

Supporting Information Structuring in a Digital Library

George R. Buchanan¹, Ann Blandford², Harold Thimbleby², and Matt Jones¹

¹ University of Waikato, Hamilton, N.Z.
{g.buchanan, mattj}@mdx.ac.uk

² UCL Interaction Centre, London, U.K.
{a.blandford, h.thimbleby}@ucl.ac.uk

Abstract. In this paper we present Garnet, a spatial hypertext interface to a digital library. Spatial hypertext systems support information structuring – the organisation of documents performed by a user to complement their information seeking. In the past, spatial hypertext systems have suffered from poor connectivity with information sources such as digital libraries. Conversely, digital libraries have provided strong support for document retrieval whilst offering little support for information structuring over the retrieved documents. Garnet provides an integrated environment for both seeking and organising information. We report on the results of a user study that elicits the response of users to a combined seeking and structuring environment. The feasibility of exploiting the information structuring of users to identify the interests of users is also investigated.

1 Introduction

The information seeking behaviour of library users has been studied for many years. This research has influenced the development of digital libraries and the facilities they provide. The focus of most digital libraries is upon supporting the discovery and retrieval of documents. Researchers such as David Ellis [7] and Carole Kuhlthau [12] have observed patterns within information seeking that are not simply about retrieving documents. For instance, Kuhlthau identifies the practice of *Collecting*, where documents are grouped together and organised. ‘Collecting’ supports a variety of tasks, ranging from the identification of topical themes to discerning ‘missing’ information. Researchers from a human-computer interaction background such as Malone [13] and Kidd [11] have observed patterns of document use in physical environments that co-ordinate and support tasks such as the identification of outstanding information needs and emerging topical strands. Examples of such patterns of use include positioning documents in such ways that they act as reminders of outstanding work, or piling together documents needed for the same task. Together, these different activities are known as information structuring.

Spatial hypertext systems support information structuring in an electronic environment. They provide a freeform visual workspace within which each document is represented by a shape. Examples of spatial hypertext workspaces will appear later in this paper. Implemented spatial hypertext systems include Pad++ [2] and VIKI [14]. Studies of VIKI have demonstrated that users employ similar patterns of document positioning in spatial hypertexts to those seen in physical environments by Kidd and Malone [15, 17].

We know that visual patterns of organisation carry between physical and digital workspaces. On the other hand, it is not clear whether information structuring's relationship with wider information seeking, and particularly with document retrieval, applies in digital libraries. Spatial hypertexts have seldom been connected to information sources such as digital libraries, and there is no reported work observing the relationship of information structuring and information seeking in an integrated environment.

Garnet, our spatial hypertext interface to a digital library, was created to allow us to observe the behaviour of users in a combined, digital, information and structuring environment. In [3] we introduced Garnet, describing its architecture and operation, an initial informal evaluation and a brief comparison of visual DL interfaces.

In this paper we undertake a more detailed formal user evaluation of our refined implementation, identifying work flows and information structuring behaviour, and update our review of Garnet's comparison to visual DL interfaces. This paper proceeds in four parts: first, the operation of Garnet is demonstrated in a simple example; second, we compare Garnet with existing visual interfaces to digital libraries; thirdly, the user study that we performed with Garnet is reported; the paper then concludes with a summary of our findings.

2 Garnet in Use

A pilot version of Garnet has been created, which is integrated with the New Zealand Digital Library Project's Greenstone software [20]. Greenstone is a comprehensive open-source Digital Library software system, supporting common actions such as full-text and index searching, and browsing in category hierarchies. Access to the digital library system is via a remote digital library protocol. As demonstrated in our earlier work [1], the Greenstone protocol can be trivially mapped to the three other common DL protocols – Dienst, Z39.50 and SDLIP – so Garnet could readily be integrated with alternative digital library systems that employ these other protocols.

We will now demonstrate the system in use.

2.1 Overview

In Fig. 1, we see a typical Garnet user session in progress; a window appears inside the main browser window. This window is a collection of materials that the user has recorded in the current, or a previous, session. Each document is represented by a rectangle containing some text that we term a 'label' for simplicity.

Within a collection, the user is free to place, size and colour each document label as they see fit – the space is entirely freeform. Labels can be moved and/or copied between collections in the usual way for similar direct manipulation environments. Document labels can be added explicitly by the user or through interaction with the digital library's search facilities.

Therefore, the user is free to use the document labels both in freeform structures of their own making inside collections, and in a more formal organisation by using the explicit hierarchical forms of a set of document collections. In Fig. 1 above, we have a collection called "Cattle", which has a column of three documents on the left-hand side and a pile of three documents near the bottom. The column and pile are structures

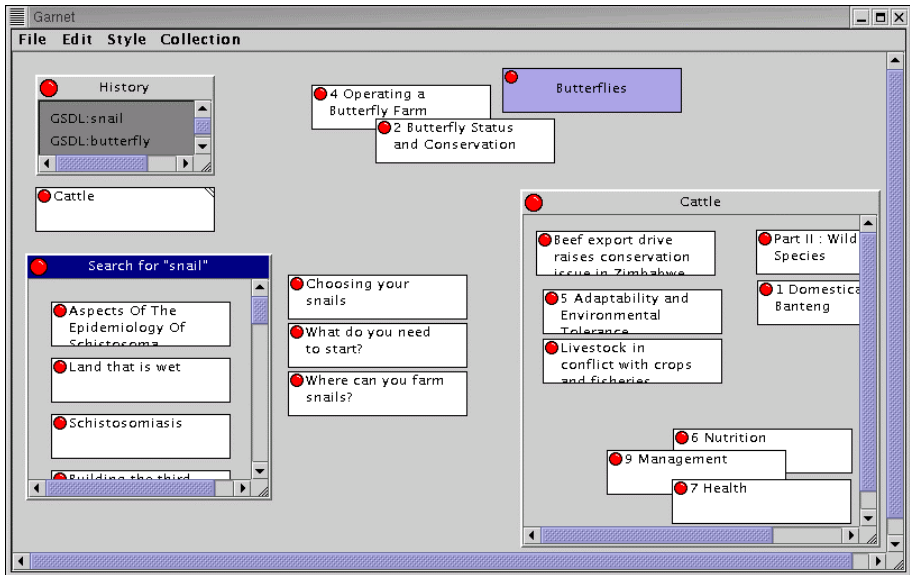


Fig. 1. A Garnet Client in Use

created by the user's exploitation of space – not features enforced by the system. The column idiom can also be seen in the root collection – on the left-hand side. Some use of colour can also be seen – e.g “Butterflies” – but the creator's intention in grouping and colouring is not clear to us as readers of the hypertext.

A search history appears on the top left-hand corner, and the current search for “Snail” is seen on the left-hand side of the main window. It appears like a normal collection, though its contents are selected and ordered by the digital library. As with most web and digital library searches, the search is ordered by relevance. Documents can be dragged from the search to the main workspace or a collection. Also, the user can delete items from the search list by clicking on the small red ‘blob’ on the top left of a document label. To read a document, the user double-clicks on its label. Garnet then displays the document in a separate window.

2.2 Demonstration of “Scatter”

Garnet can exploit the organisation done by the user in a novel manner. We can “scatter” a set of documents (including search results) from a selected window over the existing layout of documents in the workspace. A “Scatter” places the search documents near groups of existing documents which they have a strong similarity to.

In Fig. 2 we have selected a few useful-looking documents on the subject of snail farming on the main workspace, but let us suppose that a couple of questions remain unanswered. We have a plentiful supply of bananas which we would like to use, but we are not sure whether this food would be appropriate. If we did a naïve search, on “banana”, the initial results do not match our particular interest well (Fig. 2, left).

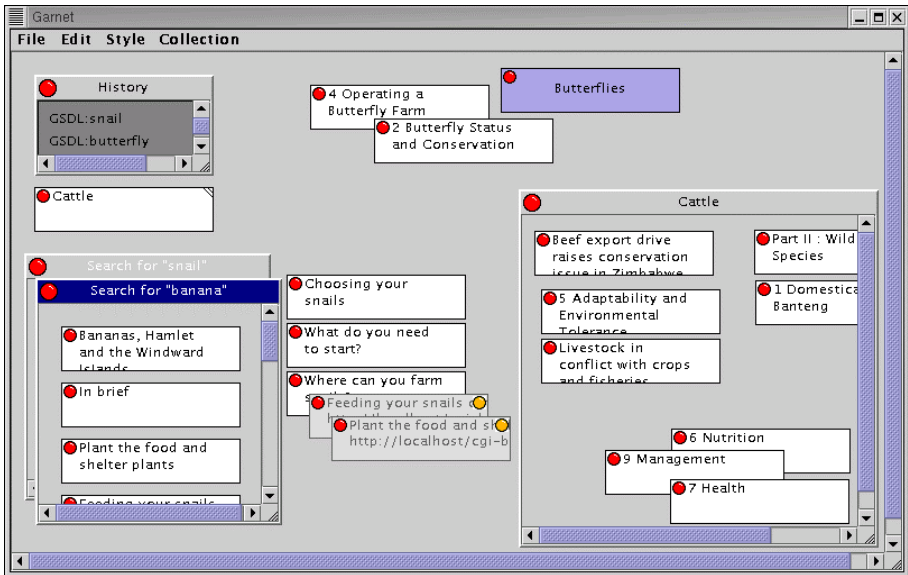


Fig. 2. An example “scatter”. Note the shaded document labels in the centre of the display

In fact, documents that relate to our interest can be found in both the “snail” and “banana” searches. However, these documents of interest may not appear at the very top of either list. Normally, we would have to try and re-work our query manually to make it more targeted. With Garnet, we can use the ‘scatter’ feature to discover any material similar to documents we have already selected. Or, in other words, Garnet can generate existing search terms or filtering to represent our user’s interests, based on the workspace layout they have already created.

Viewing Fig. 2, note the third item from the top of the “Search for ‘banana’” list: “Plant the food and shelter plants”. This item is related to the three documents on the main workspace (for clarity we’ve chosen something that is visible in this example). A “scattered” subset of the “banana” search results appears on the workspace in light grey. These documents closely matched the existing set of documents, which appear in white. Suggestions are displayed in this grey colour, and below and to the right of the group of documents that they are believed to be similar to.

We can now investigate the two suggested documents that are similar to the previously selected pair. As it happens, these documents would confirm that ripe bananas can indeed be used to feed snails. If we wanted to permanently add one or other suggestion to the workspace, we can click on the ‘blob’ which appears on the top right corner of each of the suggestions.

If we no longer wish to see the existing suggestions, or when another set of documents is scattered, the current suggestions are cleared.

In this section, the basic functions of Garnet have been briefly introduced. We will now compare Garnet to other visual DL interfaces, and follow that with a report of a user study of Garnet in use.

3 Visual Interfaces to Digital Libraries

Garnet, with its visual, graphical DL interface can be compared to other visual DL interfaces such as Daffodil [8], NaviQue [9], SketchTrieve [10] and DLITE [5]. Garnet is similar to these existing systems in certain ways, but has both new combinations of features and entirely novel ones. In this section, we review these existing systems and compare and contrast them with Garnet.

The developers of all these visual DL interfaces report that spatial hypertext systems have been an influence upon their design. The difference with Garnet is that it provides a whole range of spatial hypertext features, rather than a subset.

For example, Garnet follows spatial hypertext systems such as VIKI in providing a wide range of visual controls over the appearance of documents in their workspace. In comparison, the other DL interfaces provide limited scope for affecting the appearance of a document. This difference may seem ‘cosmetic’, but spatial hypertext research has demonstrated the importance of visual controls for the expression of the user’s perception of a document. The range of cues used by users in physical environments [11] also suggests that a rich visual control plays an important role in the information structuring that spatial hypertext supports.

The tools that support structuring of document sets are key elements of a spatial hypertext system. There are two main forms of document groups in spatial hypertexts: *explicit groups* where the group and its membership are directly and precisely defined; and *implicit groups* where the existence of the group and its membership are uncertain. Explicit groups can be found in many organisational systems – the folders of a filing system and the documents that they contain, or the subject hierarchies in a library. Implicit groups, on the other hand, are common when structure is transitory or provisional. Implicit groups are more a key feature of spatial hypertexts than explicit groups, given their role in supporting the task of structuring and organising a set of documents.

Garnet supports the identification of implicit groups of documents – e.g. documents placed close together in a pile or column. A spatial parser [17] identifies separate groups of documents in a single window or ‘collection’. No published information about the existing visual DL interfaces suggests that any of them have a spatial parser or other means of identifying implicit groups. In NaviQue [9], the user can manually identify a group of documents by selecting an area of the workspace. The group created is transitory, but the explicit selection that creates it means that it is in fact an explicit group.

Turning to explicit groups, NaviQue has only the transitory groups just mentioned. In DLITE [5] document groups only exist as the product of searches and other retrieval operations – explicit sets produced indirectly by the user’s action. Groups cannot have documents added or removed; neither can they be used as the target of a search. SketchTrieve [10] has a similar approach.

The information structuring support of existing visual digital libraries is thus weak when compared to spatial hypertexts: visual cues are restricted, implicit structuring is absent and explicit structures are often system- rather than user-controlled.

However, the representation of DL facilities such as search and browsing access are relevant to developing a spatial hypertext workspace for a digital library. Spatial hypertext systems have had little or no connection to information repositories such as DLs [18], and have not needed to represent the features of such a system in their

workspace. Garnet needs to do just that, and so we were influenced by these existing visual DL interfaces.

The Daffodil system [8] provides another novel visual interface to digital libraries. Users are presented with a number of strategies and tools which they can choose from to recover material from the digital library. An example strategy would be citation linking – which is delivered through the transparent (to the user) use of underlying search technologies. Browsing nodes and search result sets are presented in individual windows that contain vertical, textual lists of documents or child nodes – i.e. somewhat similar to a traditional web-based interface. Documents, as in DLITE for example, do not appear as individual objects that can be manipulated independently. It is unsurprising, therefore, that Daffodil does not support information structuring. How information structuring support, as provided by Garnet could be integrated with Daffodil's strategy-centred interface is a matter worthy of further investigation.

Daffodil also allows users to order search result sets in unorthodox ways – e.g. by similarity to a single selected document. This provides some common ground with Garnet's 'scatter' and 'find similar' features. NaviQue's 'Similarity Engine' provides a feature that will highlight documents that are similar to a selected set of documents. Again, this bears some similarity to Garnet's 'scatter' facility. However, Garnet matches documents in external libraries, not only in its own workspace, and also will either bring similar documents to a selected group, or conversely scatter the individual documents in the selected group to other groups that they match. Garnet also includes a spatial parser that can identify visual groups automatically.

Garnet therefore provides a much stronger set of information structuring tools, and a richer set of similarity tools than found in existing visual interfaces to digital libraries. We wished to discover the benefits that users perceived in having the facilities of a spatial hypertext interface to a digital library, and their response to the related textual similarity tools we provided. The next section will discuss the user study that we undertook to explore these issues.

4 User Study

Garnet provides a novel interface for a digital library, providing facilities for both information structuring and traditional information seeking. Though earlier studies in physical environments noted the frequent interleaving of these activities, we are not aware of any similar study in an integrated digital environment. In our earlier paper on Garnet [3] we reported the findings of informal, formative evaluations.

We present a new, formal study that we undertook to identify salient issues in integrating Garnet's digital library and spatial hypertext elements. This was a qualitative study to elicit design considerations and identify problems for further investigation. We followed a pattern of similar probing studies established in our previous DL work, e.g. [4]. A panel of ten participants was recruited, each studying a degree in psychology or computer science at final year honours level or above. Our participants were frequent users of digital documents. As information structuring has only been closely studied in skilled information workers, we believed that participants with casual information seeking skills and needs would be less realistic. The participants also had no prior exposure to spatial hypertext systems, which permitted us to capture the initial expectations of how they could benefit from an information structuring tool.

Participants were initially screened in a pre-study questionnaire to capture their information seeking and structuring skills. Then, they were introduced to Garnet in a ten-minute tutorial, followed by an open-ended period of self-directed exploration. The main study was then undertaken, with the participant and their activity on the computer being recorded on videotape. At the conclusion of the main study, a post-experimental interview and questionnaire captured the participants' impressions, views and experiences. Where users were asked to express an opinion, scoring was on a seven-point Likert scale.

Each participant was given the same task for the main study – a simple information-seeking task (to find papers that would be good source material for a literature review on digital libraries). They were given a brief description of digital libraries and a list of related topics to assist the selection of their initial queries. After completing the initial digital library topic task, a further requirement for documents upon human-computer interaction as a theme in digital libraries was introduced, and participants asked to obtain specific information on that. They first used the “scatter” tool described above to support this task, before embarking upon an independent search for this material. For this task, participants used a digital library collection of over ten thousand computer science technical reports.

4.1 Results

Participants were asked to compare their experiences of working with Garnet with a number of familiar systems that support information structuring or information seeking. We also observed their pattern of work and their organisation of documents during the study. We will first report the effectiveness of Garnet as a DL interface, before moving on to the patterns that we observed in the participants' use of Garnet and concluding with an examination of the participants' response to the information structuring support of Garnet.

Accessing Digital Library Features

In [4] we reported some potential problems when digital library functions were provided within a spatial hypertext interface. Anticipated problems included difficulties such as metaphor dissonance and the effectiveness of the presentation of suggestions from the ‘scatter’ facility. Our first goal was to identify the actual degree of problems encountered in real use.

We started by evaluating particular features of a digital library. Participants reported that basic digital library tasks such as searching and reading documents were comparable in ease-of-use with the same features in a web-based digital library. No participant reported, or was observed, experiencing problems with these features. This strongly suggests that the spatial hypertext interface of Garnet does not impede access to digital library features.

Participants were also asked whether they had problems distinguishing parts of the system that they could manipulate – e.g. documents in their own workspace – with parts where they could not – e.g. in browsing structures of the library. Given the known problems of different modes of operation in human-computer interaction, we were concerned that this could prove a major problem. However, all participants denied having a problem with this. There are some contributory factors that may have influenced this. Firstly, all parts of the workspace which included a view upon a digi-

tal library component – e.g. a search result set or a browsing node – were very regular in appearance, containing a column of documents or other items, and had a different colour background. Compared with the more freeform organisation preferred by our participants, the contrasting regularity of system items in the workspace provided a subtle distinction to the users’ own creation. The distinction between system-owned and user-owned items may also have been generally assisted by the fact that many operations could be achieved on both system- and user-owned objects of the same type, minimising the scope for unexpected behaviour.

Participants were also asked to rate the particular visual representation of documents, search lists and other items individually. All items were rated positively: however some useful and interesting ideas were suggested, as follows:

Firstly, six of the ten participants independently expressed their wish to be able to alter the title of documents. We had not allowed for this, as it is at odds with the nature of a digital library where documents are not normally editable. It is, however, very much within the nature of spatial hypertexts. Explanations included opaque titles of documents, and that titles often did not fit the immediate task of the user. This suggests that even for users who, like our subjects, have not been exposed to spatial hypertexts before, some spatial hypertext features that disrupt digital library expectations may be an important contribution of integration. Such a feature may, on the other hand, raise issues concerning copyright and authorship.

Secondly, five participants requested a more visual access to digital library features that were obtained from outside the workspace – e.g. the launch of new queries. Here, the preference could be explained both from the persistent appearance of such elements in web interfaces to digital libraries and the visual interactive style of spatial hypertext. In addition, two existing visual DL interfaces – DLITE [5] and SketchTrieve [10] – have used such a representation. For example, in DLITE each service of a library is represented by its own object in the workspace. Whether a ‘toolbar’ item – directly suggested by two users – or an object in the workspace – not actually suggested – is the appropriate design remains to be seen.

We also elicited the user response to our “scatter” facility, which matches the documents in a search result- or browsing- set against the workspace organised by the user. Seven participants rated the matches it found as “useful” or “very useful”. Two of the dissenting three participants had used a large, miscellaneous list for most of their documents. Given the text-matching approach we used to implement “scatter”, such heterogeneous groups would not result in any matches being found [21]. Our participants also approved of the search history provided by Garnet – corroborating the expectations of the designers of both SketchTrieve [10] and NaviQue [9].

To summarise, our participants found no difficulties using Garnet to access DL facilities. They were able to distinguish between system- and user- owned areas with apparent ease, quickly recognising the different behaviours of each. Our users identified areas for improvement, such as being more consistently focussed on the workspace presentation of tools, and permitting more editing of items than our digital library origins led us to anticipate. Our novel features, that exploit the user’s own organisation of documents, were positively received, and we were able to corroborate the claims made by the designers of other visual DL systems.

Patterns of Behaviour

We were also interested in how users followed their information seeking and information structuring tasks throughout the study. This was captured through both the video recording and post-experimental interview.

A first point of interest is that subjects closely interleaved information seeking and information structuring. Once a user decided to keep a document, even provisionally, it was immediately moved onto the spatial hypertext workspace. Organisation of the document was performed at the same time. This simple pattern was observed in every participant.

A document on a new subject or of uncertain role would often be placed in a particular group in the workspace before being reorganised to another position later in the user's work. This behaviour mirrors the patterns of work previously observed in physical environments [11]. However, two participants (8 and 10) focused on a single miscellaneous column, minimising their organisation work within the task. In interview, one reported that they would organise their documents more precisely at the end of their detailed reading, and before doing any final searching. The other participant stated that they would probably not organise documents within a task, though they would organise documents between separate tasks.

When the remaining eight participants identified a theme in two or more documents, this would result in the creation of new group. However, the consequences did not stop there. In half of all cases, the creation of a new group would result in the user doing a new query to the digital library to attempt to obtain similar documents to add to that group.

Given these behaviours, the organisational activity of information structuring was certainly interlinked in a manner that resonates with previous information seeking and spatial hypertext research [7, 12, 15].

Information Structuring and Spatial Hypertext

In the pre- and post-experimental questionnaires, we screened the participants for their use of information structuring features in existing digital libraries and their rating of these features in comparison to those in Garnet. Though six participants regularly used digital libraries (monthly or weekly), none used any information structuring tools provided by them. For example, the ACM Digital Library (used by all six of these subjects) provides a means of organising documents into 'binders'. Only two participants had attempted to use this feature, and neither found it useful.

Participants were also probed as to their use of bookmarks (or favourites) in their web browser. Nine participants used the bookmark facility, but only three organised their bookmarks each month or more frequently. Nine participants rated Garnet as being superior to the bookmark facility – the exception being the one who did not use bookmarks. Our participants reported that the purely visual interface of Garnet was better suited to organising work than a browser's bookmarks – the latter often being invisible, and organisation is done separately to adding a new bookmark. E.g. User 2 said: "it is nice to have it all in one area". More comments on using Garnet's workspace appear later in this section. The only advantage of bookmarks that was reported was the advantage of being able to change the title used – see Section 4.1.1.

We asked the participants to compare the informal structuring tools that they used in the experiment with the formal, explicit organisation that they could perform in other environments – e.g. the folders in a filing system. Garnet contains support for

both explicit and implicit structures, and we wished to elicit the perceived advantages of implicit structures. Seven users rated implicit structuring as being superior to explicit structures, and three rated it equally. Implicit structuring was noted as being particularly beneficial in the middle of searching for documents, and explicit structuring superior for long-term storage towards the end of a searching cycle.

Participants embraced the ability to organise documents on their workspaces. When asked what benefits they perceived in this, answers included: User 4, “I can see a document on the desktop without having to go back”; User 7, “being able to store stuff and organise them is good...this way you can have stuff that relates between a couple of areas”. Seven participants specifically mentioned the advantages of having an overview of what they collected, and eight reported storing documents as being an important benefit over traditional Web-based DL interfaces.

Participants also commented positively on the tangible, drag-and-drop interaction of the interface: e.g. User 9, “I really like the ability to manipulate here and move them around and take them off”; User 6, “You just drop stuff where you want it”.

Other advantages reported included supporting deciding which search to do next, remembering which searches had already been done and prioritising documents in perceived order of importance. All these are activities previously reported in physical environments, and claimed as potential advantages of spatial hypertext.

A DL/Spatial hypertext system similar to Garnet, is evaluated in [19]. However, the experiment uses a fixed set of documents and subjects do no queries or searches. In contrast, our study reveals the connection between seeking and structuring.

Discussion

Our study clearly suggests that spatial hypertext’s information structuring facilities are supportive of traditional information seeking in a digital library. Participants’ observed patterns of workflow matched the interleaved patterns observed in [7, 11, 12, 13, 16] and they reported the same patterns in post-experimental interviews.

Our participants also demonstrated known patterns in spatial hypertext, despite none having used any similar system before (the closest analogy was that two had used ‘MindMap’ software). This corroborates existing hypertext research and suggests that our subjects demonstrated typical rather than exceptional behaviour.

The visual, gestural interaction of spatial hypertext was particularly noted as an advantage by the participants, and suggested changes such as editing document titles and presentation of search facilities on the workspace are consistent with both spatial hypertexts and other visual DL interfaces like DLITE [5].

5 Conclusion

In evaluating Garnet, we have discovered that information structuring occurs in electronic as well as in physical environments. Given the evidence from observations such as Ellis and Kuhlthau, it is clear that users of digital libraries, as information seekers, benefit from information structuring during their searches. Support for information structuring in digital libraries is currently poor, and even for systems with traditional web-based interfaces there is a strong case for providing good information structuring support. However, the fluid organisation of documents seen in information structuring is, we feel, more readily supported by a drag-and-drop visual interface than a dialog-

centred web one. The response of our participants to web-based tools for information structuring that they had used was notably negative.

Previous systems have used existing classifications or automatically generated ones [6, 20] to organise the results of searches. From our user study, we have obtained evidence that the topical structures implicitly created during information structuring may be a further kind of classification that can be used for this purpose.

Information structuring is, however, at odds with some expectations of libraries. For instance, the ability to re-title documents was regularly requested by our subjects. Support for information structuring in the digital library is clearly worthy of much more research.

Acknowledgements

This work was supported by: Middlesex University, London; University College, London; and University of Waikato, New Zealand.

References

1. Bainbridge, D., Buchanan, G., McPherson, J.R., Jones, S., Mahoui, A., Witten, I. H.: "Greenstone: A Platform for Distributed Digital Library Applications". Proceedings of the European Conference on Digital Libraries, Springer-Verlag, pp. 137-148, 2001.
2. Bederson, B. and Hollan, J. "Pad++: A zooming graphical interface for exploring alternate interface physics." In Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '94, Marina del Rey, CA,) ACM Press, New York, 1994, pp. 17-26.
3. Buchanan, G., Blandford, A., Jones, M., Thimbleby, H.: "Spatial Hypertext as a Reader Tool in Digital Libraries", K. Börner, C. Chen (Eds.), Lecture Notes in Computer Science, Volume 2539, Springer-Verlag, ISSN: 0302-9743, pp. 13 – 24, 2002
4. Buchanan, G., Jones, M., Marsden, G.: "Exploring Small Screen Digital Library Access with the Greenstone Digital Library". Proceedings of the European Conference on Digital Libraries, Springer-Verlag: pp. 583-596, 2002.
5. Cousins, S.B., Paepcke, A., Winograd, T., Bier, E.A., Pier, K.A.: "The Digital Library Integrated Task Environment (DLITE)". Procs. ACM DL Conference, pp.142-151, 1997.
6. Cutting D., Karger D., Pedersen J., Tukey, J. W.: "Scatter/Gather: A Cluster-based Approach to Browsing Large Document Collections", Proceedings of the 15th Annual International ACM/SIGIR Conference, Copenhagen., pp. 318-329, 1992.
7. Ellis, D., "Modelling the information seeking patterns of engineers and research scientists in an industrial environment", Journal of Documentation 53(4):pp.384-403, 1997.
8. Fuhr, N., Klas, C.P., Shaefer, A., Mutschke,; "Daffodil: An Integrated Desktop for Supporting High-Level Search Activities in Federated Digital Libraries", Proceedings of European Conference on Digital Libraries 2002, Springer, pp. 597-612, 2002.
9. Furnas, G. W., and Rauch, S.: "Considerations for information environments and the NaviQue workspace". Proceedings of the Third ACM Conference on Digital Libraries (DL '98), ACM Press, New York, pp. 79-88, 1998.
10. Hendry, D. G. and Harper, D. J.: "An informal information-seeking environment". Journal of the American Society for Information Science, 48(11), pp. 1036-1048, 1997.
11. Kidd, A.: "The Marks are on the Knowledge Worker", Proceedings of the ACM CHI Conference, Boston, MA, pp. 186-191, 1994.
12. Kuhlthau, C. C.: "Seeking Meaning: a process approach to library and information services" Ablex Publishing, Norwood, New Jersey, 1992.

13. Malone, T. W.: "How do People Organise their Desks? Implications for the Design of Office Information Systems", *ACM Transactions on Information Systems*, v. 1 (1), pp. 99-112, January 1983.
14. Marshall, C., Shipman, F. and Coombs, J.: "VIKI: Spatial Hypertext supporting emergent structure". *Procs. of the ACM European Conference on Hypermedia Technology* ACM Press, pp. 13-23, 1994.
15. Marshall, C. and Shipman, F.: "Spatial Hypertext and the practice of information triage". *Proceedings of the Eighth ACM Conference on Hypertext*, ACM Press, pp. 124-133, 1997.
16. O'Day, V and Jeffries, R.: "Orienteering in an Information Landscape: How Information Seekers Get From Here to There", *Proceedings of INTERCHI*, ACM, pp. 438-445, 1993.
17. Shipman, F., Marshall, C., and Moran T.: "Finding and Using Implicit Structure in Human-Organized Spatial Layouts of Information", *Proceedings of Human Factors in Computing Systems (CHI '95)*, pp. 346-353, 1995.
18. Shipman, F.: "Seven Directions for Spatial Hypertext Research", *First International Workshop on Spatial Hypertext*, *ACM Hypertext Conference 2001*, Aarhus, Denmark, 2001. Online at: <http://www.csdl.tamu.edu/~shipman/SpatialHypertext/SH1/shipman.pdf>
19. Shipman, F., Hsieh, H., Moore, M.J., Zacchi, A. "Supporting Personal Collections across Digital Libraries in Spatial Hypertext", *Procs. JCDL 2004*, in press.
20. Witten, I., McNab, R., Boddie, S., Bainbridge, D.: "Greenstone: A Comprehensive Open-Source Digital Library Software System". *Proceedings of the Fifth ACM Conference on Digital Libraries*, ACM Press, pp.113-121, June 2000.
21. Zamir, O., Etzioni, O., Mandani, O. and Karp, R. M.: "Fast and Intuitive Clustering of Web Documents". *Third International Conference on Knowledge Discovery and Data Mining*, AAAI Press, Menlo Park, California. pp. 287-290, 1997.