and then adjust it to show only those published after 1995. One question, after requiring a setup that showed only books published before 1801, asked about possible errors to see if students could interpret the data points on the y axis (see Figure 5).

It should be noted that this is a rather non-standard kind of usability test. The subjects were computer science students who were very comfortable with technology in general, although novice users of the application. The aim of the study was a formative rather than a summative evaluation, intended to gain a richer understanding of how the system could be used and how its usability and learnability could be improved. Computer science students are not the most representative of intended users of this application. Nevertheless, with suitable discounting, we believed that we might make useful findings. For example, any aspects of the interface that our subjects found somewhat challenging to learn or use are very likely to be at least as challenging for less technologically adept users. By contrast, we must be more cautious in generalizing from aspects of the interface that our subjects found easy to use.

The study was strictly qualitative and formative. We wanted to know more about the actual and potential use of the system, and to inform subsequent redesign and future more systematic user studies. As one of the students observed, it was more like a 'fishing expedition' than the more typical rigorously controlled experiments that they had studied in class.

6. RESULTS

The study provided very rich feedback on the potential of starfield visualizations. We have organized our observations into four general areas: graphical interface features, representation and analysis of the information space, support for diverse user strategies, and background of subjects.

6.1 Graphical interface features

Although our subjects had previously read about the starfield display, had heard a talk about it, and had seen a short demo in class, many of the subjects commented that the system only made sense to them when they actually started to use it themselves. The immediate feedback of choosing dimensions, moving sliders and mousing over items to get pop-up information was very important to get a sense of what the elements of the display meant, what they could interact with, and the overall meaning of the results shown.

Most subjects were able to quickly learn how to use the zoom bar, but they did have initial confusions. These must be noted because as computer science students they are particularly expert computer users who might even be expected to have no problems at all. Aspects of the confusion were that a zoom bar looks like a scroll bar but behaves differently. It

consists of many subcomponents (the two vertical bars, the two pairs of arrows at each end and the central block) each of which has a specific function. One very common error was that in using the zoom bar to restrict a view, the subjects would correctly attempt to manipulate the vertical bar component, but would pick the wrong one. Although common, this error was not serious, as the immediate feedback of the display made it clear that the view had been restricted, but not in the desired way, and so the subjects would then try the other bar and be successful. In many ways the interactions observed were analogous to those of computer novices encountering a conventional scroll bar for the first time.

The dynamic, animated changing of the display that resulted from adjusting the zoom bars contributed to the ease with which subjects understood the meaning of refining a view. Using the magnification box led to a redraw of the whole screen, which did not seem to cause confusion, but which we suspect would be helped by a rapid animated expansion rather than a simple switch. Animation imposes yet more of a time penalty on an already overloaded computation, however it carries great benefits for the learnability of the system.

As zooming occurs, the labels on the axis change. This can potentially be confusing, but as noted, the animation of the change helps in understanding the meaning of the process. This is most simply illustrated by the consequences on the date axis when zooming in (Figures 2 and 3). The system automatically adds in more date labels as space permits. These are chosen to take greatest advantage of the available space, and are equally distributed, but are not chosen for particular significance. Thus at a certain magnification there may be a date label for every 10th year, but these need not necessarily be particular round numbers (1960, 1970, etc.). For dates, this did not cause the subjects any problem, other than occasional passing comment about the scale looking funny. It is relatively straightforward to interpret and to interpolate the meaning of the intervening unlabelled positions. However there were two situations where this variable axis scaling and labeling caused confusion. A label present in the broader view may disappear in the narrower view, even though many more intermediate labels are visible. This is illustrated in the transition from Figure 2 to Figure 3. The selection in the wider view includes books published in 1955, a date that is labeled on the x-axis in Figure 2, but in the narrower view of Figure 3, although new labels appear (for 1958, 1968, 1978 and 1988), the 1955 label disappears. This does not cause a problem with dates (none of our subjects had difficulties in that situation), but it does cause difficulties for other values such as publisher name or category. An interest in a named value can lead to the user zooming in on, say, an area with the label 'Neurology', and at an intermediate level of detail this label disappears even as new intermediate nearby labels appear. This can be very confusing until users understand the mechanism and realize that on zooming in yet further, this label and all other labels will be visible.




Figure 5. Representation of multiple books by a single dot in the starfield.

6.2 Representation and analysis of the information space

Subjects were invariably surprised to discover that a dot could represent more than one item (e.g. multiple books published on the same year on the same subject). One of our tasks deliberately tested for this issue – without it we suspect that many students would never have understood this characteristic of the interface. Furthermore, they expected that when zooming in to a smaller part of the space, a multi-item dot would ‘unpack’ in the now available space to

create a set of single-record dots. This is a reasonable but incorrect assumption. On zooming down to more detail, some records that were represented by the same dot in an overview view do separate out, but only if they have distinctly different (but nearby) values. However even at maximum zoom, a view of subject categories against date, for example, will show a single dot to represent the set of four books published in 1974 on “Spain - History” that are in the collection, as illustrated in Figure 5. In this example, the information panel can be used to traverse each of the books associated to the selected dot. In the GRIDL system [Shneiderman et al., 2000] this problem is addressed by, where space permits, 'unpacking' the relevant data into a cluster of dots. However, GRIDL then needs a different representation (a bar chart) for large numbers of clustered items and necessarily uses more axis-space – e.g. for a range of dates [Shneiderman et al., 2000].

A problem that is unique to the subject category axis deserves further discussion. All other potential axes have a clear indication of their ordering. Dates are obvious, and publisher, title, etc are clearly in alphabetical order, even if this selective process of labeling is confusing. However categories can appear to be in no particular order. Most of our subjects remarked on this although they did not seem particularly concerned by it (perhaps because they thought it was an aspect of the prototype nature of the system). However this lack of understanding of the meaning and ordering of category and its representation did cause substantial confusion for an early task, which was to find a category for which the library had a collection extending over many years. Understanding how to interpret the visualization was not an issue. All groups quickly realized that they were looking for a long horizontal line as shown in Figure 2 by the truncated label for ‘Portugues’. Mousing over the label (which some but not all subjects discovered) reveals that this is actually the category “Portuguese Literature [PQ:9000–9999]”; this confirms the effectiveness of the 'Tool Tip' strategy for managing overcrowded axes [Shneiderman et al., 2000]. Zooming in to get more detail caused considerable confusion in cases where the zoomed-in view lost the label ‘Portugues’. Some subjects did not understand the idea of category at all, and others (determined by subsequent questioning and discussion) thought that the category labels were hierarchical, with zooming in revealing labels of subcategories. This is entirely reasonable, and indeed is our redesign recommendation for addressing this problem. However with this assumption, the current interface becomes confusing and seemingly arbitrary. The call number (e.g. PQ:9000–9999) seemed to be ignored by most subjects. When asked about it, most were able to guess what it might be, but in many cases subjects seemed surprised when it was drawn to their attention, implying that they had not given it prior thought. The meaning of categories and classification systems requires a certain level of background knowledge, which can be assumed for librarians and skilled library users, but as the study re-emphasized, not for all users. Despite the complexities of designing usable categorical zooming interfaces they have the potential to supporting 'overview plus detail' visualizations [Shneiderman et al., 2000].

The identification of data quality issues was achieved by nearly all subjects, but was far less smooth than those tasks involving interpretation of the qualitative patterns of the collection. We believe that this is because as students, our subjects normally consider a library as a set of resources to be accessed rather than as a database to be managed for the benefit of others. In passing we should note that many students remarked that they rarely used libraries and their catalogues for information. Thus as computer science students they were unrepresentative in that although they were expert users of computer applications, they were relatively naïve users of libraries.

6.3 Support for diverse user strategies

We had designed our tasks to test aspects of the system that we suspected would be problematic and would benefit from a more careful analysis to inform design (the meaning of the category selection being a prime example). However in running our tests we discovered that our subjects were extraordinarily creative and diverse in how they solved our tasks. They managed to invent ways to interact with the system to obtain the information required that we had not anticipated when designing the experiment. We consider this to be an unintended benefit of the current system design that we should work to enhance. There are often several different ways to refine a search. For example, to get more detail one can manipulate the zoom bars in different ways, or use the magnification rectangle, or use the filters to eliminate unwanted detail from the display. Some subjects discovered that the central portion of the zoom bar could be moved back and forth like a conventional scrollbar and that in conjunction with narrowing the focus with the vertical handles of the zoom bar this allowed an overview-plus-detail view by rapidly sliding the bar back and forth. When trying to zoom into the collection, a pair of subjects selected the magnifying glass option but instead of drawing a rectangle around the items to be included in a new view they clicked once on an area of the starfield, which resulted in an automatic zoom-in by 30%. From that point on they relied on this functionality whenever zooming in was needed.

The provision of many alternate ways of interaction with a novel interface seems highly desirable, although there is a potential trade-off in confusing users with too many options. The provision of immediate, often animated feedback appears to mitigate that problem in this case, although we must also be aware of the technical sophistication of our subjects.

Several subjects remarked that they would have preferred to ‘play around’ with the application before trying to do the assigned tasks. This is noteworthy, as a common recommendation in running end user studies is to always assign tasks or the activity will be meaningless to the subjects. As computer scientists, it may be that our subjects are anomalous in finding new computer applications inherently interesting and in feeling confident in learning new applications by exploration. However we also suspect that the direct manipulation aspects of the information visualization afford learning by exploring (confirming previous observations [Ahlberg et al., 1992]), and want to raise this as an issue for further study.

6.4 Background of subjects

The post-test discussions proved extremely successful. Texts on usability testing often caution against asking end users for design recommendations on the grounds that users are not computer scientists and so although very expert in what they do and what they might want, are not skilled in coming up with feasible, practical, and creative design ideas. However our subjects *were* computer scientists, and indeed had been exposed to ideas of usability and information visualization. They were more than happy to offer analyses, and design suggestions for improvements, even sketching out ideas as they went along. As such our test was less like a traditional end user experiment and more like a multi-expert technical review. Some of the many suggestions that the subjects provided are referred to in Section 7.

Subjects were also asked to comment on the kinds of situation when a starfield-like visualization would be useful. Again, this is not a recommended activity for regular end-users but proved valuable with computer science students. Comments included that the visualization was useful in getting a sense of the ‘weight’ of the collection, that it helped in seeing particular strengths of the collection, but also for spotting weaknesses or gaps, that by using a date filter it was possible to get a sense of recent books in the collection, and that in general it would be useful for browsing activities where either the user did not really know what they were looking for, or knew of one book and were interested in books ‘like’ it. The analogy of the old-fashioned card catalogue was used to describe how the visualization allows the discovery of items nearby to a known item.

7. DESIGN IMPLICATIONS

The user study generated a large number of ideas for design improvements. Some of the suggestions from the subjects were: provide multi-step undo, since when exploring an interface it is very frustrating when you do something wrong by accident; improve the categories labeling by explicit use of hierarchies; provide fisheye views for areas of interest while maintaining an overview of context; provide a thumbnail map of the complete collection with a rectangle showing the focus of the current screen; clarify the limits of zooming so that one cannot accidentally zoom in too much and see nothing; revise the scaling-labeling algorithm so that labels do not disappear on zooming in; and distinguish between dots that stand for one and for many books perhaps by color, shade or using a 3D representation; allow the user to display on the interface only those filters in use to avoid clutter; and provide a “hand tool” to enable panning of zoomed-in starfield.

As with all user studies, a careful analysis of the underlying causes of confusion yielded insights into possible design improvements. The confusions about labels in general and categories in particular as examined above illustrate the process. Other design ideas informed by the user study are: changing the appearance of the zoom bar to emphasize its different component parts: increase the animation of changing controls to emphasize learnability; yoking the changing of different controls so that when using one way to obtain a view (such as the magnify box selection), the other controls (such as the zoom bar) clearly change at the same time. A design trade-off to be explored is between additional, more emphasized and perhaps slower animation to improve learnability, but at the expense of slower interaction and a need for greater processing power that could be irritating for more expert users. An explicit learner mode is one option, but introducing a modality itself adds more complexity.

The ongoing question of “what is this kind of visualization and interaction good for?” has generated some new answers. The unexpected use of the visualization for data quality analysis and user-driven data mining deserves greater exploration, and the development of features to explicitly support it. Figure 6 gives a simple example of how powerful

visualizations can allow a skilled user to spot and investigate patterns in the data. In this screenshot, subject is plotted against title. By zooming in on interesting parts of the display and examining particular records it is easy to determine that the large diagonal line shows books whose title begins with the subject name and the smaller line shows books that begin with “the” followed by the subject name. The discovery that books on a subject often have a title that begins with the subject name may not be a particularly exciting finding, but it does dramatically illustrate how patterns observed in a visualization can be used as a spur to try and investigate and account for the underlying meaning in the data that produce them. A less conspicuous but potentially more interesting example is one that results from spotting a particularly dense cluster on the top left region of the starfield (enclosed by a rectangle). By zooming into this area the user obtains the starfield illustrated in Figure 7, which shows a significant collection of books dealing with diverse aspects of the United States.




Figure 6. Spotting patterns in data (note the dense area in top left corner).




Figure 7. Using EVA2D zooming functionality for visual data analysis.

8. CONCLUSIONS AND FUTURE WORK

The usability study, the analysis preceding the study, and the subsequent analysis of the findings has provided us with a very rich source of information to inform future redesign and further use tests. We had expected that the evaluation process would give us ideas for improvements to the interface and the interaction process to improve learnability and usability. It certainly did, but it also gave us an indication of additional potential uses of a novel visualization that can be exploited by further functionality and interface development.

Testing on computer science students, which was done for reasons of expediency (but with an awareness of problems of unrepresentativeness), proved in reality to be an extremely valuable source of insights and recommendations. By a focus on confusions and relatively minor difficulties that this rather expert group had with the application, we believe that we have identified some (but not all) problems to be addressed. Once the resultant improvements to the learnability and usability of the system have been made, the next stage is to run tests on librarians and skilled library users, graduate students in the arts and social sciences. This will allow us to detect and understand other more subtle learnability and usability problems and also to gain a greater understanding of the potential of novel kinds of information visualization to complement interacting with a large dataset via a conventional search-based interface. One appealing hypothesis to explore is whether librarians might be a better focus group for starfield visualizations than the library's end users. After all, they maintain the collections and might find the visual data mining aspects interesting.

ACKNOWLEDGMENTS

We would like to thank the students who participated in the study. We would also like to express our appreciation to researchers at the Human-Computer Interaction Lab of the University of Maryland for providing access to the documentation and source code of visualization tools. In particular, thanks are due to Ben Shneiderman and Catherine Plaisant for their support. This work has been partially supported by Conacyt (Project No. 35804-A).

REFERENCES

1. R. Akscyn, D. McCracken, and E. Yoder, "KMS: A distributed hypermedia system for managing knowledge organizations", *Comm. of the ACM* **31**(7), pp. 820-835, 1988.
2. C. Ahlberg. 1996. "Spotfire: An information exploration environment", *ACM SIGMOD Record*, **25**(4), 25-29.
3. C. Ahlberg and B. Shneiderman, "Visual Information Seeking: Tight coupling of dynamic query filters with starfield displays", *Proceedings of the Conference on Human Factors in Computing Systems (CHI'94)*, pp. 313-317, ACM Press, New York, NY, 1994.
4. C. Ahlberg, C. Williamson, and B. Shneiderman, "Dynamic queries for information exploration: An implementation and evaluation", *Proceedings of the Conference on Human Factors in Computing Systems (CHI'92)*, pp. 619-626, ACM Press, New York, NY, 1992.
5. R. Beale, A. Pryke and R.J. Hendley, "Evolutionary Approaches to Visualisation and Knowledge Discovery", *Proceedings of 6th Asia-Pacific Conference on Computer-Human Interaction (APCHI'04)*, pp. 30-39, Springer, Berlin, 2004.
6. C. Chen and K. Börner, "Top ten problems in visual interfaces to digital libraries", in *Visual Interfaces to Digital Libraries*, LNCS 2539, C. Chen and K. Börner, eds., pp. 226-231, Springer, Berlin, 2002.
7. K.C. Cox, S.G. Eick, G.J. Wills, R.J. Brachman, "Visual Data Mining: Recognizing Telephone Calling Fraud", *Data Mining and Knowledge. Discovery*, **1**(2): 225-231, 1997.
8. J-D. Feteke, and C. Plaisant C., "Interactive information visualization of a million items". in *The Craft of Information Visualization – Readings and Reflections*, B. Bederson, and B. Shneiderman, eds., Morgan Kaufman, San Francisco, CA, pp. 279-286, 2003.
9. B. Shneiderman, D. Feldman, A. Roseand and X.F. Grau, "Visualizing digital library search results with categorical and hierarchical axes", *Proceedings of the 5th International Conference on Digital Libraries (DL'00)*, pp. 57- 65, ACM Press, New York, NY, 2000.
10. N. Silva, J.A. Sánchez, C. Proal and C. Rebollar, "Visual exploration of large digital libraries collections", *Proceedings of the Latin American Conference on Human-Computer Interaction (CLIHC 2003)*, pp. 147-157, ACM Press, New York, NY, 2003.
11. T. Smith, "High Density Interfaces: More than Meets the Eye", *Tech Update*, June 2nd 2004, <http://techupdate.zdnet.com/hyperviewer.html?tag=tu.arch.link>.