Improving Navigation Interaction in Digital Documents

George Buchanan
Future Interaction Laboratory
Swansea University
Swansea, UK
g.r.buchanan@swansea.ac.uk

Thomas Owen
Future Interaction Laboratory
Swansea University
Swansea, UK
cstomo@swansea.ac.uk

ABSTRACT

This paper investigates novel interactions for supporting within–document navigation. We focus on one specific interaction: the following of figure references. Through this interaction we illuminate factors also found in other forms of navigation. Three alternative interactions for supporting figure navigation are described and evaluated through a user study. Experimentation proves the advantages of our interaction design, and the degree to which the interaction of existing reader software can be improved.

Categories and Subject Descriptors

H.3.7 [**Information storage and retrieval**]: Digital Libraries—*User issues*

General Terms

Human Factors

Keywords

Document Readers, Computer-Human Interaction

1. INTRODUCTION

Within–document navigation is a common action performed by users when reading texts in many circumstances: initial, quick readings, the deep analysis of a selected document, and checking remembered details.

Cathy Marshall [6] has demonstrated that digital document reader software provides significant impediments to users' interaction with documents. Our work rises to the challenge, raised by Marshall, of providing "library materials that not only capture the affordances of paper, but also transcend paper's limitations". This paper examines navigation in digital documents from a narrow focus, studying the following of figure references within a text. While our underlying interest is in arbitrary navigation, such unpredictable activity makes comparisons between systems problematic. In contrast, the predetermined start and end points of figure navigation make it highly suitable for laboratory evaluation. Figure navigation also shares key aspects of interaction with arbitrary navigation. For example, a user may flip between start– and end–points. When figures are numbered in sequence, with no direct relationship to the numbering of text sections, the section structure fails to assist navigation. Consequently, users must scan the text to locate the item – a behavior often seen during arbitrary navigation.

Building on our own research this paper commences with the applicaton of "click-through" navigation [1] to figure navigation, contrasting it with both hyperlink– and page–navigation, as used in many document reader applications. The three navigation techniques of hyperlink, click–through and page navigation are then evaluated in a laboratory–based experiment focusing on figure navigation. We then discuss click–through interaction in the light of existing designs. The paper concludes with a summary of our contribution and directions for future research.

JCDL'08, June 16–20, 2008, Pittsburgh, Pennsylvania, USA.

2. CLICK-THROUGH NAVIGATION

To improve user interaction with digital documents, new interfaces are required. In [1] we introduced a novel interaction technique "Click Through Navigation". This simple method uses snapshots of a particular part of a document to act both as a method of extracting information, and rapidly returning to it. This paper studies the adaptation of click–through navigation for following figure references in a document. This focussed interaction will serve to demonstrate the wider opportunity to improve navigation techniques in digital documents.

Figure 1 shows the implementation of click–through navigation for figure references. References to figures are highlighted with a surrounding rectangle. This snapshot shows a thumbnail of the figure, and a click on this navigates to the figure itself. We will now discuss three different interactions, including the thumbnail one seen here.

2.1 Hyperlink Navigation

One clear candidate for the interaction to be activated by clicking on a figure reference is to navigate the document view to the destination figure and to place it in the centre of the user's field of view. This is similar to the navigation options available within the Adobe Acrobat reader. In the default Acrobat navigation a click on a link activates a particular *view* of the document. In many cases the scale of the document display will change, and a particular area on a particular page is placed on display on the screen.

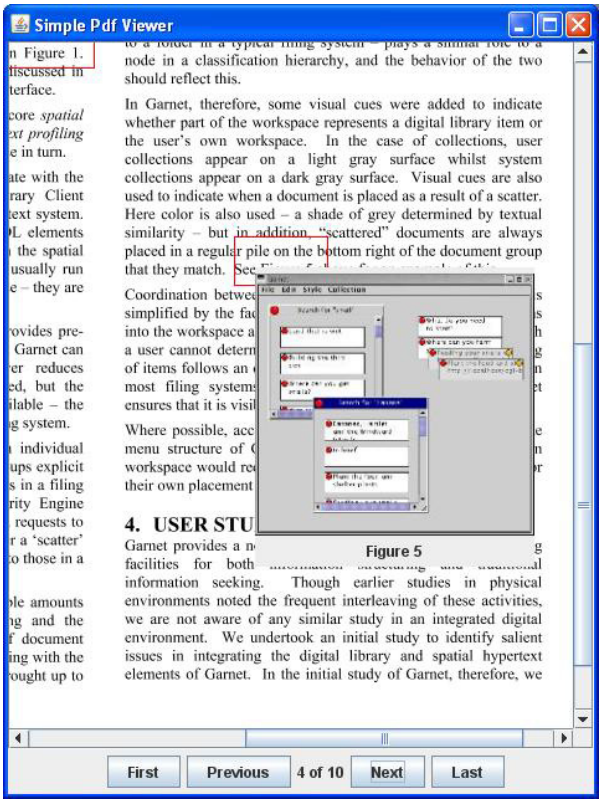

Figure 1: Click–through navigation for document references: Thumbnail of referenced figure appears right–centre.

Hyperlink navigation suffers many problems in current document reader software (e.g. Acrobat). One difficulty is that using a link is not an undo-able action. This means that while the user can readily use a hyperlink to move *to* a location, returning *from* a location requires the user to recall their previous location and navigate explicitly to it. Also, there is no support for automatic reference linking: authors must add explicit markup, and many fail to supply this.

2.2 Thumbnail Navigation

Thumbnails have been used to assist within–document navigation [2] and to provide document previews for small screen displays. Click–through figure navigation differs from other navigation thumbnails in portraying part of a page, and in providing navigation to specific content at a particular location on the destination page rather than to a whole page.

Thumbnails present a reduced view of the target image, giving a first impression of the item being referenced. Where the user has already seen the figure, this may support recognition, and for new images, a general impression can be given. A thumbnail therefore initially retains the reader's current document view. If the reader wishes to proceed to the figure, they simply click the thumbnail. Our thumbnail navigation, seen in Figure 1, supports two-directional navigation. When viewing the figure, a floating palette permits the user to return to the reference that led the user there.

Thumbnails have foreseeable disadvantages: larger and more complex figures suffer a greater loss of information and detail, and confusion may arise when thumbnails of different figures appear to be the same.

2.3 Traditional Page Navigation

In addition to hyperlink and thumbnail navigation, we used a baseline standard page-based navigation. This interaction provided scrolling, text search and page-based navigation controls such as forward, backward and go to page.

All the interactions discussed in this section were built using the JPedal Java–based PDF decoder software. After basic implementation, incremental improvments were made with feedback from pilot studies.

3. EVALUATING FIGURE NAVIGATION

In this section, we will discuss the design and execution of the experiment we conducted to compare the three navigation methods introduced above.

3.1 Experimental Design

As we wanted a detailed comparison between interfaces, we adopted a laboratory experiment on closed tasks. The tasks required the user to follow references at a range of distances between reference and the corresponding figure. Each task related to a specific document, to eliminate learning effects between tasks. For each navigational distance a set of 6 tasks was created, to ensure that the interaction was evaluated, rather than effects from individual documents.

We used a within–groups design for the experiment, as recruitment was a potential problem. Each participant would use the three interfaces in turn. The order of interface use was balanced across participants to neutralise the scope for ordering and learning effects. The pool of tasks was also assigned to interfaces using squared–design method, to decouple tasks from particular interfaces in a controlled manner.

The structure of each session started with an induction questionnaire followed in turn by training and familiarisation, experiment and semi–structured post–completion interview. Information on age, use of document reader software, computing experience etc. were gathered during the induction interview. The participant's assigned ordering of interfaces was followed, with a training and familiarisation period preceding the period of experimentation with each interface. Each tutorial session started a brief demonstration of the navigational method. The participant was then given two tasks to complete using the navigation method, followed by a period for open–ended exploration.

Twelve participants were recruited from the graduate school of a computer science department. The task documents were from the ACM Digital Library and represented a wide range of CS topics. Whilst the general format of the documents would be familiar, this is unlikely to help predict the location of a figure relative to its referencing text. The familiar format should assist general navigation.

Screen capture software and observational notes recorded the participants' behavior. Timings were taken from the beginning of each task until the point when the participant verbally confirmed that they had found the target figure. This time may occur some time after the figure is first seen by the reader: e.g. when the caption or other identifying element has become visible and has been read.

The post–completion interview gathered subjective data from the participants. The stuctured part obtained ratings for each interface from selected criteria from the IBM us-

ability satisfaction questionnaire [5] including effectiveness, ease of use and an overall utility score,.

3.2 Hypotheses

We had two main hypotheses. First, that for references to images that appeared in close proximity to the referenced figure, the traditional scrolling interaction would be most effective in terms of both time and user satisfaction. The thumbnail presentation would be disadvantaged by visually interfering with the document display and obscuring the content to which it referred, and the hyperlink presentation would suffer from the low visibility of its effect. Second, longer movements (several pages) the thumbnail and hyperlink navigation would be much more effective than scrolling.

There were also a number of specific questions that we wished to examine from the experimental process. First, to determine the impact of navigational distance on the relative performances of the interaction (and thus the "break even" points between our first and second hypotheses). Second, to determine which of the novel navigations was superior to the other. Thirdly, to identify what were the interactive and cognitive effects upon the user of each interaction.

3.3 Quantitative Findings

The quantitative data from the study were evaluated using either Student's t-test or ANOVA with two factors. In Figure 2 we see the overall time taken by participants for six tasks in each of the three interfaces. It is clear that the traditional scrolling navigation takes a longer time to complete the set of tasks than the hyperlink and thumbnail methods (average time/tasks: scrolling=7.41s, thumbnail=1.46s, hyperlink=1.94s). There is a strong statistical significance in the gap between scrolling navigation and either novel technique ($p<0.01$): significance is unsurprising even in a small sample when the gap in performance is large.

Our hypothesis that scrolling would prove superior across short distances was clearly undermined in the case of thumbnail navigation, which was consistently if not significantly, faster ($p=0.12$). Link navigation was potentially ($p=0.22$) slower than scrolling for movement of one page or less.

Similar differences are also seen in the subjective feedback of the participants shown in Table 1. Again $p<0.01$ for comparison between either novel interaction and the paging/scrolling mode. The traditional method scored poorly with all users, and was consistently ranked third. The time performance of the thumbnail– and hypertext– methods is comparable, though there is a slight advantage to the thumbnail method, at $p=0.12$.

The time performance of the three interactions across different navigation distances unveils further detail. Unsurprisingly, scrolling navigation becomes increasingly ineffective as the navigation distance increases (Fig. 2). In our study, the maximum distance between reference and target was two pages. Even this short distance demonstrated that the

Interface	Practicality	Effectiveness	Ease of Use
Scrolling	2.75 (1.46)	2.50 (1.69)	2.81 (1.31)
Thumbnail	5.75 (2.76)	6.63 (2.62)	7.00 (2.56)
Hyperlink	8.50 (1.20)	9.00 (1.07)	8.63 (1.03)

Table 1: Subjective participant ratings for three interfaces; scores out of 10, S.D. in parentheses

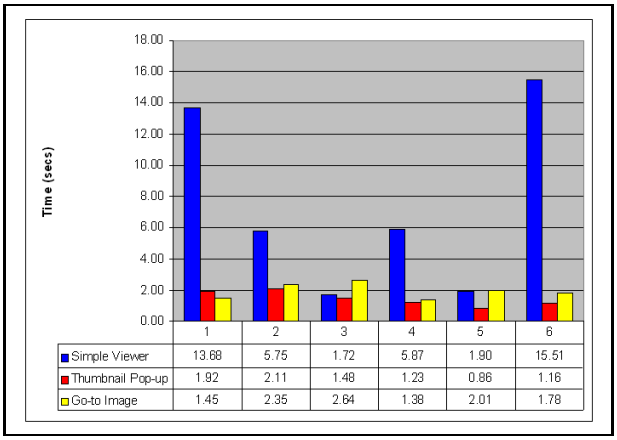

Figure 2: Mean task completion times in three interfaces; two each of within–page, neighbouring page and two page navigations

navigation time rises steeply as the distance increases. Conversely, the typical performance of the link and thumbnail navigation systems remains consistently under two seconds.

Despite the consistent if small performance advantage of the thumbnail method, the subjective judgments of participants favored the hyperlink navigation. Whilst for most participants the difference between their ratings was small, for two there was a substantial dislike for the thumbnail technique. There was no evidence from other quantitative data to explain this preference, but some answers will be seen as we turn to the participants' subjective responses.

3.4 Qualitative Feedback

In the post–completion interview, one goal was to identify both differences in demands on the user between the navigational modes, and the tasks that each method was best suited to. Six participants complained that the thumbnail interaction could obscure the display of the main document, particularly when the navigation target was on the same page. Nonetheless 7 participants gave a score of 7 out of 10 or higher for the design. In contrast, the hyperlink navigation received complaints of a feeling of disorientation: on arriving at a new page, users felt uncertain where to look, or what to look for, particularly when the diagram was unfamiliar (e.g. the figure appeared after the reference).

Specific participant comments explain much their reactions. A comparison of the thumbnail and hyperlink navigations is particularly revealing. As we noted above, despite the thumbnail view being faster, its subjective scores were significantly poorer than those for hyperlink navigation. There were a number of detail issues with the implementation of the thumbnail navigation that were raised by participants.

First, the thumbnail always appeared to the bottom right of the current mouse position, and as a result it was sometimes only partially visible. This problem was repeatedly mentioned by participants, and is a clear error in the implementation. Whilst in standard click–through navigation this position allows the user to easily "hide" a thumbnail, it is clearly unhelpful in this case.

Second, thumbnails could suffer poor legibility, particularly for larger figures. Higher legibility over selected parts of

the figure content can be achieved through "semantic zooming"[7] or "content aware cropping"[8] techniques.

Problems were not limited to thumbnail navigation. Relocating one's position in the document after a movement was a particularly significant problem in the hyperlink mode. Both the new systems received better feedback than the traditional scrolling method of navigation.

4. DISCUSSION

The experiment confirmed the superiority of the two novel interactions, particularly for movements of one page or more. An analysis of within–document figure references from 100 ACM papers determined that 47% of references spanned longer distances that favor the improved interaction.

The preview provided in the thumbnail navigation had a mixture of positive and negative comments. There are number of simple steps that can be taken to remedy its worst deficiencies (e.g. adjusting the thumbnail position). One surprise was that hyperlink navigation performed poorly when the navigation distance was short. Interview and observation data suggests that the problem was that when the distance travelled was very small, the visual feedback of the limited movement was not sufficient to communicate that the navigation had worked.

The issue of the navigation distance at which our novel methods would demonstrate an advantage was partly answered. While navigation within the same page or to the nearest edge of the neighbouring edge could favor scrolling navigation – if by only a small margin – even the modest distance of the further edge of a neighbouring page gave the new methods a reliable advantage.

Common document reader applications (e.g. Acrobat) lack effective tools for bi–directional movement. Instead, users must deploy adept navigaton skills to enable them to move rapidly between two points. XLibris has shown this need not be so [3], though unfortunately that solution was not evaluated systematically. We further improve on the bi–directional support provided in XLibris. First, we automatically extract references by a combination of textual and spatial parsing [4], without explicit markup. Second, by providing a thumbnail of the part of the document referred to, the user can view a figure and fluidly continue their reading from the reference. Our spatial parsing technique is also used to ensure that the caption and other key details are immediately visible when a reference is followed.

Though we cannot draw direct lessons about arbitary navigation, we can draw some conclusions regarding click–through navigation. When reading printed texts, humans often use the insertion of fingers or bookmarks to support their rapid return to a chosen location. Click–through navigation implements bi–directional hyperlinks to reconstitute this simple and powerful interaction.

Cathy Marshall has called for new interactions that transcend the limitations of paper [6]. Our rapid navigation uses means not available on paper. While the gap between paper and screen may not be eliminated, such advantages can mitigate the shortcomings of digital texts.

5. FUTURE WORK

Ultimately, this current paper demonstrates that it is possible to create interactively lightweight approaches to within–document navigation, that also satisfy users in practice. Click–through navigation reproduces useful attributes of paper, such as facilitating quick movement between different pages, but it also transcends paper by support of direct navigation to figures.

There is much research to do. While the thumbnail navigation method performed best, feedback uncovered shortcomings in our implementation. We can expect to improve further the performance and user experience of click–through navigation, narrowing the gap in user satisfaction and increasing its advantage in time efficiency when compared to both scrolling and hyperlink navigations. The presentation of the thumbnails themselves needs further examination.

6. CONCLUSION

Document triage is a critical point in any user's information seeking, where users compare documents in detail and navigate rapidly between different points. Inspired by Marshall's challenge to provide effective document navigation tools, we have demonstrated that it is possible to provide significant improvements in digital navigation.

Thumbnail previews proved to interfere with reading and interaction in some circumstances, so refinements are clearly necessary. Our click–through navigation proved an effective competitor to hyperlink navigation in a controlled study, and significantly superior to page–based navigation techniques such as scrolling. Click–through navigation has significant advantages over hyperlink navigation, particularly in regards to arbitrary and bi–directional movement.

7. REFERENCES

[1] G. Buchanan. Rapid document navigation for information triage support. In *Proc. ACM/IEEE Joint Conference on Digital Libraries*, page 503. ACM Press, June 2007.

[2] A. Cockburn, C. Gutwin, and J. Alexander. Faster document navigation with space-filling thumbnails. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 1–10, New York, NY, USA, 2006. ACM.

[3] G. Golovchinsky. Going back in hypertext. In *HYPERTEXT '02: Proces. of the 13th ACM conference on Hypertext and hypermedia*, pages 82–83. ACM, 2002.

[4] F. Lakin. Visual grammars for visual languages. In *Proceedings AAAI-87*, pages 683–688, July 1987.

[5] J. R. Lewis. Ibm computer usability satisfaction questionnaires: Psychometric evaluation and instructions for use. *International Journal of Human-Computer Interaction*, 7:57–âĂŞ78, 1995.

[6] C. C. Marshall and S. Bly. Turning the page on navigation. In *JCDL '05: Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, pages 225–234, New York, NY, USA, 2005. ACM.

[7] K. Perlin and D. Fox. Pad: an alternative approach to the computer interface. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 57–64, New York, NY, USA, 1993. ACM.

[8] B. Suh, H. Ling, B. B. Bederson, and D. W. Jacobs. Automatic thumbnail cropping and its effectiveness. In *UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 95–104, New York, NY, USA, 2003. ACM.