Middleware to Expand Context and Preview in Hypertext

Simon Harper, Carole Goble, Robert Stevens, and Yeliz Yesilada
Information Management Group
Dept of Computer Science, University of Manchester
Manchester UK
http://augmented.man.ac.uk
(sharper|carole|stevensr|yesilady)@cs.man.ac.uk

ABSTRACT

Movement, or mobility, is key to the accessibility, design, and usability of many hypermedia resources (websites); and key to good mobility is context and preview by probing. This is especially the case for visually impaired users when a hypertext anchor is inaccurately described or is described out of context. This means confusion and disorientation. Mobility is similarly reduced when the link target of the anchor has no relationship to the expected information present on the hypertext node (web-page). We suggest that confident movement with purpose, ease, and accuracy can only be achieved when complete contextual information and an accurate description of the proposed destination (preview) are available. Our past work (1) deriving mobility heuristics from mobility models, (2) transforming web-pages based on these heuristics, and (3) building tools to analyse and access these transformed pages; has shown us that a tool to expand context and preview would be useful. In this paper we describe the development of such a middleware tool to automatically and dynamically annotate web-pages with additional context information present within the page, and preview information present within hypertext link destinations found on the page.

Categories and Subject Descriptors

H.1.2 [**Models and Principles**]: User/Machine Systems—*Human factors / Human information processing*;
K.4.2 [**Computers and Society**]: Social Issues—*Assistive technologies for persons with disabilities*;
J.m [**Computer Applications**]: Miscellaneous;
I.7.2 [**Computing Methodologies**]: Document Preparation—*Hypertext / hypermedia*

General Terms

Human Factors, Algorithms, Experimentation

Keywords

Web, travel, mobility, visually impaired, practical implementation

ASSETS'04, October 18–20, 2004, Atltana, Georgia, USA.

1. INTRODUCTION

Movement through and around complex hypermedia environments, of which the web is the most obvious example, has long been considered an important and major issue in the Web[1] design and usability field [8, 13]. The commonly used slang phrase 'surfing the web' implies rapid and free movement, pointing to its importance among designers and users alike. It has also been long established [4, 9] that this potentially complex and difficult movement is further complicated, and becomes neither rapid or free, if the user is visually impaired[2]. Previously we have shown that virtual mobility can be likened to real world mobility, that solutions to real world problems can be applied to the web, and that models of mobility can inform our creation of mobility heuristics. We have established that the richness of visual cues presented to a sighted user are not appropriate or accessible to a visually impaired user [15]. Further, both our work and others have shown that lack of *context*[3] and *preview*[4] are a hindrance to visually impaired users and users of small devices where screen real-estate is limited [12, 5](like PDAs[5]).

1.1 Motivational Example/Problem

To fully realise the problems involved we suggest that our readers start their browser and limit the window size to the top left fifth of the screen (see Figure 1). Now browse a series of simple and complex Web sites by jumping from hyperlink to hyperlink (Using the `tab` key or mouse if you prefer) and disregard everything that is not described in that hyperlink. In this way a better understanding of worth of the context of the surrounding information and preview of the upcoming information can be appreciated. In our example we have used our own university's news site. Note the problems you have; we believe you will find that:

1. You cannot get an overall feel for what's on the page.
2. You do not know where you are on the page or if you've been there before, consequently, you become disoriented.
3. You cannot tell where you are going to move to once the hyperlink is selected.

[1] Or hypermedia.
[2] Here used as a general term encompassing the WHO definition of both profoundly blind and partially sighted individuals [27].
[3] The part of something written which precede or follow a hyperlink – in this case – which are so intimately associated with it as to throw light upon its meaning.
[4] A view of upcoming information that cannot be readily detected also known as preview by probing.
[5] Personal Digital Assistants.

Figure 1: Where will I go? (http://news.man.ac.uk/)

4. You cannot tell if the information on the target page is the information you are expecting or require.

5. You cannot tell if the link is referential, associative, or structural [19].

6. You often find you tab from link to link all with the same name and in effect become lost in the repeated 'Read More...', 'Read More...' anchors.

7. The whole movement and travel experience is neither satisfactory nor enjoyable.

We wish to tackle this problem by allowing readers to *preview* link targets during their journey around the Web. And to do this effectively *context* must feature in their decision process.

1.2 Synopsis

We describe the development of our system, a middleware tool to automatically and dynamically annotate web-pages with additional context information present within the page and preview information present within hypertext link destinations found on the page. We have created a 'Perl'[6] based Apache[7] content handler which transforms simple links on a web-page to complex links that contain information about the context of the link and a preview of the destination page. The rest of the paper can be summarised as follows:

What is Browsing We investigate the browsing behavior of Web surfers to support our assertion that context and preview are important. And find from related literature that browsing behavior suggests that both context and preview are not only involved in the activity, but are also important for the successful completion of the activity.

Context and Preview Next, we investigate the issues surrounding context and preview in hypermedia and we establish why these twin concepts are important for better universal access to Web-pages. We explain how context influences the interpretation of the link destination by the reader; and assert that users need some form of detection and avoidance schema based on accurate and appropriate previews.

[6] `http://www.cpan.org`
[7] `http://www.apache.org`

Transcoding Transcoding is a technology used to adapt Web content so that it can be viewed on any of the increasingly diverse devices on the market. We briefly review the state of current research in transcoding as it has been used for a number of years in the context of making incomplete or badly written hypertext accessible to visually impaired users and their accessibility technologies. We can use aspects of this technology to transform web pages on-the-fly to include more context and preview information.

Discussion We describe the concepts, rational, and techniques behind our experiment, a middleware system to automatically and dynamically annotate Web-pages with additional context information present within the page and preview information present within hypertext link destinations found on the page. And show how we made anchor descriptions accurately represent the anchor destination within the context of the current information search.

Conclusion Finally, we focus on our conclusions from the work undertaken and look at future work including system evaluations.

2. WHAT IS BROWSING

It is well known that users do not read Web pages, they scan them [23] and links are important elements of Web pages that facilitate scanning and browsing. In consequence, we investigate browsing and scanning behavior to support our assertion that context and preview are important dimensions of these activities. The literature on scanning and browsing activities suggest that both context and preview are not only involved in the activities, but are also necessary for its successful completion. If Web surfers did not use context and preview while browsing hypertext then it would be difficult to suggest that enhanced context and preview would be useful in link descriptions.

Browsing is an activity that is difficult to define [6], but there is general agreement that "we all browse in various contexts to make sense of the world around us" [7]. Some researchers also describe it as a process of "picking out bits and pieces$\cdots$selecting worthwhile information need or interest" [11]. Different disciplines look at browsing from different perspectives [7]. Various reviews suggest that browsing is a kind of searching, in which initial search criteria or goals are only partly defined or known in advance. Browsing involves scanning, which has been previously described as looking, examining, or sampling, during which the person's body or eyes move smoothly at will. Browsing also involves distinct consumer shopping behavior that is related but not equated with buying behavior. Methodologically, eye movement can be a useful indicator of browsing and has been used to test the effect of different page layout or catalogues on browsers' attention. Browsing is fundamentally scanning and has been related to environmental perception and cognition. For example, sightseeing is environmental browsing as perceptual experience [7].

Although all of these views of browsing have various approaches and provide different definitions, there seems to be an agreement on the essential characteristic of browsing which is movement. Browsing can be thought of as travel in information space, and in fact many users refer to real world metaphors to describe browsing [21]. Indeed research suggests that browsing is made up of *Searching*, the task of looking for a known target. *Inquiry*, the task of looking to see what is available in the world. *Querying*, using a search engine to submit a description of the object being sought and receiving relevant content of information. And *Navigation*, moving

oneself sequentially around an environment, deciding at which step where to go.

In investigating the nature of browsing, several researchers have attempted to establish different types of browsing [17]. Fundamentally, these types have been established by considering the goal, purpose or the information need. Search strategies have been defined as "a set of ordered tactics or behavioral moves that are consciously selected, applied and monitored to solve the problem" [22]. Analytical search strategies are formal, discrete and deterministic, in contrast, browsing search strategies are informal, continuous and heuristic. Indeed the four distinguishable browsing strategies of scanning, observing, monitoring, and navigating have been identified.

Finally, five dimensions that can be used to distinguish browsing from other information seeking behaviors and to characterize types of browsing [7] have been suggested. *Context:* Organisation (how resources are organised and presented), Interface (the display perceived by the user), Feedback (relevance or content related, and orientation feedback), Economics (access costs, resources available such as time). *Behavioral:* Scanning (orientation or exploratory scanning) and movement [25]. *Motivation:* Purpose (why people engage in browsing) and goal (what they intend to accomplish). *Cognitive:* Plan (accomplishing a goal can be planned or unplanned) and knowledge/experience (content or structure knowledge) [25]. *Resource:* Form (object or representation) and focus (content or path) of resource. These five views are important to our research because motivation influences context and link description supports context. Influences link description composition supports behavior, and good information preview supports cognition and resource destination.

In summary, browsing is movement in the information space and the user is in control of what to read or examine. Many studies have addressed different browsing types and strategies [21] however movement is the essential characteristic of browsing [3]. We all browse in various contexts picking out bits and pieces of information and selecting worthwhile information [20]; and we accomplish this by using a searching and scanning behavior over organisations of the material [26, 23], interfaces to that material [2], and feedback about the material [22].
Because scanning is extensively used [18, 17, 23]; because context is required for accurate decision making [25]; and because preview is needed to decide on changes of focus [25]; including context and preview in hyperlinks *should* assist browsing behavior.

3. CONTEXT AND PREVIEW

3.1 Context

In general, context means the part or parts of something written or printed which precede or follow a text or quoted sentence, and are so intimately associated with it as to throw light upon its meaning.
Information contextualisation is important for mobility within the docuverse[8] because users presume they will be staying within the context of the current information space. This means that movement becomes difficult if context changes when the user does not want to. Obviously there may be times when a user wishes to change context but without more information than a hyperlinked keyword the decision to move along that link is more difficult. For instance creating an anchor point described as 'Georgia' will enable linking to an associated resource, however if the context of the discussion is not known then the user cannot tell if they will move to another resource describing Georgia the former eastern bloc country, Georgia the US state, or Georgia our dog.

Put simply, in terms of our research, context influences the interpretation of the link destination by the reader.

3.2 Preview and Probing

In Web-mobility, the lack of previews of upcoming hyperlink destinations and information relating to movement across those destinations suggests that some degree of 'probing' must be implemented so that a limited preview can be obtained [14]. Indeed a user observed traversing the Web, can be seen to select a hyperlink, preview the contents (by clicking or placing the caret over the link to see the destination) and return if the contents are not applicable. This probing is continued until each hyperlink is previewed, and interesting contents are found which suggests that to avoid irrelevant information encountered 'on-the-fly' a user needs some form of detection and avoidance schema based on accurate and appropriate previews.

3.3 Transcoding

Simply, transcoding is a technology used to adapt Web content so that they can be viewed on any of the increasingly diverse devices on the market. It has been used for a number of years in the context of making incomplete or badly written hypertext accessible to visually impaired users and their accessibility technologies. Transcoding in this context normally involves: (1) Syntactic changes like shrinking or removing images [16], (2) Semantic rearrangements and fragmentation of pages based on the mean of a section [24, 10], (3) Annotation of the page created by a reader [5], (4) and Generated annotations created by the content management system [5].

There are a number of different ways that transcoding can take place. In one example, the original material (an HTML document, for example) is analysed by a program that creates a separate version containing annotations. The annotations include information that will instruct the reformatting process, and inaccessible elements will be removed or altered. Systems are often based along similar lines and address set problems, some are annotation based [16], others generate text only versions [24, 10], some filter the content [1], and others are specifically used for small scale device interaction [5]. Whatever system is used it invariably does not transform all inaccessible elements but just a subset leaving holes in the accessibility of their transcode.

Each of these types of transformations are fraught with problems with regard to the acceptability of the resulting generation. This is especially the case when sighted users as well as visually impaired users wish to use the same page. Automatic transcoding based on removing parts of the page results in too much information loss and manual transcoding is near impossible when applied to dynamic websites. Most systems use their own bespoke proxy-servers or client side interfaces and these systems require a greater setup cost in-terms of user time. Finally, some systems require bespoke automatic annotation by a content generator and so are not useable by every user and all systems.

[8]A named set of documents, in this case, joined by hyperlinks which provide a pathway through the discourse. For text, a docuverse is the analog of a database.

Often transcoding systems lean towards solving the problems of one user group (the problem that got us here in the first place) and so destroy the content / structure / context for other non-target groups. This directly challenges the nature of the World Wide Web and of other Open Hypermedia Systems in general.

4. DISCUSSION

The question we faced was:

> 'How can we make anchor descriptions accurately represent the anchor destination within the context of the current information search, therefore making the user interaction easier? At a minimum we must also preserve the experience of all users, and if possible enhance their experience as well.'

We analysed a number of implementation paradigms and came up with three options to implement our system.

Client Side Extension By creating an extension to a conventional browser – like Netscape or Internet Explorer – we could enable users of these systems to generate context and preview information themselves. While this is initially an attractive option there are some difficulties. Firstly, many visually impaired surfers use systems like emacspeak, lynx, and other non conventional browsers for their web interaction and so extending each would be impractical. Secondly, our system requires that remote resources are scanned to decide on their content. This occurs once for each hyperlink on the page and so the time from initial page request to a fully generated page output would be prohibitive on a slow conventional connection.

JavaScript Extension We could generate JavaScript code and descriptions on the server of each link and it's context. Then, insert the results of this generation into the page so that the page appears unmodified however additional information would be displayed when anchors were 'rolled-over'. The major problem with this system is that many browsers used by visually impaired surfers do not have a JavaScript ability, and those that do often have out-of-date and incomplete versions.

Document Extension Extending the document on-the-fly may be the best solution. By re-engineering the document a user can have access to information in a form that is viewable on all browsers, is generated quickly as it is part of the server engine, and is on a fast connection. Speed and time advantages could also be gained by periodically processing all files on the site using this method as opposed to generating them on-the-fly. An update function could also be provided on each page to generate a new view on-the-fly if the user required. Problems do exist with this solution as the original document is altered and the look may deviate from the designers original concept due to the document modifications.

After further discussions and experiments we decided that the Document Extension option was the most appropriate for the majority of users.

We decided to create our system as an external application / tool to the Web Server but linked to it by a Content Handler (see Fig. 2). In this way it would run on most Web Servers – although we used Apache. This solution also has the benefit of making the tool code-base updateable on-the-fly without affecting the server itself. And, it allows designers and web site administrators to specify which directories or file types (see Fig. 3 – here using the extension `.cphtml`) they would like to be processed in this way so that the 'look and feel' of design critical material can be maintained. We also decided to code the tool in 'Perl' so that it would run on most operating systems and could be translated to 'c/c++' if required. Perl also has the advantage of being a very good glue and file parsing language, both useful traits in our solution.

```
<Location /context_preview>
    SetHandler perl-script
    PerlHandler Apache::add_context_preview
</Location>
```

Figure 2: Content Handler on a Directory

```
AddType text/html .cphtml
<Files ~ "\.cphtml$">
    SetHandler perl-script
    PerlHandler Apache::add_context_preview
</Files>
```

Figure 3: Content Handler on a File Type

4.1 Expanding Context and Preview

Our task was to add context to links without context and to place preview information into links which accurately describe the destination of the link; and to make these two 'concepts' universally accessible – because at the moment they are not. We are not concerned with the type of link in this system (Structural, Associative, or Referential) however process time and server load could be reduced if a differentiation process could be formulated as the preview of a referential link could in effect be derived from the context of the referring page.

4.1.1 How do we Expand Context?

Visually impaired users often have problems deciding on a link's context when that link is not descriptive. Many current websites only link one or two words when trying to associate pieces of information on their page with associated information on a destination page. Problems also exist when content management systems generate summary pages and clusters of summary links – to referential information – but only provide phrases like 'More' or 'Click Here' to link these pages. Some content management systems just use the title of the referenced material as the link anchor however this does not work for associative links.
Expanding context is a non-trivial task and our solution provides only a partial answer. Our system uses the text surrounding the link to give the link itself a better context (see Fig. 4) when the original text ('Read Me...') is expanded based on the content of the summary above it.

Our basic algorithm is simple but effective and moves through the page in a systematic way thus:

1. Search the document for anchors – `<a...> </a>`.

2. Make sure the hypertext link uses the hypertext transfer protocol – `<a href="http://....">` – we don't give context to mail and ftp protocols etc.

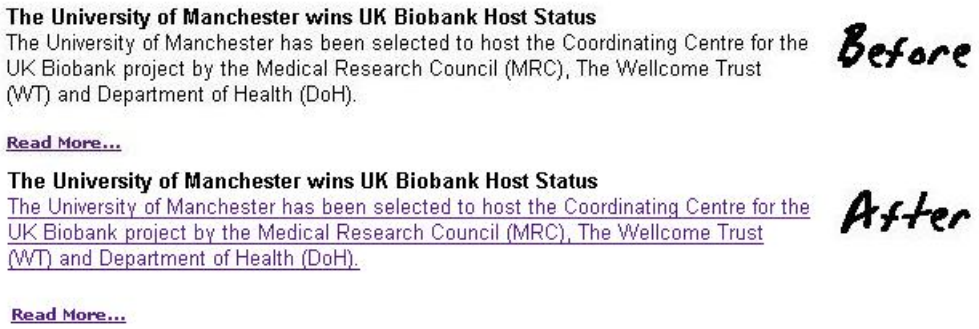

Figure 4: Complete Context – Example 1

3. We now move forwards and backwards expanding the anchor text until we meet either:
   - '.', '.' ',' ';', '(', ')', or ':'
   - `<start tag>` or `</end tag>`
   - However, if the tag is a line break (`<br>` or `<br />`) it is ignored.
   - All tags and whitespace are ignored that occur directly before the anchor tag.

However, our system doesn't always provide links that are fully descriptive. In fact our second example (see Fig. 5) processes an anchor from 'Read More...' into 'Instead, they cite conduct and emotional problems as more likely causes. Read More... '. We are intending to investigate this further by looking at ways to modify the anchor based on the sentence structure by deciding if it is a noun phrase built on a pronoun or if words within that phrase are personal pronouns (like 'they'). In which case we could iterate once more to try to find the subject of the pronoun. However, we are aware that the anaphoric reference is a notorious problem in natural language processing.

However, deciding on the amount of information to place in a link is often difficult and intractable. In fact anchor size is usually decided on in an adhoc fashion by the developer. In our case we don't have this problem as our anchor size is arbitrary but we do realise that the need to generate expanded context must be balanced with information overload – if the anchor size becomes too large.

4.1.2 How do we Expand Preview?

Expanding preview is our most compute intensive activity, because it requires each link destination to be fetched and processed before the complete document is returned to the user. Some systems (as we've already seen) attempt to cull only the title of the link destination in an attempt to avoid this processing overhead. However, from experiments performed with visually impaired users this method often does not produce effective link descriptions because titles often do not fully represent the content or context of the page. In our system we process each link target looking for a set of key description information.

1. We look for the text in the first paragraph `<p>...</p>`.

2. If this is not present we look for text inside the `<body>` element but outside other elements – in the document object model (DOM).

3. When Processing the previous elements we remove all tags within the paragraph and process to the first '.', '.' ',' ';', '(', or ')'.

4. We also look at `<meta ..>` tags within the document header and process these to provide a keyword set.

5. Finally, we look for our richest meta data named element 'DC.Description.Abstract' [9]. This element is the Dublin Core Description element and can be found in newer `XHTML` type pages. This element normally gives an abstract description of the page that has been reached.

The page would obviously become overly complex if we wrote the preview information as part of the anchor text along with the context text. Therefore, we add this preview information to the `title` attribute of the anchor so that accessibility tools can read it and the information is also displayed when the cursor is rolled over the link text. Our example (see Fig. 6) shows this happening with an associative link from the World Wide Web Consortium site. The page is processed, links (in this case for the 'IsaViz' tool) are expanded with information found in the first paragraph (in this case there is no Dublin Core data).

4.2 Analysis

We can see the results of a generation on the World Wide Web Consortium site. The example (see Fig. 7) shows how the page is rendered after being processed through the tool. The one and two word anchors (see Fig. 6) have been expanded as per our context processing algorithm and the title attributes of each anchor have been augmented with preview information from their destination page.

When tabbing through the document we can see the before and after text (listed below) spoken to visually impaired surfers when using assistive technologies to view the page.

Before Process ⟶ After Process (Context — Preview)

1. W3C RDF Validation Service ⟶ (The W3C RDF Validation Service — Note: this online service has been updated and now supports the Last Call Working Draft specifications issued by the RDF Core Working Group, including datatypes.)

2. Announcement ⟶ (has been updated to support all of the specifications described in the RDF Last Call Working Drafts

[9] `http://dublincore.org/`

Figure 5: Incomplete Context – Example 2

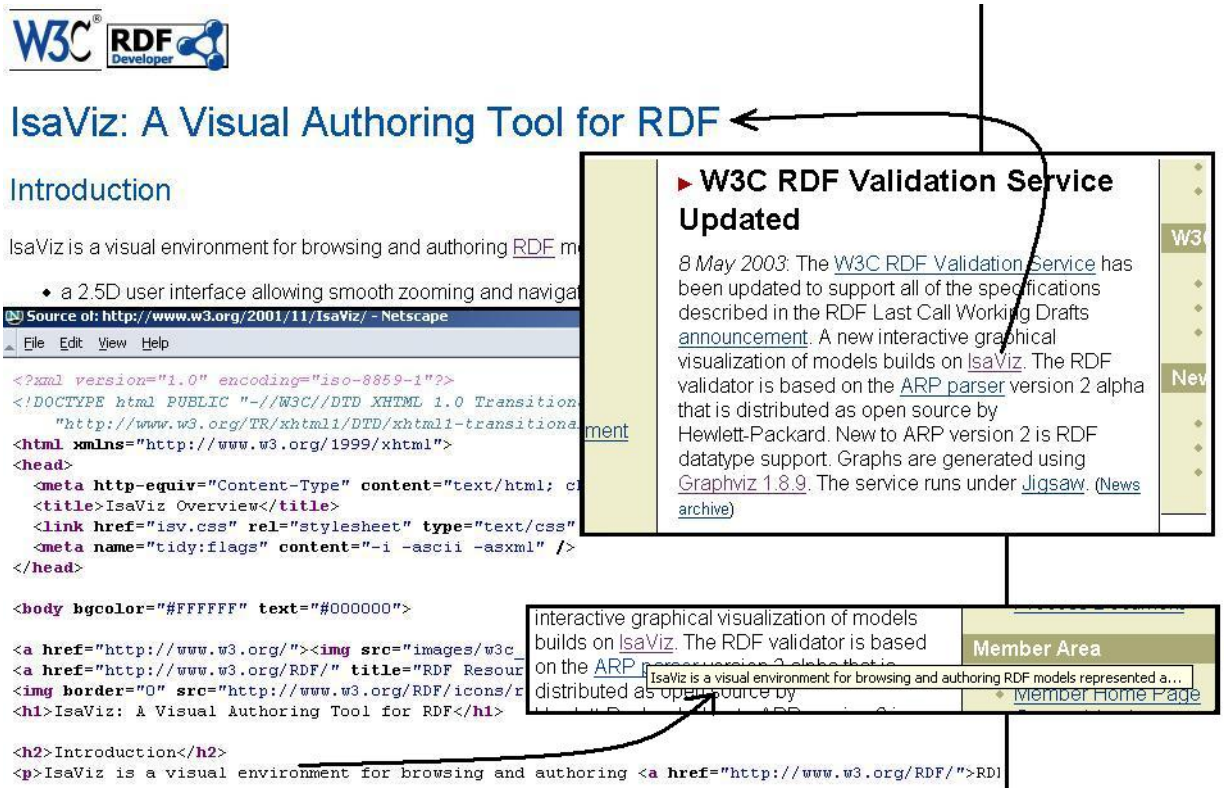

Figure 6: Preview – Example 3

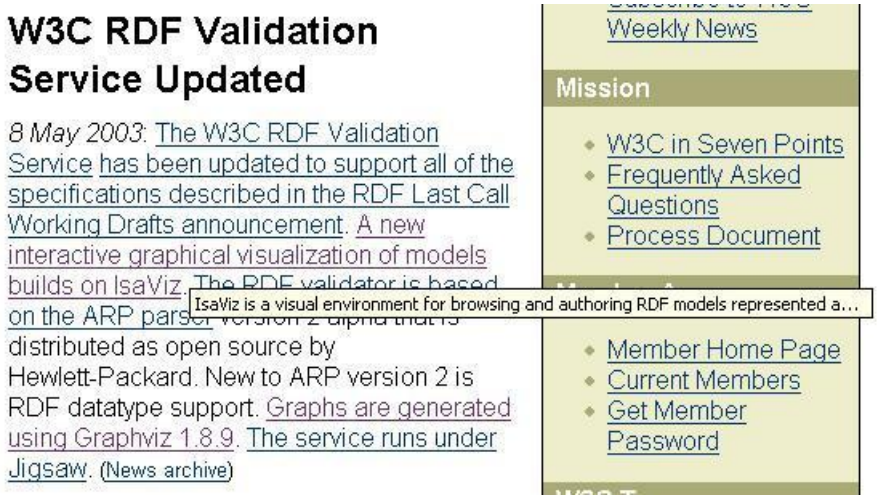

Figure 7: Combined Context and Preview – Example 4

announcement — Six RDF Last Call Working Drafts Published 24 January 2003: The RDF Core Working Group has released six Last Call Working Drafts. Comments are welcome through 21 February.)

3. IsaViz ⟶ (A new interactive graphical visualization of models builds on IsaViz — IsaViz is a visual environment for browsing and authoring RDF models represented as graphs.)
4. ARP parser ⟶ (The RDF validator is based on the ARP parser — For more information on ARP see the documentation included within the download package or see the main ARP web site http://www-uk.)
5. Graphviz ⟶ (Graphs are generated using Graphviz 1.8.9 — Open Source graph drawing software including web, Java, Linux and Windows support, and graph layout / stream filters.)
6. Jigsaw ⟶ (The service runs under Jigsaw. — Home page Mailing List Download Jigsaw)
7. News archive ⟶ (News archive — W3C RDF Validation Service Updated 8 May 2003: The W3C RDF Validation Service has been updated to support all of the specifications described in the RDF Last Call Working Drafts announcement.)

We can see that some 'anomalies' do exist within the generation. Item 6 concerning 'Jigsaw' lists the first 3 structural anchor texts "Home page Mailing List Download Jigsaw" which do not provide very much useful information regarding the information search. Also, the size of some of the preview texts are longer than we expected and this could lead to information overload in some users; and possibly in users with mild cognitive impairments (Agnosia, Dementia, etc). Finally, some users may have problems with the docuverse due to the additional clutter of preview components. The process was however successful because we managed to make anchor descriptions accurately represent the anchor destination within the context of the current information search. This may make the user interaction for visually impaired users easier than before. Additionally, we also managed to preserve the document for sighted users and attempted to enhance their experience by providing preview information as part of the title attribute of the anchor element.

5. CONCLUSIONS

We cannot claim that transcoding a web page in the way stated will increase a visually impaired person's cognition of context and preview. This is because we have not completed a full and statistically significant user evaluation. We have tested the transformed pages with the JAWS screen reader and find the additional information useful when 'tabbing' but such a small evaluation cannot be taken as significant. We have only shown that addressing the context and preview issues of visually impaired users *may be beneficial and is achievable*, we also feel that our work can be useful to sighted users as they too cannot accurately see the the destination of an anchor without good preview.

We **Conclude** that the use of context probing to assist preview and appropriate knowledge feedback may increase the mobility of users within many virtual journeys, and therefore solve many of the mobility problems encountered frequently in web based travel. However, we note that the tool adds another stage of complexity to standard web systems, but by automatically processing hypermedia resources and augmenting it for better web-mobility we consider this additional stage to be worthwhile. This additional stage can be justified but simpler alternatives may also be found.
Originally, we felt that extensions to the proposed system should be made. These would incorporate manual marking-up of pages with mobility information, browser mobility-application integration, and use of the Extensible Mark-up Language (XML) so that context and preview information could be accurately placed and described. However, we now think that due to the nature of the web this approach is neither flexible nor appropriate. Eventually, we feel that it will be possible to process hypertext information on-the-fly based on a user pre-defining a set of context and preview options that represent their preferences based on their cognition requirements, however this is not yet the case. We conclude that for a truly universal solution we need to mark our document alterations and allow users to control what is displayed at the client-side. Again this is not possible at present as it would require agreement on client-side functionality between developers, and this again is unlikely to occur.

We propose that **further work** needs to be undertaken to extend the tool such that the problems of incorrect preview, information overload, and clutter in the docuverse are overcome; we can then perform a complete user evaluation.
Finally, we feel this work has been useful in examining the redress required of writers – of pages, web designers, and implementors – of web focused systems, when their efforts do not provide sufficient context and preview information.

6. REFERENCES

[1] WebCleaner. http://webcleaner.sourceforge.net.

[2] R. Barrett, P. P. Maglio, and D. C. Kellem. How to personalize the web. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 75–82. ACM Press, 1997.

[3] M. Bates. The design of browsing and berrypicking techniques for the online search interface. *Online Review*, 13(5):407–424, 1989.

[4] M. Brambring. Mobility and orientation processes of the blind. In D. H. Warren and E. R. Strelow, editors, *Electronic Spatial Sensing for the Blind*, pages 493–508, USA, 1984. Dordrecht, Lancaster, Nijhoff.

[5] O. Buyukkokten, H. G. Molina, A. Paepcke, and T. Winograd. Power browser: Efficient web browsing for PDAs. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 430–437. ACM Press, 2000.

[6] E. Carmel, S. Crawford, and H. Chen. Browsing in hypertext: a cognitive study. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(5):865–883, 1992.

[7] S.-J. Chang and R. E. Rice. Browsing: a multidimensional framework. *Annual Review of Information Science and Technology*, 28:231–276, 1993.

[8] C. Chen. Structuring and visualising the www by generalised similarity analysis. In *Proceedings of the 8th ACM Conference on Hypertext and Hypermedia*, New York, USA, 1997. ACM Press.

[9] A. Chieko and C. Lewis. Home page reader: IBM's talking web browser. In *Closing the Gap Conference Proceedings*, 1998.

[10] Codix.net;. *Textualize*;. http://codix.net/solutions/products/textualise/index.html.

[11] J. F. Cove and B. C. Walsh. Online text retrieval via browsing. *Information Processing & Management*, 24(1):31–37, 1998.

[12] U. Cress. Previews in hypertexts: effects on navigation and knowledge acquisition. *J Comp Assist Learn*, 19(4):517–517, 2003.

[13] R. Furuta. Hypertext paths and the www: Experiences with walden's paths. In *Proceedings of the 8th ACM Conference on Hypertext and Hypermedia*, New York, USA, 1997. ACM Press.

[14] C. Goble, S. Harper, and R. Stevens. The travails of visually impaired web travellers. In *Proceedings of the Eleventh ACM on Hypertext and hypermedia*, pages 1–10, New York, USA, 2000. ACM Press. Towel Project (`http://towel.man.ac.uk`) - Winner of the Doug Engelbart Award for Best Paper at ACM Hypertext 2000.

[15] S. Harper, R. Stevens, and C. Goble. Towel: Real world mobility on the web. In J. Vanderdonckt and A. Puerta, editors, *Computer-Aided Design of User Interfaces II*, pages 305–314. Kluwer Academic, 1999. Towel Project (`http://towel.man.ac.uk`).

[16] M. Hori, G. Kondoh, K. Ono, S. ichi Hirose, and S. Singhal. Annotation-based web content transcoding. In *In Proceedings of 9th International World Wide Web Conference*, 2000.

[17] S. Jul and G. W. Furnas. Navigation in electronic worlds: a CHI 97 workshop. *ACM SIGCHI Bulletin*, 29(4):44–49, 1997.

[18] B. H. Kwasnik. A descriptive study of the functional components of browsing. *Engineering for Human Computer Interaction*, pages 191–203, 1992.

[19] D. Lowe and W. Hall. *Hypermedia and the Web: An Engineering Approach*. Wiley, 1999.

[20] P. Maglio and R. Barrett. Intermediaries personalize information streams. *Commun. ACM*, 43(8):96–101, 2000.

[21] P. Maglio and T. Matlock. Metaphors we surf the web by. In *Workshop on Personalized and Social Navigation in Information Space*, 1998.

[22] G. Marchionini. *Information Seeking in Electronic Environments*. Cambridge Series on Human-Computer Interaction. Cambridge University Press, 1995.

[23] J. Morkes and J. Nielsen. Concise, scannable, and objective: How to write for the web, 1997. `http://www.useit.com/papers/webwriting/writing.html`.

[24] W. Myers. *BETSIE:BBC Education Text to Speech Internet Enhancer*. British Broadcasting Corporation (BBC) Education. http://www.bbc.co.uk/education/betsie/.

[25] K. L. Norman. *The Psychology of Menu Selection: Designing Cognitive Control at the Human/Computer Interface*. Ablex Publishing Corporation, 1991. ISBN: 0-89391-553-X.

[26] H. Obendorf and H. Weinreich. Comparing link marker visualization techniques: changes in reading behavior. In *Proceedings of the twelfth international conference on World Wide Web*, pages 736–745. ACM Press, 2003.

[27] V. RNIB. A short guide to blindness. Booklet, Feb 1996. http://www.rnib.org.uk.