Navigating, Discovering and Exploring the Web: Strategies used by People with Print Disabilities on Interactive Websites

Christopher Power[1], Helen Petrie[1], David Swallow[1], Emma Murphy[2], Bláithín Gallagher[2], Carlos A. Velasco[3]

[1]Human Computer Interaction Research Group, Department of Computer Science, University of York, York, YO10 5GH, United Kingdom
[2]National Council for the Blind of Ireland, Whitworth Road, Drumcondra, Dublin 9, Ireland
[3]Fraunhofer Institute for Applied Information Technology (FIT), Schloss Birlinghoven, D53757 Sankt Augustin, Germany

{christopher.power, helen.petrie, david.swallow}@york.ac.uk
{emma.murphy,blaithin.gallagher}@ncbi.ie
carlos.velasco@fit.fraunhofer.de

Abstract. The majority of research into web accessibility has focused on identifying and eliminating the problems that people with disabilities encounter when interacting with the Web. In this paper we argue that we need to move away from studying user problems to studying how people with disabilities apply interaction strategies while browsing the Web. In this paper we present a study of 19 print disabled users, including blind, partially sighted and dyslexic people, interacting with a variety of interactive Web 2.0 web applications. The participants undertook tasks using concurrent and retrospective protocols to elicit information about how they interact with web content. The result of this study was a collection of 586 strategic action sequences that were classified into seven different types of strategy. Differences in the application of strategies between the user groups are presented, as well as the most frequent strategies used by each user group. We close the paper by discussing some implications for the design of websites and assistive technologies as well as the future directions for empirical research in accessibility.

Keywords: Web accessibility, user study, user strategy, print disabled Web users, blind Web users, partially sighted Web users, dyslexic Web users.

1 Introduction

The Web provides many opportunities for users to take part in work and leisure activities online. The evolution of Web 2.0, in which websites allow users to participate in creating, contribute and share content [9], makes our connections to both other individuals and organizations stronger in the virtual world. It is important that all users be able to participate equally in these activities, including people with disabilities who use different assistive technologies and interact with the Web in very different ways from mainstream users. However, recent studies show that people with disabilities still encounter large numbers of problems on the Web [1,7,12,14]. While existing web accessibility good practice covers some of those problems encountered by disabled users, there are still a substantial number of reported problems that have no clear solution. What are the design principles that can help close that gap and better support users with disabilities?

When trying to address problems encountered by mainstream users on the Web, researchers and practitioners fall back on empirical research and design principles about user interactions with the Web and other interactive systems. However, when trying to address problems encountered by disabled users, relevant empirical studies about user interactions are rare. As a result, we are unable to go back to "first principles" when trying to either remove or, at least, reduce the impact that problems have on users with disabilities. We need more empirical work eliciting and understanding interaction strategies that disabled users apply when interacting with the Web. With that information, we can construct a framework for analysing why disabled users encounter the problems on websites that are currently beyond our understanding.

In support of this goal, we present a study with blind, partially sighted and dyslexic users, users who are collectively part of the group referred to as print-disabled users [6]. In the study, we elicit the strategies that users undertake while interacting with interactive Web 2.0 applications.

2 Related Literature

There is an increasingly large body of evidence demonstrating that the Web has not been very accessible to people with disabilities throughout its history. In 2001, Coyne and Nielsen [5] conducted a study with 20 blind, 20 partially sighted and 20 mainstream users. Each user undertook one task on each of four different websites, including one search engine. Blind users were only able to complete 12.5% of tasks; while partially sighted users fared only slightly better, with a task completion rate of 21.4%. Given that mainstream users completed 78.2% of tasks, these results show the challenge of providing an equivalent experience on the Web for disabled users. In the same paper, Coyne and Nielsen presented user-based evaluations of websites conducted with 18 blind people, 17 partially sighted people and 9 people with physical disabilities. They recorded the problems encountered by participants, and the analysis of those problems produced 75 guidelines for building accessible web applications.

The Coyne and Nielsen study was in the early days of web accessibility, and one would expect that the number of problems encountered by disabled users on the Web would decrease as developers gained awareness and knowledge of web accessibility. In 2004, the Disability Rights Commission (DRC) of Great Britain undertook a Formal Investigation into the state of Web accessibility [1]. This investigation included the largest user study of web accessibility to date. In the DRC study, users with a range of different disabilities undertook tasks on 100 websites. From 913 tasks that were undertaken on the websites, 76% of tasks were successfully completed across all disability groups. However, blind users only succeeded 54% of the time, while partially sighted users achieved a success rate of 76%. So, while web accessibility did improve between Coyne and Nielsen's study in 2001 and the DRC study in 2004, many problems still remained. The DRC report provided a comprehensive list of the types of problems that were most commonly encountered by different disabled users and demonstrated there was a large amount of overlap between the different groups. For example, even though blind and partially sighted users navigate websites in very different ways, and explore information on pages using their own distinct strategies, both groups struggled with page structure and unclear navigation mechanisms.

Lazar et al. [7] conducted a diary study with 100 blind users regarding the frustrations they encountered during their day-to-day interactions with the Web. The results echoed many of the DRC findings, with poorly designed forms, navigation structures, misleading links and page layout problems all being prominent frustrations listed by participants.

In 2007, Petrie and Kheir [12] conducted a study with six blind and six mainstream users. These users undertook seven tasks on each of two mobile phone websites. Users were asked to undertake a concurrent verbal protocol during the tasks in which they identified problems as they occurred and rated those problems for their severity. The results showed that despite having very different means of interacting with websites, the two user groups shared a number of problems.

A recent study by Power et al. had 35 blind Web users undertake two to three tasks on 16 websites, with 10 users testing each website [14]. Users encountered nearly 1400 problems, ranging from lack of feedback, through to alternatives not being provided for inaccessible content. Notably, a large number of the problems related to trying to understand the layout of the content in webpages, how navigation structures were organized and finding information within the website.

When discussing Web accessibility, the problems that disabled users encounter have been central. Certainly, the examination of those problems has raised awareness about the importance of web accessibility and provided some solutions [2,10,11]. However, we argue that the current dominant approach of encouraging developers to implement websites that simply avoid those problems is not an effective way to create websites that people with disabilities can use. While some problems, such as conflicts between assistive technologies and web browsers, may be addressed by informing developers on how to avoid them, the range and complexity of problems experienced by people with disabilities is too high to be addressed by a problem-based approach.

Given the recurring problems encountered by disabled users of exploring web content, discovering the layout of pages and understanding the navigation structures in a website, it is intractable to list of all of the problems that can occur, what their causes are, and how to address them through specific implementation techniques. We must move to an approach where we can accomplish good interaction design for websites that support the interaction strategies of disabled users.

However, information regarding how disabled users interact with the Web, as opposed to what problems they encounter, is surprisingly thin.

As early as 1995 the ACCESS project elicited requirements from 150 blind users about how to improve accessibility. Their report emphasized the need for flexibility and support in getting overviews of material on webpages, allowing blind users to jump to key landmarks and providing signposting to key paths through a website [13].

Berry [2] conducted one of the earliest in-depth interview studies about blind users' interactions with the Web. He presented a number of themes that emerged from the analysis of the interviews and he makes a key observation comparing novice and competent disabled Web users. He states that competent Web users had:

"…a strong mental image of what a webpage is and how it works, combined with a structured and confident approach to information location and retrieval."

This shows that having a good mental model of a webpage is important, and therefore we need designs that support disabled users in building those models.

In 2003, Theofanos and Redish [16] presented an observational study of 16 blind screen reader users undertaking a variety of tasks on government websites. From this study the authors developed a more concrete understanding of how people use their screen readers to understand a webpage. Among the key findings was that users skip quickly through content and listen to only the first few words. Such skipping was often sequential, but users also used lists of links and headings to skip through content in a more structured way. At that time, users were reluctant to use 'skip navigation' links, as they largely did not understand where the link would take them. Users also did not typically use functionality in the screen reader to find forms, and once forms were found, users were reluctant to move between reading mode and form mode.

A follow-up study by Theofanos and Redish [17] investigated the interactions of 10 partially sighted screen magnifier users when interacting with the Web. 5 of their participants changed from high to low magnification repeatedly, while 3 others changed font sizes in the page. Others compromised between how much they could read in a lens, versus how much of the page they could see.

Takagi et al. [15] conducted an investigation of blind people using online shopping websites. They had users undertake tasks on real online shopping sites and recorded their behaviour in regards to navigation within a page. A set of search pages for the online shop had a set of accessibility enhancements added to them including: alternative texts, properly marked up headings and 'skip navigation' links. Results from 5 participants revealed that some users used links to explore a page while others moved sequentially through content on the page. On the search pages with improved accessibility, there were no major changes in the browsing behaviour of users, with the users employing their own idiosyncratic strategies as opposed to adopting new ones based on the content. Of particular interest to this work, the authors say that users "[Stopped scanning] and then crawled around to check the content" when they found information that was relevant to their task, possibly indicating a strategy being applied. While interesting, there is no information on how often this happens, or why users were undertaking that activity. The authors conclude their paper with a statement that increasing the number of landmarks users can apply in their navigation will improve their success in navigating within webpages. Work in 2010 by Trewin et al [18] with 3 screen reader users also had indications that both established and ad hoc landmarks were important.

Wantanabe [19] examined how blind and partially sighted participants used headings on websites. He had 16 sighted users and 4 blind users undertake tasks on two different sets of recipe websites: one set had heading markup and presentation for key sections of webpages, while the other set had no structured headings. Both sets of participants had shorter task completion times when working with headings present than without. While interesting, it is not clear what actual role the headings played in the overall strategies of the users, or whether they were instructed to use the headings in the study.

Bigham et al. [3] used a proxy to record the browsing behaviours of 10 blind and 10 sighted users on their home computers. Similar to findings from previous studies, users seldom used skip navigation links, and interacted almost exclusively with links that were labelled with descriptive text. Blind users also employed probing more than their sighted counterparts, in which a webpage is visited through a link and then followed by a quick return to the previous page. What is unclear is why users were employing this probing interaction? Bigham et al. propose that users are trying to cope with inaccessible Web content; however, with the data available it is impossible to tell why users are applying those accelerators.

Borodin et al. [4] provided a high level overview of the ways that blind and partially sighted users interact with the Web with screen reader technologies. Their survey of strategies was based on "… general browsing strategies that were observed in the course of several user studies." It is difficult to determine from the information provided how often the strategies discussed were used, or how well the strategies represented the overall interactions of people with visual disabilities on the Web.

Finally, Lunn et al. [8] observed nine blind and partially sighted users browsing the Web for between 2 and 4 hours in a naturalistic setting. During this study, researchers recorded occasions that users encountered problems and identified strategies for overcoming those problems. 7 different strategies were identified for dealing with accessibility problems. Importantly, this work described in detail the strategic behaviour of users. However, the strategies identified did not include many of the common behaviours discovered in previous research. There are a number of reasons why this may have happened. Lunn et al. describe using a peripheral membership role typical to ethnographic studies to avoid influencing the observations. However, in such short sessions, with no opportunity for the researchers to integrate themselves into the setting, it is possible that there was a Hawthorne effect that influenced how participants behaved. In addition to this, the researchers collected data in a very short time while trying to observe many individuals concurrently in one classroom, and so it is unlikely that this set of strategies is complete. Finally, the approach used by Lunn et al. still involves studying the problems that users have, albeit with a different lens. In order to understand the broader contributory factors that impact on users, we must study their broader interactions on the Web in the absence of problems.

Given that there has been almost two decades of research in web accessibility, it is surprising how little empirical research there is into how disabled users interact with the Web. In comparison to the hundreds of papers that have been published on how and why mainstream users interact with the web, the participant numbers on which we form the foundations of design for web accessibility is very small. Further, the most comprehensive work on how disabled users apply strategies to navigate, discover and explore the web is now 10 years old [16,17]. While useful, there has been no recent work that demonstrates that what we knew at the very beginning of the web still holds true now. Given the evolution of assistive technologies and web application design in that ten-year period, we would expect that the strategies of users would have changed and become more sophisticated; however, we have no evidence to support that supposition. Finally, and most notably, we have no framework for describing the strategies that disabled users use when working with the Web. Almost all of the empirical studies examine user interactions only at the level of sequences of key presses and other low-level actions in the interface. We must understand how these action sequences relate to what the user is trying to accomplish in the interface. By linking the action sequences to higher-level user strategies, we will begin to have a foundation on which we can solve the continuing problems disabled users are encountering on the Web

In this paper we present a study of print disabled users, including blind, partially sighted and dyslexic users, regarding the strategies they apply when interacting with web applications. We seek to answer the following questions:

1. What are the high-level strategies used by people with print disabilities when interacting with the Web?
2. What are the most frequent strategies used by each of blind, partially sighted and dyslexic user groups?
3. What are the similarities and differences in how these different user groups apply strategies?

3 Method

3.1 Design

Participants from 3 print disability groups (blind, partially sighted and dyslexic) were asked to undertake a number of tasks on Web 2.0 applications. Participants were asked to conduct a concurrent verbal protocol while undertaking each task, concentrating on the strategies they were using in doing the task. After each task they were asked to provide a retrospective verbal protocol, providing any further information about their strategies. Information about the strategies was extracted from transcripts of the protocols for detailed analysis.

3.2 Participants

19 people took part in the study, 10 men and 9 women. Ages ranged from 18 to 60 years. Participants had between 7 and 40 years experience of using computers, and each of them used a computer on a daily basis. 5 participants were blind: 3 were congenitally blind, and the other two lost their sight more than 20 years ago. All 5 blind participants used the JAWS screen reader. All participants considered their familiarity with JAWS to be good (4) or very good (5) on a 5-point Likert scale. The

participants' experience of using computers ranged from 7 to 40 years, with a mean experience of 23 years. Each of the participants used their computer very frequently and indicated their familiarity with computers on a 5-point Likert scale from very poor (1) to very good (5) to be above the midpoint (3).

7 participants were partially sighted: 5 since birth and 2 since early childhood. The partially sighted participants used various screen reader/magnification technologies: 1 participant used JAWS; 2 used SuperNova; 2 used ZoomText (specifically for the magnification and colour contrast functionality); and 2 used the zoom settings of their browser or operating system. The participants' experience of using computers ranged from 10 to 30 years, with a mean experience of 20 years. Each of the participants used their computer very frequently and indicated their familiarity with computers to be good (4) or very good (5) on a 5-point Likert scale.

7 participants were dyslexic. Their dyslexia had been diagnosed between 1 and 6 years previously. None of the participants used any assistive technologies for using the Web. The dyslexic participants' experience with computers ranged from 9 to 14 years, with a mean experience of 11 years. Each of the participants used their computer very frequently. All participants rated their familiarity with computers to be at or above the mid-point on a 5-point Likert scale.

3.3 Equipment and Software

The study was run on a number of personal computers, each with a screen size of 17 inches. Each computer ran Windows XP and had the Morae® screen recording software installed. Each computer had the following assistive technology installed for users: JAWS (v10) and WindowEyes (v7) and Supernova (Access Suite v12.08).

3.4 Websites and tasks

In order to elicit a wide variety of strategies from the participants, we surveyed web applications that covered the breadth of Web 2.0. We surveyed 2-5 highly interactive, participatory web applications from the following domains: social networking (e.g. Twitter, Facebook), blogging (e.g. Blogger, Wordpress), e-commerce (e.g. online banking and shopping), video sharing and viewing (e.g. YouTube, online television), e-government (e.g. government portals, participation portals).

For each web application we collected tasks that allowed users to create, retrieve, update and delete content in the web application. For each domain, we identified the most frequently occurring tasks that were shared across the web applications contained within it. These frequently occurring tasks were then decomposed through hierarchical task analysis to provide a comprehensive list of subtasks.

We selected 6 of the surveyed websites for use in the study, with a minimum of one for each domain, each with 6-10 tasks that could be completed by users. These websites were: 2 e-banking sites (Smile Online and Egg), 1 social networking site (Facebook), 1 e-participation site (Citizenscape), 1 blogging site (Wordpress) and 1 internet television site (BBC iPlayer).

The complete list of tasks along with the number of subtasks identified in the decomposition (indicated in brackets) were as follows:

- Egg/Smile Banking (e-banking): Login (2), Check Balance (3), Transfer Money to a friend (4), Manage money in accounts (9), Manage contact details (3), Arrange an overdraft (4)
- Facebook (Social networking): Sign-up (6), Add a friend (3), Update status (3), Leave a message on a friend's wall (5), Create an event (5), Upload and tag a photo (3), Chat with a friend (3), Write a note (2), Find places nearby current location (1), Update privacy settings (4)
- Citizenscape (e-participation): Find a webcast (5), Watch and navigate a webcast (5), Post a comment on a webcast (2), Send a tweet about the webcast (3)
- BBC iPlayer (Internet Television): Find last watched programme (3), Find a previous episode (6), Watch a television programme (5), Listen to a radio programme (6), Manage favourite programmes (4), Manage parental guidance features (3)
- Wordpress (Blogging): Login and view a blog (3), Create a new blog post (6), Edit an existing blog post (6), Share a video on a blog (7), Liking and comment on a post (4), Delete a blog post (3), Adjust blog settings (3)

For each of the tasks, a realistic scenario was written that would provide the participants with the contextual information about the website and the task they were undertaking and to increase the ecological validity of the study, making each task as natural as possible. An example of one such scenario from the BBC iPlayer protocol was as follows:

Due to having a busy week last week, you missed an episode of your favourite TV programme, Top Gear. You are aware that older episodes of Top Gear are available on BBC iPlayer and you would like to see if you can find the episode you missed. Please locate the previous episode of Top Gear.

As we were trying to draw conclusions about overall types and numbers of strategies applied by each user group, and not a direct comparison of strategies applied between the individual websites, users did not undertake all tasks on all websites. Further, due to the fact that we were using websites in the wild, and not a predefined website under our control, it was impossible for us to predict exactly what users would do within a given task. As a consequence of this, the comparison of user strategies between particular websites is not a research question we have pursued in this work.

3.5 Procedure

Due to the level of control that was required and the recording software needed for data analysis, observations were undertaken in controlled environments. They took place in the HomeLab at the University of York, the offices of the National Council for the Blind of Ireland (NCBI) in Dublin, and the offices of the Foundation for Assistive Technology (FAST) in London. In order to preserve ecological validity, and avoid participants acting in ways that would be inconsistent with their home use, participants used the same brands of assistive technology and user agents that they used in their home environments. Participants were given an opportunity to set up the computer display, sound and related software to preferences that approximated their home use. Great care was taken to ensure that the experience was as close to their home computing experience as possible in a controlled laboratory study.

Participants were asked to undertake each task with a concurrent verbal protocol. They were instructed to concentrate on explaining what they were doing on the website to achieve the tasks given to them. The evaluator stopped participants when they reached the end of a sub-task. Participants were asked to retrospectively describe the strategies they were using, and in particular what they hoped to achieve by using that strategy. When the user had described the strategies supporting the previous sub-task, they were asked to continue from that point to complete their task. The videos for these two stages were analysed together for the key strategies that people use to interact with the Web. Evaluation sessions lasted a maximum of 2 hours. Within that time participants performed tasks on 3 to 8 websites each, undertaking between 4 and 19 tasks.

3.6 Data Analysis

The video recordings of each participant were reviewed to identify strategic action sequences. A *strategic action sequence* is a series of operations in the Web application interface that the user applies to achieve a goal. For each recording, we noted all utterances that could be interpreted as representing an action sequence. Each utterance was transcribed and time-coded along with an indication as to whether the utterance was made during the concurrent or retrospective verbal protocol The distinction was made to avoid the possibility of double-counting a single occurrence of a particular strategic action sequence that had been mentioned in both protocols.

The action sequences were classified, using an emergent-grounded theory approach. Each action sequence was assigned a description and label regarding its composition. For example, an action sequence where a blind user traverses a set of headings using their screen reader was labelled *Heading Traversal*, whereas a partially sighted user scanning the area above the navigation bar and then jumping to the bottom of the page to scan the footer was labelled *Top and Bottom Scanning.* Action sequences with the same composition were assigned the same label.

When an action sequence composition was identified, the participant's comments about the action sequence and the strategy they were applying when using it were recorded. This process built up a collection of action sequences related to particular strategies. These comments were also analysed to produce a set of broad interaction strategies that were applied by users in the web applications.

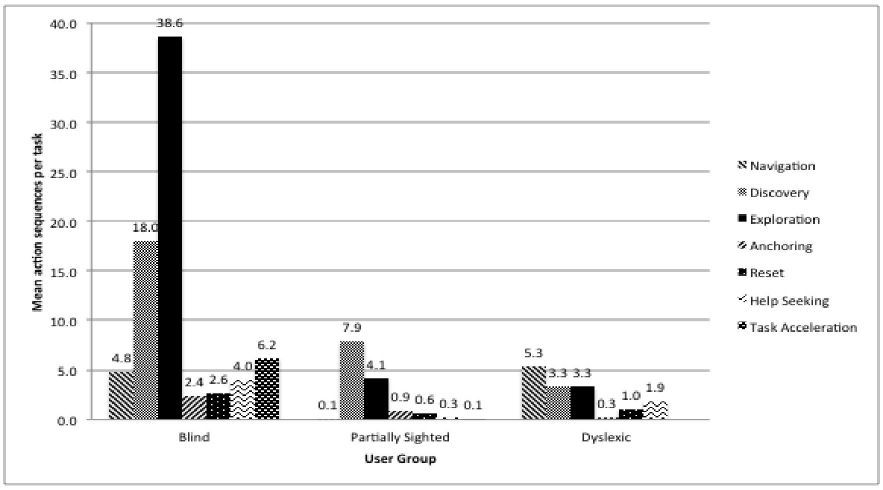

Figure 1: Mean instances of strategy types per task for each user group.

4 Results

A total of 586 instances of strategic action sequences were recorded across 89 different compositions of action sequences. Blind users accounted for 383 action sequences of 62 compositions. Partially sighted users had 98 sequences of 35 compositions. Dyslexic users had 105 action sequences of 25 compositions.

From the comments of users regarding how the action strategies were applied in the web application, 7 types of strategies were identified:

- Navigation: supports users moving from page to page within a website. This includes identifying where navigation bars and menu items lead and understanding the overall site map. Examples of action sequences for this strategy include probing navigation items and then returning to the calling page, or examining tool tips for more information about where a navigation item leads.
- Discovery: is users gaining and overview of the overall structure of webpage content. An example of action sequences for this strategy would be a user reviewing the headings of a webpage to understand what are the key sections.
- Exploration: is users extracting the meaning or context of a particular piece of content or the functionality of an interactive component. An example action sequences for this strategy would be reading a heading and its related content.
- Anchoring: is users limiting the area of the webpage in which they interact. Actions sequences for this strategy includes users avoiding particular areas of the page (e.g. right hand side) or focusing on a specific area (e.g. top of page).
- Help seeking: is users actively seeking help from online documentation or from another people.
- Reset: is where a user abandons what they are doing and starts again from a safe point. This strategy could refer to restarting the whole computer or a piece of software.

- Task Acceleration: is where users are trying to speed up tasks. This includes things like keyboard shortcuts in assistive technology and web browsers.

Figure 1 presents the mean number of instances of each strategy type per task for each user group. To investigate whether the three user groups applied the strategies at different frequencies, a 2-way mixed ANOVA was conducted on the rates of strategy types per task for the strategy types Navigation, Discovery and Exploration, (the rates of the other strategy types were too low for quantitative analysis). There was a significant main effect for user group ($F = 16.25$, $df = 2, 16$, $p <0.001$), a significant main effect for strategy type ($F = 16.90$, $df = 2, 16$, $p < 0.001$) and most interestingly, a significant interaction between the user group and the strategy type ($F = 15.20$, $df= 2,32$, $p < 0.001$).

The user group main effect showed that there was a significant difference between the overall number of strategies per task, with blind participants using more strategies per task (5.47, SD=2.36) than either partially sighted users (2.50, SD=0.82; Bonferroni post hoc: $p<0.005$) or dyslexic participants (1.53, SD=0.59; Bonferroni post hoc: $p<0.001$).

The user group by strategy type interaction showed that blind participants applied significantly more strategies for Exploration than for Navigation (Bonferroni post hoc: $p < 0.002$) and for Discovery ($p < 0.002$). Further, partially sighted people used more strategies for Discovery than for either Navigation ($p < 0.001$) or Exploration ($p < 0.01$). Dyslexics did not show any significant differences in how they used their strategies.

4.1 Frequently Occurring Action Sequence Forms

The classification of action sequences related to strategies produced a total of 89 different action sequence compositions used by participants. Table 1 presents the 12 most common compositions applied by blind users.

Blind participants spent a substantial amount of time traversing different types of content in the order it occurred on a webpage. Links and headings were commonly traversed (sequence B1 and B2 in Table 1); however, participants seldom expressed that they were trying to get an overview of the page using these sequences. Only 15 of the 94 strategy instances in these two most frequently used action sequence compositions were classified as being used in support of the Discovery strategy. Instead, participants were usually exploring content in a local area on the webpage. In the case of headings, participants checked headings to see if they related to a specific piece of content they were looking for on the page. In the case of links, participants were often looking for a cue as to what the next action would be to complete their tasks.

Table 1: Most frequent action sequence compositions for blind participants.

Action Sequence Description	Instances N (%)
B1) Traversal of links on a page	51 (13.3)
B2) Traversal of headings on a page	43 (11.2)
B3) Keyword search within page or site	26 (6.8)
B4) Sequential traversal through page content	24 (6.3)
B5) Task accelerator in assistive technology or user agent	19 (5.0)
B6) Traversal of form controls within one form to understand form	14 (3.7)
B7) Fishing for controls with the cursor in a localized area of the page	11 (2.9)
B8) Search for form controls with assistive technology	11 (2.9)
B9) Probing controls to identify what they do	10 (2.6)
B10) Jump to known page control	9 (2.3)

There were two sequence compositions (B3 and B10) where blind participants would try to jump to a specific piece of content. When blind participants felt that there should be a form to aid them in their task on a webpage (e.g. login, money transfer, status update), they would use their screen reader to jump to the first form. Similarly, when specific content was expected on a page (e.g. a television programme), participants would do a keyword search within the page to try to jump to that content.

Blind participants also applied an action sequence composition for Exploration that we refer to as *fishing for controls*. In this sequence, participants would suspect that a control for a form was in the area they had been exploring but they could not locate it. Participants would activate the screen reader cursor to try to trigger any control in the area (B7). If they were unsuccessful, they would move the cursor a small amount, and then try again. This behaviour would continue until they activated a control or gave up.

The most frequent action sequence compositions for partially sighted users are presented in Table 2. The majority of the action sequences were used for Discovery.

Table 2: Most frequent action sequence compositions for partially sighted participants.

Strategy	Instances N (%)
P1) Checking tooltips for further information about link purpose	9 (9.2)
P2) Scrolling to the bottom of the page looking for content	8 (8.2)
P3) Moving viewport vertically on visible edge to understand hierarchical structure of content	8 (8.2)
P4) Checking for expected content in locations consistent with web conventions	7 (7.1)
P5) Moving viewport vertically while highlighting content to understand	6 (6.1)

relationships	
P6) Moving viewport horizontally to locate content or read information	6 (6.1)
P7) Zooming out for overview and then zoom in for detail	6 (6.1)
P8) Moving viewport horizontally to scan links in top navigation bar	5 (5.1)
P9) Moving viewport vertically to scan links in left hand navigation bar	4 (4.1)
P10) Moving viewport to explore upper limit of a region of content, and then moving viewport to explore lower limit	4 (4.1)

Partially sighted participants would often align their screen magnifier viewport to an edge of a content area (e.g. top navigation bar) and then move along that edge, either vertically or horizontally. Sometimes, this would be very easy, such as situations where they investigated the left navigation bar where the edge of the webpage would keep the magnifier aligned (sequence P9 in Table 2). In other cases, participants manually tried to follow an edge of an arbitrary region of the web page to investigate content on the page (P3 and P6). We observed that at the beginning of tasks, partially sighted participants quickly moved their screen magnifier down the page all the way to the bottom in order to get an overview of the whole page structure (P2).

Interestingly, the action sequence composition with the largest proportion for partially sighted participants was the use of tool-tips for Exploration. They used this form extensively in order to understand the purpose of a link or control, sometimes in the absence of any other contextual information in the viewport of the screen magnifier.

Partially sighted participants would often check areas of websites where they find content on other websites (P4). For example, participants would check the upper right hand corner of a site for a search form, or look for navigation links on the left hand side of the page.

The most frequently occurring action sequence forms elicited from dyslexic participants involved navigation around the website and are presented in Table 3.

Table 3: Most frequent action sequence compositions for dyslexic participants.

Strategy	**Instances N (%)**
D1) Carefully selecting links based on link label and link probing	25 (23.8)
D2) Returning to the website homepage and beginning task again	14 (13.3)
D3) Checking web convention locations for expected content	12 (11.4)
D4) Keyword searching within page, or within site	7 (6.7)
D5) Random probing of links	6 (5.7)
D6) Looking for feedback that actions have had desired effect or information on what errors have occurred	6 (5.7)
D7) Looking for familiar icons	4 (3.8)
D8) Scrolling to bottom of the page looking for content	4 (3.8)

D9) Checking the top of the heading area of the page above the navigation bar, and then the footer of the page	4 (3.8)
D10) Scanning links in left navigation bar	3 (2.9)

Of particular interest in the results for dyslexic participants was that they employed a variety of action sequences in support of Navigation. Dyslexic participants were very selective about what links they would follow for completing their task (sequence D1 in Table 3). Often this was through careful reading of a set of link labels and occasional, but purposeful, probing of links to investigate where they led. In other cases, when dyslexic participants became unsure of how to proceed with their task, they would select links at random to probe in the hopes of finding something useful (D5). Dyslexic participants also readily returned to the home page of a website when they became lost during a task (D2).

A number of action sequences used by dyslexic participants related to familiar or expected webpage content (D3). For example, dyslexic participants looked for content in locations that adhere to web conventions (e.g. a link to the user's account being in the top-right hand corner of the page), which was quite similar to partially sighted participants. Where dyslexic participants had an idea of what they were looking for, they would often use keyword search (D4), either within a webpage using the browser search controls or across webpages using the website search facilities.

5 Discussion

The results indicate that blind, partially sighted and dyslexic Web users all apply a wide variety of different action sequences that map onto seven different types of strategy to support their interaction on the Web. The strategies that were most frequently observed in the study were those that supported finding content within a website (Navigation), obtaining an overview of the contents of a webpage (Discovery) and understanding the meaning of individual pieces of content (Exploration). Interestingly, there were cases where action sequences, such as heading traversal, could be used for more than one strategy type: Exploration or Discovery, depending on the situation.

Many of the top strategies applied by blind and partially sighted participants match those that were elicited in the Theofanos and Redish studies [14, 15] nearly a decade ago. In some ways, this is positive as it gives us confidence in the methodology used in this study. On the other hand, it is very odd that after a decade of experience of Web development, during which both assistive technologies and the Web itself have evolved and supposedly improved, these two user groups are still largely doing the same things they were when using much less sophisticated technology. This could be interpreted as users sticking with strategies that they know work. Alternatively, it could be argued that interaction design for blind and partially sighted people has stagnated because we do not fully understand how they interact with the web and how they understand webpage content and structure.

The emphasis by blind participants on understanding individual pieces of content through Exploration strategies suggests that many blind Web users may work from a bottom-up approach in building their mental models of websites. The information gathered through use of Exploration strategies is then supplemented by Discovery strategies to understand an individual webpage in a more top-down manner. The intensive use of links and headings has important implications for the design of both websites and assistive technologies.

Current screen readers provide functionality to access a list of headings on the webpage. This functionality presents users with the headings and their level (e.g. h1, h2) in the order they occur on the webpage. Users who are building a mental model bottom-up will work from the individual headings and then integrate them into an overall model of how the webpage is structured. The current screen reader functionality supports, to some extent, both Exploration and Discovery strategies. However, if the structure of the headings in a website is poor, there will be a mismatch between a user's mental model that is constructed from the individual headings and the intentions of the designer. This shows the importance of emphasizing to web developers the use of good heading structures. It also lends support to helping users through the application webpage transcoding to add or correct headings on webpages [8, 15].

However, in our results traversing links was as common as traversing headings. Yet, in current screen readers, the support for forming an understanding of how the links relate to the rest of a page is almost non-existent. When users interact with links in screen readers it is usually through functionality that presents all of the links on the page sequentially from the beginning to the end. All structure and context for the links is lost. This functionality supports Exploration, but not Discovery. This has implications for future research and the development of screen readers. Specifically, how can we support Discovery strategies for blind users and help them build an appropriate mental model of a webpage? It is possible that new assistive technology features could be designed to do this, but there are also likely opportunities for changing the design of web content. The web is becoming more flexible in how it is implemented, especially with the advent of responsive design, and taking advantage of that to better support the strategies found in this study needs to be investigated urgently.

When looking at the action sequences of partially sighted users, the use of Discovery strategies by screen magnifier users is certainly understandable. Enlarging text and other content comes at the cost of how much of the screen a user can see, and therefore they must reconstruct the overall page. However, the action sequences used by partially sighted users provide insights into future designs. Going beyond previously reported results, we observed users working with edges and borders of content to constrain where they moved their magnifier viewport. Moving vertically on the edge of a screen was relatively easy, while moving horizontally proved to be more difficult. In order to improve the effectiveness of those action sequences, we may be able to transform, or enhance, websites to better support following those edges. One solution may be to employ current responsive design approaches to websites, where the horizontal dimension of a website can be shrunk, moving the content into a single long narrow column. This might result in a very long webpage, which would need to be augmented further with links that allow partially sighted users to move directly to sections of content that are of interest, or back to the top of the page.

Dyslexic participants employ a range of strategies in reasonably equal proportions. One aspect of the most frequently applied strategies by dyslexic participants is that many relate to moving from one page to another. Very selective Navigation strategies, returning to home pages when lost and the use of keyword searches to find content in a site are all indicators the importance of good navigational support on websites to these users.

One interesting aspect of all user groups is how many of their strategies relied on having designs that are externally consistent between websites. Both blind and dyslexic participants used keyword searches to find words that regularly appear in websites (e.g. login, search), or words commonly used in the website domain (e.g. a bank sort code). Further, partially sighted and dyslexic participants used strategies that relied on controls or content being in “typical places” on websites. The data supports the idea that good general design practices related to consistency for usable websites will benefit all users, including users with print disabilities.

6 Conclusions

In this paper we have presented a study of the interactions that people with print disabilities, specifically blind, partially sighted and dyslexic users, have with websites. This study has produced a number of important findings that add to the very thin body of empirical literature that exists about users with disabilities interacting with the Web.

The key result that comes from this study is a set of 7 key strategy types that are applied by print disabled users when interacting with web applications. These strategies provide a framework, in which we can analyse the interactions of users, above the level of individual operations in their user agent, or assistive technology, which was the focus of most previous research. Using these strategies, we have identified a number of interesting qualities about the differences in how users from different user groups approach websites.

In general, blind users appear to exhibit far more instances of applying strategies than their partially sighted or dyslexic counterparts. They have a heavy reliance on action sequences supporting the Exploration strategy. This type of interaction implies that blind users are not actively seeking out information about webpage structure, but instead letting the structure emerge.

In comparison, partially sighted users applied action sequences that were for active Discovery of webpage structure. Further, while dyslexic users have no one strategy that they prefer over others, they have a wide range of strategies that they apply.

There are very clear, distinguishing features for each of these user groups regarding how they interact with the Web. However, they all share a large number of web accessibility problems [1]. The implication of our findings on strategies is that: it is very unlikely that there is one solution that can be prescribed to solve any given problem that is encountered by all user groups. As a result, the current approach of avoiding accessibility problems in web design is not going to be sustainable in the future. As we begin to understand more deeply the interactions of users, we need to define design principles and user-validated design patterns that support the strategies of different user groups. Once we have such a set of design principles, we can begin to personalize web applications for people with disabilities in order to provide them with a truly equal experience.

However, before we can do that, we need a great deal more empirical research about disabled users and their interactions on the Web. Our future research will examine how different action sequences and their strategies are combined within the interface during user interactions. It is hoped that these strategy combinations will reveal important insights into the contributory causes of the problems users encounter on the Web, and point the way towards solutions for different user groups.

7 Acknowledgments

The research was conducted as part of the Inclusive Future - Internet Web Services project. We are grateful to European Commission (FP7; Grant: 257623) for part funding the project and all the project partners for their hard work and support. We thank all the participants in the study for their time and effort.

8 References

1. The Web: Access and Inclusion for Disabled People. Disability Rights Commission, 2004.
2. Berry, J. Apart or apart?: access to the Internet by visually impaired and blind people, with particular emphasis on assistive enabling technology and user perceptions. Information Technology and Disabilities, 6, 3, 1-16 (1999)
3. Bigham, J. P., Cavender, A. C., Brudvik, J. T., Wobbrock, J. O. and Lander, R. E. WebinSitu: a comparative analysis of blind and sighted browsing behavior. In: 9th International ACM SIGACCESS Conference on Computers and Accessibility, pp. 51-58. ACM, New York (2007).
4. Borodin, Y., Bigham, J. P., Dausch, G. and Ramakrishnan, I. V. More than meets the eye: a survey of screen-reader browsing strategies. In: 2010 International Cross Disciplinary Conference on Web Accessibility (W4A), pp. 1-10. ACM, New York (2010)
5. Coyne, K. P. and Nielsen, J. Beyond ALT Text: Making the Web Easy to Use for Users with Disabilities - Design Guidelines for Websites and Intranets Based on Usability Studies with People Using Assistive Technology. Nielsen Norman Group (2001).
6. Engelen, J. J., & Bauwens, B. Large Scale Text Distribution Services for the Print Disabled: The Harmonisation and Standardisation Efforts of the TIDE-CAPS Consortium. In: Rehabilitation Technology: Strategies for the European Union: Proceedings of the 1st Tide Congress, 22 pages. IOS Press, (1993).

7. Lazar, J., Allen, A., Kleinman, J. and Malarkey, C. What frustrates screen reader users on the web: a study of 100 blind users. International Journal of Human-Computer Interaction, 22, 3, pp. 247-269 (2007)
8. Lunn, D., Harper, S. and Bechhofer, S. Identifying Behavioral Strategies of Visually Impaired Users to Improve Access to Web Content. ACM Trans. Access. Comput., 3, 4, pp. 1-35 (2011)
9. O'Reilly, T. What is Web 2.0: Design patterns and business models for the next generation of software. Communications & strategies, 1, 17 pages (2007)
10. Paciello, M. Making the World Wide Web Accessible for the Blind and Visually Impaired. Florida Libraries, 39, 4 (1996)
11. Paciello, M. Making the Web Accessible for the Deaf, Hearing, and Mobility Impaired. Florida Libraries, 39, 5 pp. 83-91. (1996)
12. Petrie, H. and Kheir, O. The relationship between accessibility and usability of websites. In: 2007 SIGCHI Conference on Human Factors in Computing Systems, pp. 397-406. ACM, New York, (2007)
13. Petrie, H., Morley, S., McNally, P. and Graziani, P. Authoring hypermedia systems for blind people. In*:* IEEE Colloquium on Authoring and Application of Hypermedia-Based User-Interfaces, pp. 6/1-6/3. (1995)
14. Power, C., Freire, A., Petrie, H. and Swallow, D. Guidelines are only half of the story: accessibility problems encountered by blind users on the web. In: 2012 SIGCHI Conference on Human Factors in Computing Systems, pp. 433-442. ACM, New York, (2012)
15. Takagi, H., Saito, S., Fukuda, K. and Asakawa, C. Analysis of navigability of Web applications for improving blind usability. ACM Trans. Comput.-Hum. Interact., 14, 3, pp. 13/1-13/36. ACM, New York, (2007)
16. Theofanos, M. F. and Redish, J. Bridging the gap: between accessibility and usability. Interactions, 10, 6 pp. 36-51. (2003)
17. Theofanos, M. F. and Redish, J. Helping Low-Vision and Other Users with Websites That Meet Their Needs: Is One Site For All Feasible? Technical Communication, 52, 1, pp. 9-20. (2005)
18. Trewin, S., Richards, J, Bellamy, R., John, B.E, Thomas, J., Swart, C. and Brezin, J. Toward Modeling Auditory Information Seeking Strategies on the Web. In: 2010 Extended Abstracts on Human Factors in Computing Systems, pp. 3973-3978. ACM, New York, (2010)
19. Watanabe, T. Experimental evaluation of usability and accessibility of heading elements. In: 2007 International Cross Disciplinary Conference on Web Accessibility (W4A), pp. 157-164. ACM, New York, (2007)