

Shooting the information rapids

Ann Blandford & Hanna Stelmaszewska

Interaction Design Centre
School of Computing Science, Middlesex University,
Bounds Green Road, London, N11 2NQ
a.blandford@mdx.ac.uk

SUMMARY

Terms such as 'navigation' and 'information orienteering' have been applied to users working in large information spaces such as the Web or digital libraries. Such terms – and their descriptions – imply that the user is in control of the interaction, moving deliberately through the information space. In practice, as recognised in the work on situated cognition, users often behave much more reactively than this, responding to external stimuli in a fluid way. In this paper we report on user behaviour when interacting with a collection of digital libraries, focusing particularly on situations where users were switching between multiple windows.

KEYWORDS : Digital Libraries, navigation, usability.

INTRODUCTION

Within the literature on how users work with large collections of information, there is a dominant view of users as navigating purposefully from one state to the next. This general description is applied whether the users have only a single goal or multiple goals that emerge through the interaction. O'Day & Jeffries [2] discuss this in terms of 'information orienteering', using the analogy of the runner who navigates across unknown terrain with some kind of navigation aids (such as map and compass) to reach a goal and then adopts the next goal and repeats the process for that goal until the final goal – the end of the course – is reached. Similarly, Pirolli [3] presents a computational model of what he terms 'information scent-following', within a framework that accommodates both cognitive and ecological perspectives, considering user behaviour as largely rational and intentional.

Clearly, such accounts are adequate descriptions for many interactions, but not all. Here we report on some more extreme findings, which we refer to as 'shooting the information rapids', taking the analogy from Suchman's early work on situated cognition [4], in which she uses the example of a canoeist shooting the rapids to discuss the way that people's behaviour is often more a consequence of the environmental conditions than of their personal goals, though both play their part.

The work reported here aimed to study users in as natural

a setting as possible. Rather than studying the design of a particular environment, we were interested in how users worked with a variety of digital libraries of their own choosing. The details of the study are presented elsewhere [1]. In summary: five individuals – four research students and one established academic – were given access to six widely available digital libraries, and invited to perform tasks of their own choosing with the libraries, and to think aloud while working. A few days later, they were interviewed, to clarify any points that were unclear from the transcripts. Their activities were videotaped, and subsequently transcribed and analysed. The main finding reported in [1] were the identification of particular hurdles to be overcome by novice users learning to work with new libraries: library designers make strong assumptions about the expertise of their users that fail to recognise the difficulties of the initial learning phase. The study has since been extended to a further two users, a research assistant and an administrator with a particular interest on which she wished to acquire more information. Here, we focus on one aspect of the results for four of these seven users, namely the behaviour of the interactive system when there were multiple browser windows opened simultaneously. The other three users did not get into a situation where they were working with multiple windows.

These four users found themselves – intentionally or otherwise – in a situation where there were multiple windows open on the display. In the case of users B and E, the opening of a second browser window was intentional, a way of dealing with slow network response times. In the case of users C and G, the second window was generated automatically by the system when the user followed particular links.

The interactions

For each of the users discussed here, we provide summaries of key extracts from their interactions (Figures 1 – 4). In the figures, standard text is as spoken by the user; text in square brackets with a '>' indicates a user action, while text in brackets with a '<' indicates a system response; and italicised text is a summary of an extended period of interaction. The black, grey and white bars denote window in use, window in background and

no window, respectively. ‘...’ indicates extended interaction omitted, and numbers are used to refer to phases of interaction in the accompanying text.

First, we consider the two users who unintentionally had a second window open: users C and G. At 34 minutes in to the interaction, user C received a search result from the Emerald library that consisted of just one item, a book. He had reviewed a description of the book; at the bottom of the review page were two links, offering the option to purchase the book from two different internet booksellers. He clicked on the link to order the book from Amazon (Figure 1, phase 2), which opened a new browser window. He scrolled down then up the page, then returned to the calling page (by selecting the other window) without comment. Having reviewed a description of the book again, he returned to the Amazon window just to close it, although the book was available via this bookseller. He then repeated these steps with the second bookseller, although this bookseller displayed a message immediately that the book was not available. He commented that “I want to order the book but if I want to order the book I need to enter another web page.” He

I'm trying the book. [> scrolls down the page] [> clicks on 'order the book (Amazon)']	1	
[< 'amazon 'page displayed – new window'] [> scrolls down then up the page]	2	
[> activates 'full review article'] [> scrolls up then down the page]	3	
[> activates 'amazon'; closes window]	4	
[> activates 'review article'] [> scrolls down the page] [> clicks on 'Order the book' from 'The Internet Bookshop']	5	
[<'The Internet Bookshop' page displayed: 'Selected title unavailable'] I want to order the book but if I want to order the book I need to enter another web page. [> scrolls down and up the page]	6	
[> activates 'review articles'] [> scrolls up the page] [> copies the book title]	7	
[> activates 'The Internet Bookshop'] [> clicks on 'Basket' button] [< 'The Internet Bookshop: Shopping basket empty' page displayed] [> clicks on 'search' button] This shop is very difficult to use. I'm giving up. [> closes the window]	8	
[< 'review article' page displayed] I gave up to order this book [> clicks on browser 'back' button]	9	

Figure 1: summary of C's interaction, times 34:32 - 36:35.

made a little effort to order the book, copying the details from the library page to paste somewhere on the bookshop page, but then gave up when he could find nowhere to paste the description, commenting that “This shop is very difficult to use”. This behaviour would be difficult, if not impossible, to give an account of using established models of rational behaviour: the user's behaviour appears to be driven by a half-formed commitment to do something (buy the book) which is actually achievable, and which he has seen is achievable (at step 2 of the interaction shown), but he does not follow through and perform the actions that would be rational if the commitment were a fully formed goal. The behaviour is more exploratory: the user is following links that present themselves as possibilities on the screen, without a clear purpose.

In contrast to user C, who was fully aware of the fact that new windows were automatically generated each time he followed a 'purchase the book' link, and switched between windows and closed them at will, user G was confused and misled by the sudden appearance of new windows when she was not expecting them. This happened towards the end of her interaction (after 1 hour and 1 minute); she was surveying the list of conference proceedings within the ACM digital library, and followed a link – presumably to get more information on conferences. This link generated a new page, but she was, and remained, unaware of this, as shown in phase 2 of figure 2, where she tried to press 'back' to undo the

[< 'CHI' page displayed] [> clicks on 'CHI100' link]	1	
[<ACM Events & Conferences page, April] I don't want the list of conferences [> clicks browser 'back' button] [<ACM Events & Conferences page, Jan.] [> tries to click browser 'back' button] [> clicks browser 'forward' button] [<ACM Events & Conferences page, April] [> clicks browser 'back' button] [<ACM Events & Conferences page, Jan.] [> clicks ACM 'home']	2	
[< ACM home page displayed] [> clicks 'digital library' link] That's better [works for some time in this third window] [> closes ACM DL window]	3	...
[< ACM DL 'Events & Calendar' window] [> closes window]	4	
[< AMC DL Conference on page] [works briefly in this window] I've had enough	5	...

Figure 2: summary of G's interaction, times 60:51 - 70:04.

effect of following the link, which took her to a page she had never seen before; further pressing of ‘back had no effect, so she retraced her steps ‘forward’, then back again. To try and escape from this unfamiliar situation, she selected the link to ‘home’, which opened yet another new window – this time to the ACM home page, from which there was a link to the digital library. At last, she knew where she was again: “That’s better”, and could continue working for a few more minutes. It was only when she tried closing windows (phases 4 and 5 of figure 2) that the user uncovered the earlier windows, and by this point she appeared to have given up on trying to make sense of the system behaviour. The system was both unpredictable (it was not possible for the user to determine when following a link caused a new window to be created and when it just updated the page within the current window) and also lacked appropriate feedback to the user (she was not aware that a new window had been created). In this case, the user’s behaviour was rational and intentional in the traditional sense of those terms, but the design of the system forced the user to start behaving reactively in a situation where she had apparently performed a non-reversible action. Her goal of finding relevant articles was necessarily subsumed by the device-oriented goal of returning the system to a familiar state.

In contrast to users C and G, users B and E intentionally opened additional browser windows, both of them to try and overcome poor response time problems when downloading files from remote servers. User B started working with two browser windows about 20 minutes into his interaction, when a particularly large document was taking a very long time to download (it eventually took about 5 minutes – see figure 3, phase 7). In the following 33 minutes, user B switched windows no fewer than 30 times, an average of just over a minute in each window session. At some points (e.g. phases 3 and 5), the window switch was simply to monitor the current state of downloading in the hidden window – a very self-directed (intentional) activity. He was interleaving two main goals – of finding articles and maintaining awareness of the state of the other window. At other points (e.g. phases 16 – 20) the switch was more provoked by the system, as he had by this point become habituated to switching window every time a download was not instant. Perhaps surprisingly, user B did not appear to lose track of the state of the interaction at any point. For example, he was aware enough of what was going on in each window to use ‘back’ appropriately, indicating an adequate memory for earlier states of each separate interaction; however, he clearly found his chosen strategy demanding, as he commented (phase 29) that “It’s difficult to run between two engines because each one I did have different search requests”. Inevitably, he relied heavily on the display as he

[> minimises the ‘download’ window] ...let’s try another one...	1	
[> clicks on browser ‘new window’ option]		
[<home page displayed; <i>selects link to EBSCO</i>]	2	...
[> activates ‘download window’]	3	
[<i>works with EBSCO then initiates a search</i>]	4	...
...I will finish it up the ...	5	
[<EBSCO search results displayed]	6	
It took about 5 minutes to download... ...so I can continue working with ACM...	7	...
[< ACM search results page displayed]		
[>activates EBSCO window; <i>initiates a search</i>]	8	...
[<i>works in ACM then initiates a search</i>] let’s go back and find...	...	9
[EBSCO] there are no articles on ‘text mining’ [<i>initiates a new search</i>]	10	...
[<i>works in ACM; starts to download article</i>] ...let’s go back to EBSCO...	...	11
[> activates EBSCO window; <i>short interaction</i>] [> clicks on ‘back to article search’]	12	...
[> activates ACM DL window; page empty] ...still downloading.	13	
[> activates EBSCO] ...takes a lot of time. I will try another search engine...	14	...
[<i>selects CSTR; CSTR page is loading</i>]		
[<i>works in ACM</i>] ...let’s try another engine...	...	15
[> selects ‘ingenta’; ‘ingenta’ page loading] [<i>initiates a search in CSTR</i>]	16	...
[<i>explores Ingenta; page loading</i>]	...	17
[<i>works briefly in CSTR; page loading</i>]	18	...
[<i>works briefly in Ingenta; page loading</i>]	...	19
[<i>reviews results in CSTR; initiates new search</i>]	20	...
[<i>works briefly in Ingenta; page loading</i>] ...it takes a lot of time	...	21
[<i>reviews results in CSTR; selects IDEAL library</i>] ...it’s taking a long time to download...	22	...
[<i>works briefly in Ingenta; page loading</i>]	...	23
[activates ‘IDEAL’; <i>initiates a search</i>]	24	...
[<i>explores Ingenta; search result displayed</i>]	...	25
[<i>initiates a search in IDEAL</i>]	26	...
[<i>works in Medline then initiates a search</i>]	...	27
[> activates IDEAL window; <i>works briefly</i>] ...I think I’ll download this one...	28	...
[<activates ‘ingentaJournals’ search results] It’s difficult to run between two engines because each one I did have different search requests [<i>working in Ingenta; starts to download doc.</i>]	29	...
[> activates ‘download article’ window] OK, that’s one I want to save.	30	

Figure 3: summary of B’s interaction, times 19:36 - 52:24.

[- blank web browser window] We're going to work much more efficiently if we have a new navigator window... [> clicks browser file -> new navigator window]	1	
[< New window appears with home page] ... and have that one running in the background. [works for some time in the second window] Of course my other window might have done something by now. [continues working in same window until downloading another paper] Of course now we have time to go back to our other window which by now should have...	2	...
[> clicks on 1st window] ... finished. She says hopefully. [> scrolls to bottom of page; blank page appears] Nope. Great.	3	
[switches back to do more work in second window then switches again]	4	...
OK. So I seem to have the rest of it now. [... file sent to printer] Also, now worth going back to my first window. Hmm. Screen's probably...	5	...
[> clicks] Yes. It's all tied up for the minute, so I'll go and collect some things from the printer. [goes to printer; returns; selects first window]	6	...
[< 1st window showing] OK. So I don't know how I will know if I've got everything I've printed out. Because I'm losing track of it completely. But I think that's been sent and not needed anymore.	7	
[> clicks 2nd window] [proceeds to view and then print out paper]	8	

Figure 4: summary of E's interaction. 34:00 – 41:00.

switched windows, and this interaction more than any other leaves the viewer with the impression of someone 'shooting the rapids': reacting to a slow download by switching windows, then having to rapidly take stock, restore context and resume purposeful activity. Although the interaction is not paced by the system – or if it is, it is the slow pace that places demands on the user, rather than fast pace – the rapid switching places heavy demands on the reacting user.

Whereas user B did not appear to lose track of the state at any point, the same cannot be said for user E, who had an additional demand on her attention: that of tracking which articles had been sent to a remote printer and which had not yet. It was 34 minutes into her interaction before user E started working with a second window – an action again provoked by slow downloading of a file. Within the short interaction shown in figure 4, there is

evidence of losing track twice: in phase 5, where she referred to the second window as the first (not a serious problem, but an indication of losing track), and later in phase 7, where she commented that "I'm losing track of it completely". Unlike user B, she soon reverted to working almost entirely with a single window. Like user C, it is possible to detect a low commitment at one point – in phase 2, when she surmises about the state of the other window, but does nothing to check it.

A rational style of analysis, as conducted by Pirolli [3], gives a good account of much of the users' behaviour. For example, user G is using her knowledge of the state and possible operations to try and 'pick up the scent' that she lost in moving into phase 2 of the interaction (figure 1). In this case, it is possible to identify design requirements such as avoiding spawning new windows, or giving better feedback on the state to the user. However, the other interactions discussed here have a complexity that is difficult to discuss simply in terms of one or two main goals and the user selecting actions to achieve those goals. We have seen partial commitments that influenced behaviour but did not result in goal achievement, and also a cost-benefit assessment by users who were frustrated by the pace of the interaction, and chose to modify their strategies to place greater cognitive demands on themselves, and hence also more reliance on the system to support them in maintaining their awareness of the system state. An account of behaviour based at least partly on a 'situated' view appears to capture some of the important properties of these interactions. Most importantly, users frequently have multiple goals, engage in satisficing behaviour, and depend on the system for maintaining their knowledge of the state and of interaction possibilities.

ACKNOWLEDGEMENTS

We are grateful to the seven subjects who gave their time to participate in this study. This work is funded by EPSRC Grant GR/M81748.

BIBLIOGRAPHY

1. Blandford, A., Stelmaszewska, H. & Bryan-Kinns, N. Use of multiple digital libraries: a case study. To appear in *Proc. JCDL 2001*. In press.
2. O'Day, V. L., & Jeffries, R. (1993). Orienteering in an Information Landscape: How Information Seekers Get From Here to There. *Proceedings of InterCHI '93*, pp. 438-445.
3. Pirolli, P. Computational Models of Information Scent-Following in a Very Large Browsable Text Collection. In *Proceedings of CHI'97* ACM Press, New York, 1997, pp. 3-10.
4. Suchman, L. (1987). *Plans and situated actions*, Cambridge: Cambridge University Press.