Spatial Hypertext as a Reader Tool in Digital Libraries

George Buchanan Ann Blandford
Middlesex University
Bounds Green Road
London, N11 2NQ, UK
+44 20 8411 5939, +44 20 8411 6163
{g.buchanan, a.blandford}@mdx.ac.uk

Matt Jones
Waikato University
Hamilton
New Zealand
+64 7 358 5115
always@acm.org

Harold Thimbleby
Middlesex University
Bounds Green Road
London, N11 2NQ, UK
+44 20 8411 6061
harold@mdx.ac.uk

ABSTRACT

In this paper, we introduce Garnet, a novel visual interface for Digital Libraries. Garnet aims to provide a visual workspace in which the user can structure and organize documents of interest. This structure is then used to organize and filter further documents which may be of interest, such as search results.

Keywords

Digital Libraries, Spatial Hypertext, Visual Interfaces

1. INTRODUCTION

There is typically unequal support for different parts of the information seeking process [10]. For example, there is obvious support for the task of retrieving documents from the Digital Library (DL), whereas there is little or no support for the tasks of organizing and collating material discovered while searching and browsing into user-generated structures. (Note that we use the term "document" to refer to a discrete item of information recorded in any medium: a journal paper; video excerpt etc.)

Our system, Garnet, is intended to assist users in these latter phases of document collation and organization. It shares features in common with VIKI [7, 8]. Marshall and Shipman created the VIKI system to facilitate this sort of work, introducing the concept of a Spatial Hypertext. A Spatial Hypertext is a visual workspace where each document is represented by a shape, and cues such as position and color are used to indicate document relationships. In VIKI, the freeform use of document shapes is complemented by a more formal organization into a hierarchical set of document collections.

Marshall and Shipman intended, and expected, a number of benefits in the use of VIKI. One was that by using the hypertext, information workers could articulate their findings, expectations and conclusions to others. Secondly, it provided the opportunity for users to clarify their own thinking through the process of organizing, selecting and rating documents. Their studies into the use of VIKI within information work had positive outcomes, so it appeared that a similar approach was appropriate in supporting DL users in document organization.

The First ACM+IEEE Joint Conference on Digital Libraries '01, June 24-28, 2001, Roanoke, Virginia.

VIKI has itself not been connected directly to an information discovery system such as a digital library, so its own support for information seeking as a whole has a complementary role to existing digital library systems. This, therefore, suggested that an integrated system consisting of both "parts" would bring benefits not found in two separate systems.

In addition, for many information workers information seeking is a long-term task. If documents are organized in a workspace over a period of time, they become an ever more complete and precise reflection of the information work the worker is engaged in. Given the time cost of the work which is represented in the workspace, we wondered whether it could be used to amplify the effectiveness of users at a later date through, say, improved precision, recall, or organization of new documents. This is similar to the economic concept of the positive externality, where one person's actions indirectly benefit another, but here a user's past activity benefits their present work.

Therefore, in Garnet we are introducing features which exploit the implicit knowledge discovered in the contents of a workspace combined with the advantages of a direct connection between the workspace and digital libraries.

2. SAMPLE SCENARIO

A pilot version of Garnet has been created, which is connected to the New Zealand Digital Library Project's Greenstone software [11]. We will now demonstrate the system in use, starting from a "bare" hypertext. The library material we will use is the Humanity Development Library of the United Nations, one of the widely available examples of a Greenstone library collection, which consists of several thousand pages.

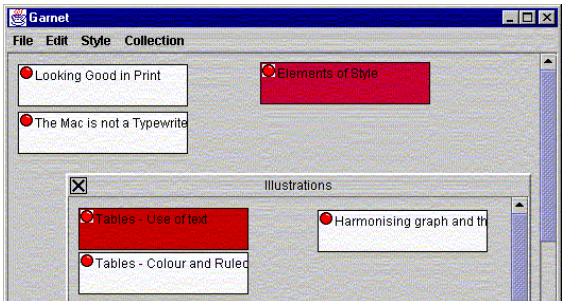

Figure 1 – A Garnet Client in Use

2.1 Overview

In Figure 1, we see a "typical" Garnet user session in progress; a number of "windows" appear inside the main browser window. Each of these is a collection of materials which the user has recorded in the current, or a previous, session. Each

document is represented by a shape containing some text, as indicated in the diagram, which we term a 'label' for simplicity.

Within a collection, the user is free to place, size and color 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 a digital library's search facilities.

2.2 Example search

Let us now follow a simple sequence of interactions, starting with an example search. For our purposes, we are going to investigate snail farming, in an attempt to discover whether we have the appropriate resources to consider that form of agriculture. With Garnet loaded, we start a new search in the Greenstone system (one of three sources Garnet supports), and we enter the simple query "snail". As shown in figure 2, a simple collection window appears with a number of document labels appearing one beneath the other, similar to a typical web-based result list.

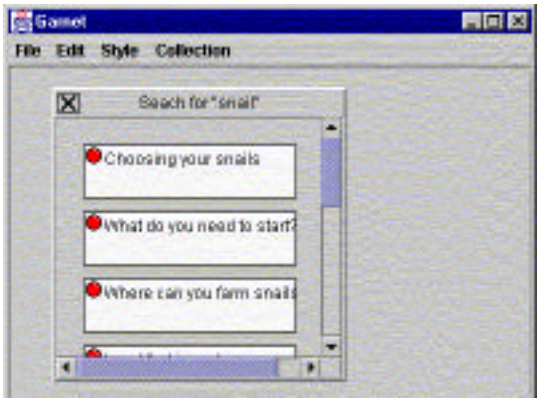

Figure 2 – A simple search

On reading the first two documents (achieved by a simple double-click on the appropriate documents), we decide that we'd like to keep the second document ("What do you need to start"), and we move it to our main workspace window. The first document, however, seems a bit advanced, and we can delete it from the list. As a result of this, the later documents move upwards. Should we wish to return to the search results at a later date, by default these changes would be retained.

2.3 Demonstration of "scatter results"

In the previous search, we performed a plain search. Garnet, however, can exploit the organization done by the user in a novel manner. We can "scatter" a set of documents (including search results) over the existing layout of documents in the workspace. "Scattering" places the search documents near to groups of existing documents with which they have a strong similarity.

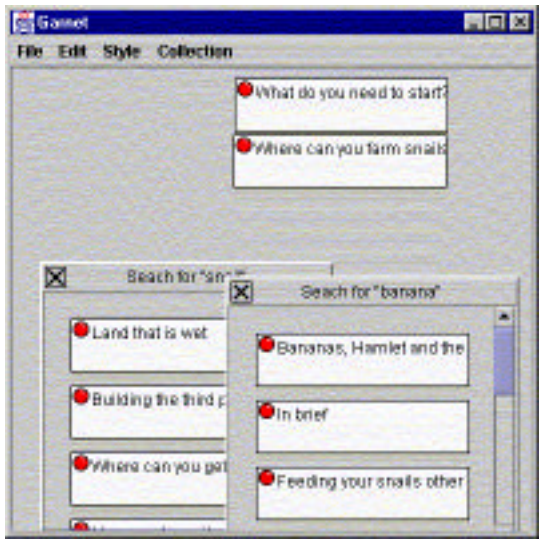

Figure 3 - before scatter

Continuing our previous example, we have now selected a few more useful-looking documents, but a couple of questions remain unanswered. Suppose we have a plentiful supply of bananas which we'd like to use, but we're not sure would be appropriate. We do a naïve search, "banana" and the results don't look too helpful (Figure 3); however, note the third item – for this example we've chosen something that is visible (for clarity) – if we now do a "scatter", (Figure 4), and only a subset of the search results appear on the main collection. This small subset, which appears in a light gray above, has been found by Garnet to be a close match to the existing pair of documents, which appear in white. We can now investigate the two suggested documents which are similar to the previously selected pair. and confirm that ripe bananas can be used to feed snails.

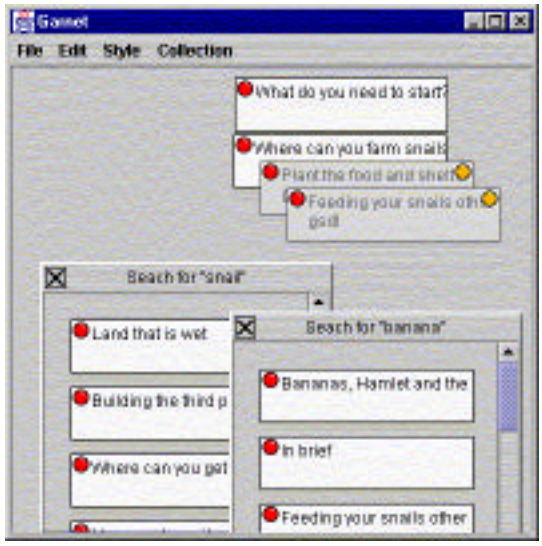

Figure 4 - after scatter

3. PREVIOUS WORK

Before describing the particular properties of Garnet, we will review a small sample of related systems and how Garnet differs in comparison to them.

3.1 Visual Interfaces for Digital Libraries

Visual interfaces specifically designed for Digital Libraries already exist. Three such interfaces which have similar objectives to Garnet are DLITE [2], NaviQue[4] and SketchTrieve[5]. All three systems are intended to give coherent access to a number of information services (e.g. content search, author indexes etc) and sources (different collections from a number of libraries), and represent separate searches as discrete objects in a 2-dimensional workspace.

To compare these systems with Garnet, we will consider how much control the user has over the appearance of objects in their workspace, the relative significance of different objects in the workspace, and in what manner they can use the objects in their workspace to perform further work.

3.1.1 DLITE

DLITE [2] is primarily intended to integrate heterogeneous services and sources in a visual environment and provide a persistent 2-dimensional visual workspace for a reader in a digital library.

Compared to Garnet, the focus visually and interactively is primarily upon search result sets, with individual documents having a very simple, small, iconic representation whereas sets appear as large, complex visual objects.

The appearance of collections of documents and individual documents is system defined and very uniform (color being the one exception). The placing of document icons within document sets is not controllable by the user – collections are only created by the system. As with Garnet, document sets can be created independently of search actions, but unlike Garnet there is no capability to build a hierarchy of sets.

The workspace behaves as a simple repository for artifacts of past actions, and to provide an interface for current activity; previous activity can be recalled and continued through the corresponding object(s) in the workspace.

3.1.2 SketchTrieve

SketchTrieve [5] is a tool for coordinating complex search activity. Like DLITE, one goal of its design was to connect a number of services to a visual workspace in a visual workspace, though in the case of SketchTrieve, search services only are supported.

Like VIKI and Garnet, SketchTrieve the representation of individual documents is substantial and, to a degree, user-customizable. Unlike Garnet, VIKI, DLITE or NaviQue, document sets are not represented as separate objects, indeed document sets do not exist in any explicit form at all except as search result lists.

In common with DLITE, SketchTrieve's workspace is an artifact repository; no advanced facilities exist to leverage these artifacts with current work.

3.1.3 NaviQue

NaviQue [4] has a number of superficial similarities to DLITE. In addition, however, it is influenced by the work of Bedersen et al on the Pad++ zoom-able spatial hypertext.

The workspace can itself directly hold documents, which are displayed in their normal rendered manner (though the use of zooming may make a document invisible or virtually iconic – one of which would normally be the case).

Search result sets can be displayed in a number of presentations, all of which are subject to being scaled and zoomed in the same manner as the document representations. Surprisingly, user-created document sets have no formal existence, though groups of documents can be identified by selecting an area of the workspace..

The use of document sets (search- or user-selected) is very flexible (e.g. they can themselves be searched), though there is no formal hierarchy (though sets may posses a logical relationship through the tiered use of search).

Like Garnet, NaviQue has methods for identifying similarities between documents and document groups, however this is exploited to assist navigation rather than retrieval.

3.1.4 Summary

NaviQue and DLITE often indicate individual documents in a very minimalist manner, whereas VIKI, Garnet and SketchTrieve would use more substantial representations. Interestingly, the visual customization of objects is also generally less varied in both NaviQue and DLITE. The expressive power available over individual documents, and the focus upon concrete documents are known to be key facets of VIKI, and therefore we have chosen to retain these.

Moving onto document sets, NaviQue, offers a more flexible system than DLITE or SketchTrieve, However, VIKI's combination of formal and informal groups offers the "best practice" of all three, so again we have also adopted this.

Finally considering the exploitation of old work in the performance of current tasks, NaviQue also has some methods for identifying document similarities, which can be used to relate new material to older material, but this is used for navigation on a document-by-document basis only. No other exploits of the workspace content are reported in the literature.

3.2 VIKI

VIKI is clearly a system which begs comparison with Garnet. Their basic appearance and interactions are very similar. However, unlike VIKI, Garnet actively exploits the workspace layout to identify semantic themes in the user's work artifacts, and similarities between visually unconnected documents or sets.

This identification of themes is done in two parts; spatial analysis and lexical analysis. The technology for identifying document groups by their visual properties is found in both Garnet and VIKI (though implemented differently). In VIKI the primary use of this information is in supporting visual interaction and direct manipulation in the workspace; no related lexical analysis is performed.

Garnet's lexical analysis is used to identify semantic similarities between documents. The "scattering" facility described above is built upon Garnet's ability to provide both spatial and lexical analyses. Many other capabilities can be generated from this combination, such as mapping between different workspaces, or assisting in the consolidation of isolated but closely related groups.

Furthermore, VIKI has only had an elementary and indirect connection to information systems, producing simple result

sets. Garnet, on the other hand, has direct access through the Greenstone digital library software to the entire gamut of facilities of the digital library (in addition to search results, for example, browsing structures, document profiles and metadata can be accessed).

4. SUPPORTING TECHNOLOGIES

At the heart of Garnet is a spatial parser [6,10], which identifies visual patterns within the arrangement of shapes (i.e. based on shape/position similarity rather than textual similarity). For each identified visual pattern, the contents of the corresponding documents are used to generate lexical classifiers to represent the semantic theme of the pattern.

One problem which could occur is the creation of classifiers based on groups of documents which are highly heterogeneous (e.g. those yet to be sorted and reviewed by the user). These sorts of sets were often observed by Marshall and Shipman, and are due to the provisional nature of the organization and collation task. Our classifiers are based on clustering techniques [3] which include the identification of incoherent clusters, and we use this common clustering function to avoid building active classifiers for heterogeneous groups. Clustering techniques can also permit Garnet to assist in the organization of disorganized groups by providing a "first-cut" at partitioning them.

When the workspace is changed, the spatial parser must re-generate the corresponding classifiers. At present, this evaluation is done eagerly, but this could cause Garnet's response to slow down when the workspace is changing rapidly; therefore, we are moving towards lazy evaluation. A related problem occurs in identifying documents arbitrarily positioned by the user in the ongoing flow of their work, and how best to indicate uncertainty of the user, rather than any uncertainty of the classifiers.

5. USER INTERFACE

Having successfully implemented a working version of Garnet, a number of design questions have emerged regarding the user interface. Firstly, there is the question of user versus system control of certain features; secondly, of how to represent information to the user; and thirdly, there are potential problems of metaphor dissonance, as discussed below.

On the question of control, with the scatter feature, we chose to place the control of this feature fully into the hands of the user. If scattering were applied automatically, document labels representing the suggested position of the documents could appear and disappear continuously, removing the user's sense of control and continuity. A related question was also which search results should be scattered at one time. Again, placing this within the user's control should provide a more user-centered approach.

One representational problem is the placing and appearance of scattered documents in their suggested places. We currently use a very simple technique (using a shaded label and placed at the bottom left of the group), but many other possibilities exist. One unresolved issue is how to demonstrate the strength of the correspondence; the current display gives no indication of this. A further difficulty is how to inform the user of the positioning of scattered documents which are in areas of the workspace which are not currently visible.

A second representational problem is how to provide browsing of the digital library. The integration of the simple structures found in search facilities is reasonably straightforward, but browsing structures are often much more complex.

Some common desktop/GUI program metaphors do not seem to hold well, creating metaphor dissonance; for example, above we deleted a document from the search results at one point. We don't normally destroy documents in a library, so this metaphor is inappropriate. Furthermore, should we bar the document from appearing in the current result set only, or from appearing in any later result sets (i.e. "blacklist" it)?

6. CONCLUSIONS

At present we have a simple pilot system working; our initial system has shown the potential for making sensible suggestions within the context of our trial Digital Library sources. A number of the issues faced so far have been discussed above. Some parts of the system are as yet incomplete, but we expect soon to perform a formative evaluation of the system through a user study, and will then develop Garnet further to discover the issues which arise when using Garnet within a co-operative, multi-user context.

7. ACKNOWLEDGMENTS

Our thanks to the University of Waikato's New Zealand Digital Library Project, especially David Bainbridge, Stefan Boddie and Ian Witten. This work was funded by Middlesex University.

8. REFERENCES

[1] 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, Nov.) ACM Press, New York, NY 17-26.

[2] Cousins, S. B., Paepcke, A., Winograd, T., Bier, E., and Pier, K. 1997. The digital library integrated task environment (DLITE). In Proceedings of ACM Digital Libraries '97 Conference. ACM Press, New York, NY, 142-151.

[3] Cutting, D. R., Karger D.R., Pederson, J.O. and Tukey, J.W. Scatter/Gather: a cluster-based approach to browsing large document collections. In proceedings of ACM/SIGIR 1992, 318-329.

[4] Furnas, G., and Rauch, S., Considerations for information environments and the NaviQue workspace. In Proceedings of the Third ACM Conference on Digital Libraries (DL '98. Pittsburgh,PA, June) ACM Press, New York, 79-88

[5] Hendry, D. G., and Harper, D. J., An informal information-seeking environment. Journal of the American Society for Information Science, 48(11):1036-1048, 1997.

[6] Lakin, F., Visual Grammars For Visual Languages, ProceedingsAAAI-87, Seattle, Washington, July 12-17, 1987, pp 683-688

[7] Marshall, C., Shipman, F. and Coombs, J., Spatial Hypertext supporting emergent structure. Proceedings of the 1994 ACM European Conference on Hypermedia Technology (ECHT '94, Edinburgh, UK) ACM Press, NY, 13-23.

[8] Marshall, C. and Shipman, F., Spatial Hypertext and the practice of information triage. In Proceedings of the Eighth ACM Conference on Hypertext (Hypertext '97, Southampton, UK) ACM Press, New York, NY 124-133.

[9] Marchionini, G. Information Seeking in Electronic Environments. Cambridge Series in Human Computer Interaction, Cambridge University Press, Cambridge, UK, 1995

[10] F. Shipman, C. Marshall, and T. Moran "Finding and Using Implicit Structure in Human-Organized Spatial Layouts of Information", Proceedings of Human Factors in Computing Systems (CHI '95), 1995, pp. 346-353.

[11] 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, June 2000, ACM Press, 113-121.